

# TRUTHFUL REPUTATION MECHANISMS FOR ONLINE SYSTEMS

THÈSE N° 3955 (2007)

PRÉSENTÉE LE 29 OCTOBRE 2007

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

Laboratoire d'intelligence artificielle

SECTION D'INFORMATIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

**Radu JURCA**

ingénieur en informatique diplômé de l'Université Polytechnique de Timisoara, Roumanie  
et de nationalité roumaine

acceptée sur proposition du jury:

Prof. T. Henzinger, président du jury

Prof. B. Faltings, directeur de thèse

Prof. K. Aberer, rapporteur

Prof. C. Dellarocas, rapporteur

Prof. T. Sandholm, rapporteur



ÉCOLE POLYTECHNIQUE  
FÉDÉRALE DE LAUSANNE

Lausanne, EPFL

2007



## Résumé

L'internet constitue aujourd'hui un milieu interactif où les communautés et les économies virtuelles gagnent de l'importance par rapport à leurs contreparties traditionnelles. Tandis que ce décalage crée des occasions et avantages qui ont déjà améliorés notre vie quotidienne, il apporte également une toute nouvelle série de problèmes. Par exemple, le manque d'interaction physique qui caractérise la majorité des transactions électroniques rend les systèmes beaucoup plus susceptibles à la fraude et à la tromperie.

Les mécanismes de réputation offrent une manière nouvelle et efficace d'assurer la confiance qui est essentielle au fonctionnement de chaque marché. Ils rassemblent les informations sur l'histoire (c.-à-d. les transactions antérieures) des agents qui participent dans le marché, et publient leur réputation. Les futurs associés guident leurs décisions en considérant l'information sur la réputation, et sont ainsi capable de faire les meilleurs choix. Les mécanismes de réputation en ligne connaissent un succès remarquable: ils sont présents dans la plupart des systèmes commerciaux déployés aujourd'hui, et sont sérieusement consultés par les utilisateurs humains.

La valeur économique de la réputation en ligne soulève des questions concernant la fiabilité des mécanismes eux-mêmes. Les systèmes actuels ont été conçus en supposant que les utilisateurs partageront honnêtement leurs avis. Cependant, des études récentes ont démontré qu'il existe des utilisateurs qui dénaturent la vérité pour manipuler la réputation.

La présente thèse décrit différentes manières de rendre les mécanismes de réputation en ligne plus dignes de confiance, en encourageant les participants à communiquer honnêtement les informations qu'ils détiennent. Différents types de mécanismes de réputation sont étudiés, et pour chacun, des mécanismes pour récompenser les agents qui rapportent la vérité sont présentés. Les problèmes liés à la complicité (c.-à-d. la coordination de la stratégie de plusieurs agents afin de manipuler le système) et à la robustesse sont également étudiés. De plus, cette thèse décrit une nouvelle application des mécanismes de réputation pour surveiller la qualité livrée par des fournisseurs de services, et étudie les facteurs qui motivent et influencent des utilisateurs humains qui postent leurs avis dans des forums existants.

**Mots clés:** mécanismes de réputation en ligne, feedback, incentive-compatibility, la collusion, mechanism design, la théorie des jeux



## Abstract

The internet is moving rapidly towards an interactive milieu where online communities and economies gain importance over their traditional counterparts. While this shift creates opportunities and benefits that have already improved our day-to-day life, it also brings a whole new set of problems. For example, the lack of physical interaction that characterizes most electronic transactions, leaves the systems much more susceptible to fraud and deception.

Reputation mechanisms offer a novel and effective way of ensuring the necessary level of trust which is essential to the functioning of any market. They collect information about the history (i.e., past transactions) of market participants and make public their *reputation*. Prospective partners guide their decisions by considering reputation information, and thus make more informative choices. Online reputation mechanisms enjoy huge success. They are present in most e-commerce sites available today, and are seriously taken into consideration by human users.

The economical value of online reputation raises questions regarding the trustworthiness of mechanisms themselves. Existing systems were conceived with the assumption that users will share feedback honestly. However, we have recently seen increasing evidence that some users strategically manipulate their reports.

This thesis describes ways of making online reputation mechanisms more trustworthy by providing incentives to rational agents for reporting honest feedback. Different kinds of reputation mechanisms are investigated, and for each, I present mechanisms for rewarding the agents that report truthfully. Problems related to collusion (i.e., several agents coordinate their strategies in order to manipulate reputation information) and robustness are also investigated. Moreover, this thesis describes a novel application of incentive compatible reputation mechanisms to the area of quality of service monitoring, and investigates factors that motivate and bias human users when reporting feedback in existing review forums.

**Keywords:** Online reputation mechanisms, feedback, incentive-compatibility, collusion, reporting behavior, mechanism design, game theory



## Acknowledgements

First and foremost, I want to thank my thesis advisor, Boi Faltings, for his guidance and support along these years. I have learned from Boi to look at the “broader” picture, to ask new questions and offer novel answers, and to pursue intuitive insights into mature results. Boi’s enthusiasm and vision set an impressive example, being a constant motivating factor along the PhD process. I also thank Boi for transparently handling the financial aspects of my PhD, which allowed me to worry exclusively about research.

I am especially thankful to the other members of my committee, Karl Aberer, Chris Dellarocas, Thomas Henzinger, and Tuomas Sandholm who took the time to assess my work, and gave me valuable feedback on this dissertation. I could not have imagined a better jury.

I am indebted to the past and present members of the Artificial Intelligence Lab (Marita Ailomaa, Arda Alp, Omar Belakhdar, Walter Binder, Monique Calisti, Jean-Cédric Chappelier, Cristian Ciressan, Ion Constantinescu, Emmanuel Eckard, Carlos Eisenberg, Thomas Léauté, Djamila Sam-Haroud, Vinh-Toan Luu, Santiago Macho Gonzalez, Mirek Melichar, Nicoleta Neagu, Quang Huy Nguyen, Brammert Ottens, Adrian Petcu, David Portabella, Martin Rajman, Marius-Calin Silaghi, Vincent Schickel-Zuber, Michael Schumacher, Radek Szymanek, Paolo Viappiani, Xuan-Ha Vu, Steve Willmott) for the priceless comments provided during the many practice talks, for their willingness to help and for creating a pleasant working atmosphere. Special thanks to Mme. Decrauzat for her generous help with various administrative tasks.

Many thanks to the researchers I had the pleasure to meet at various events during the last five years. The discussions during the coffee-breaks and the sketches on the back of the envelope were always a great source of inspiration. While any list will inevitably be incomplete, I would like to mention here Maria Balcan, Ricardo Baeza-Yates, Rajat Bhattacharjee, Felix Brandt, Svet Braynov, Vincent Conitzer, Yiling Chen, Zoran Despotovic, Karen Fullman, Sébastien Lahaie, Kate Larson, Florent Garcin, Andrew Gilpin, David Parkes, David Pennock, Mohammad Mahdian, John Riedl, Paul Resnick, Neel Sundaresan, Arjun Talwar, Jean-Claude Usunier, Le-Hung Vu, Michael Wellman, and Makoto Yokoo.

My life as a PhD student would not have been the same without my family and friends from Lausanne. I hope they will not mind for not being individually mentioned here, but their support, company, and good spirit will be very much missed once we move away.

But most of all, I am grateful to my parents for their unconditional love and encouragement, to my brother Dan for his love, joyful company, and support, and to the two “girls” of my life, Carla and Lara, for their love and understanding, for all the joy they brought in my life, and for constantly reminding me about the important things in life. Carla and Lara, to you I dedicate this dissertation, in return for the many hours this work stole from you.





# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Summary of the Thesis . . . . .	5
<b>2</b>	<b>Trust, Reputation and Reputation Mechanisms</b>	<b>9</b>
2.1	Modelling Trust . . . . .	10
2.2	Reputation . . . . .	12
2.2.1	The nature of reputation information . . . . .	13
2.2.2	The role of reputation information . . . . .	14
2.3	Reputation Mechanisms for Online Systems . . . . .	25
2.3.1	Online Implementations . . . . .	26
2.3.2	Academic Models . . . . .	28
2.3.3	Empirical studies of Reputation Mechanisms . . . . .	30
2.3.4	Other Aspects related to Reputation Mechanisms . . . . .	32
<b>3</b>	<b>Truthful Signaling Reputation Mechanisms</b>	<b>33</b>
3.1	A Formal Model . . . . .	34
3.2	Incentives for Honestly Reporting Feedback . . . . .	36
3.2.1	Incentive-compatible Payment Mechanisms . . . . .	37
3.3	Automated Design of Incentive-compatible Payment Mechanisms . . . . .	38
3.3.1	Example . . . . .	40
3.3.2	Unknown Lying Incentives . . . . .	42
3.3.3	Computational Complexity and Possible Approximations . . . . .	42
3.4	Further Decreasing the Feedback Payments . . . . .	45
3.4.1	Using Several Reference Reports . . . . .	45

3.4.2	Filtering out False Reports . . . . .	48
3.5	Robust Incentive-Compatible Payment Mechanisms . . . . .	52
3.5.1	Dishonest Reporting with Unknown Beliefs . . . . .	52
3.5.2	Declaration of Private Information . . . . .	53
3.5.3	Computing Robust Incentive-Compatible Payments . . . . .	55
3.5.4	General Tolerance Intervals for Private Information . . . . .	56
3.6	Collusion-resistant, Incentive-compatible Rewards . . . . .	58
3.6.1	Collusion Opportunities in Binary Payment Mechanisms . . . . .	59
3.6.2	Automated Design of Collusion-resistant Payments . . . . .	63
3.6.3	Full Coalitions on Symmetric Strategies, Non-Transferable Utilities . . . . .	65
3.6.4	Full Coalitions on Asymmetric Strategies, Non-Transferable Utilities . . . . .	68
3.6.5	Partial Coalitions on Symmetric Strategies, Non-Transferable Utilities . . . . .	72
3.6.6	Partial Coalitions on Asymmetric Strategies, Non-Transferable Utilities . . . . .	74
3.6.7	Partial Coalitions on Asymmetric Strategies, Transferable Utilities . . . . .	79
3.7	Related Work . . . . .	81
3.8	Summary of Results . . . . .	84
3.A	Summary of Notation . . . . .	87
3.B	Generating Random Settings . . . . .	88
3.C	Cardinality of $Q(N_{ref})$ . . . . .	88
3.D	Proof of Lemma 3.6.1 . . . . .	89
3.E	Generating Random Binary Settings . . . . .	89
3.F	Proof of Proposition 3.6.7 . . . . .	90
3.G	Proof of Proposition 3.6.8 . . . . .	91
<b>4</b>	<b>Novel Applications of Signaling Reputation Mechanisms - QoS Monitoring</b>	<b>93</b>
4.1	Formal Model and Assumptions . . . . .	96
4.2	Interaction Protocol . . . . .	98
4.3	Implementation of a Prototype . . . . .	99
4.4	Incentive-compatible Service Level Agreements . . . . .	101
4.4.1	Example of Incentive-compatible SLAs . . . . .	102
4.5	Reliable QoS Monitoring . . . . .	104

---

4.5.1	Example . . . . .	108
4.6	Deterring Malicious Coalitions . . . . .	108
4.6.1	Using Trusted Monitoring Infrastructure . . . . .	110
4.7	Summary of Results . . . . .	112
4.A	Summary of Notation . . . . .	113
<b>5</b>	<b>Sanctioning Reputation Mechanisms</b>	<b>115</b>
5.1	Related Work . . . . .	118
5.2	The Setting . . . . .	119
5.2.1	Example . . . . .	120
5.2.2	Strategies and Equilibria . . . . .	121
5.2.3	Efficient Equilibrium Strategies . . . . .	125
5.3	Designing Efficient Reputation Mechanisms . . . . .	127
5.3.1	Probabilistic Reputation Mechanisms . . . . .	128
5.3.2	Deterministic Reputation Mechanisms using Mixed Strategies . . . . .	130
5.3.3	Deterministic Reputation Mechanisms with Pure Strategies . . . . .	135
5.3.4	Feedback Granularity . . . . .	137
5.4	A Mechanism for Obtaining Reliable Feedback Reports . . . . .	138
5.4.1	The CONFESS Mechanism . . . . .	141
5.4.2	Behavior and Reporting Incentives . . . . .	142
5.4.3	Implementation in the Reputation Mechanism . . . . .	144
5.4.4	Analysis of Equilibria . . . . .	145
5.4.5	Building a Reputation for Truthful Reporting . . . . .	149
5.4.6	The Threat of Malicious Clients . . . . .	154
5.4.7	Remarks . . . . .	155
5.5	Summary of Results . . . . .	157
5.A	Summary of Notation . . . . .	159
5.B	Appendix: Proof of Proposition 5.2.1 . . . . .	159
5.C	Appendix: Proof of Proposition 5.2.2 . . . . .	160
5.D	Proof of Proposition 5.4.2 . . . . .	161
5.E	Proof of Proposition 5.4.3 . . . . .	163

<b>6</b>	<b>Understanding Existing Online Feedback</b>	<b>165</b>
6.1	The Data Set . . . . .	166
6.1.1	Formal notation . . . . .	168
6.2	Evidence from Textual Comments . . . . .	168
6.2.1	Correlation between Reporting Effort and Transactional Risk . . . . .	174
6.3	The Influence of Past Ratings . . . . .	178
6.3.1	Prior Expectations . . . . .	179
6.3.2	Impact of Textual Comments on Quality Expectation . . . . .	180
6.3.3	Reporting Incentives . . . . .	181
6.4	Modelling the Behavior of Raters . . . . .	182
6.4.1	Model Validation . . . . .	183
6.5	Summary of Results . . . . .	185
6.A	List of words, $L_R$ , associated to the feature Rooms . . . . .	186
<b>7</b>	<b>Conclusions</b>	<b>187</b>
7.1	Directions for Future Work . . . . .	190
7.1.1	From “lists of reviews” to designed reputation mechanisms . . . . .	190
7.1.2	Signaling and sanctioning reputation mechanisms . . . . .	191
7.1.3	Factoring human behavior into reputation mechanism design . . . . .	191
7.1.4	Mechanisms for social networks and P2P systems . . . . .	192
7.1.5	Reputation mechanisms translated to other domains . . . . .	192
	<b>Bibliography</b>	<b>193</b>

# Chapter 1

## Introduction

In the beginning of the eleventh century, Khalluf ben Musa sends a letter from Palermo (Sicily) to his business partner, Yeshua ben Ismail from Alexandria. The writer is a respectable Jewish merchant, deeply unsatisfied with his partner's performance. The letter was written "on the eve of the New Year"<sup>1</sup> just after Khalluf learned much to his surprise, that the partnership with Yeshua brought him no profit in that year.

As it was customary in that period, the business partnership between Khalluf and Yeshua was based on reciprocally handling each other's goods: merchandise sent to Alexandria on Khalluf's account was to be handled (sold or rerouted to other destinations) by Yeshua, while Khalluf would provide the same services for Yeshua in Palermo. We learn from the letter that Khalluf kept his part of the deal and rigorously carried out Yeshua's instructions. Khalluf's goods, however, "remained unsold year after year" because Yeshua didn't take proper care of them. Khalluf is deeply disappointed with his partner's performance and concludes that: *"Had I listen to what people say, I never would have entered into a partnership with you"*.

The same letter carries on with a detailed account of Khalluf's services for Yeshua. Asarum<sup>2</sup>, cinnamon, clove, aloe, flax and ginger were all sold without problems. Pepper, however was in very low demand: "The price went down to 130, but still no one bought. [...] I held [the pepper] until the time when the sailing of the ships approached, in the hope it would rise. However, the slump got worse. *Then I was afraid that suspicion might arise against me and sold your pepper to Spanish merchants for 133.*".

Trading conditions for pepper changed "the night before the sailing of the ships". Boats with new buyers arrived, which took the price of pepper up. Khalluf was able to sell his own pepper for 140-142. He then surprisingly announces: *"But brother, I would not like to take the profit [from selling pepper in better circumstances] for myself. Therefore I transfer the entire sale to our partnership"*.

The letter concludes with various details about transaction costs, current prices and settlement of past payments, followed by Khalluf's intention to end the partnership with Yeshua: *"I ask you to settle my account with yourself and give the balance to my brother-in-law, for you are a very busy man."*

I chose to start this introduction with a letter from one thousand years ago because it reveals several important aspects that have dominated trading activities ever since the Middle Age. The first,

---

<sup>1</sup>The original letter is preserved by the Bodleian Library, Oxford; an English translation was published by Goitein (1973, page 120)

<sup>2</sup>a medical plant imported from the Orient

is the existence of *Word-of-Mouth*. Although geographically far apart, Khalluf had information about Yeshua even prior to their partnership. Khalluf had apparently heard rumors that Yeshua was not a reliable business partner; nevertheless, he chose to ignore them. Khalluf later discovered at his own expense that the rumors were right: “*Had I listen to what people say, I never would have entered into a partnership with you*”. Considering the delay and noise of communication channels at that time (letters and messages were often lost, or took weeks, even months, to get delivered), the fact that a merchant from Sicily had reliable information about another merchant from Egypt is remarkable.

The second aspect is a merchant’s strong incentive to carry out the instructions given to him by his business partner. Khalluf knew that circumstances for selling pepper were not right, but rather than breaking the implicit contract between him and Yeshua (a contract that required him to sell Yeshua’s pepper), he chose to sell the pepper for a lower price, avoiding thus the “*suspensions*” that “*might arise against*” him.

Third, and perhaps the most interesting, is Khalluf’s decision to act generously towards Yeshua, by sharing with him profits that resulted from providential fortune. Khalluf had no way of knowing that pepper prices would change within days, so his decision to sell Yeshua’s pepper for 133 dinars (considered a low price) was perfectly justifiable. In fact, any other merchant would have done (and apparently *did*) the same. “The night before the sailing of the ships” Khalluf was one of the few merchants to still have pepper for sale, which allowed him to obtain a much better price. Nevertheless, Khalluf generously split the profit with Yeshua, and compensated for the bad deal made a few days ago. Khalluf’s action is even more surprising considering that he wanted to end his partnership with Yeshua.

These three aspects provide strong evidence for the existence of an informal *reputation mechanism* among the traders from the Mediterranean basin (Greif, 1989). Although located in different parts of the Mediterranean coast, traders were part of the same community, sharing the same values and norms regarding appropriate business conduct. In the case of Khalluf and Yeshua, they both belonged to a group of Jewish traders mainly known as the Maghribi traders. One norm of the Maghribi was to exchange information about the performance of previous partners, such that the community would be informed about the past behavior of its members. Merchants anyway had to communicate frequently in order to react to changing market conditions, so the overhead for transmitting reputation information was insignificant. This is why Khalluf was able to learn that Yeshua is not a trustworthy partner prior to establishing a partnership.

The main role of the reputation mechanism was to establish incentives for individual traders to respect the implicit or explicit contracts with their business partners. Although local and religious laws have long before condemned the breaking of business agreements, they had little power to enforce contracts on an international level. Firstly, courts were hardly capable to verify claims about overseas activities. A merchant from Sicily, for example, could understate sicilian market prices to his business partner from Egypt, and thus pocket the difference for selling the latter’s goods for a higher price. Any complaint filed by the merchant in Egypt would probably remain unresolved, since neither the merchant, nor the court would be capable to exactly establish the prices obtained in Sicily. Secondly, the procedure for tracking emigrating merchants, if at all possible, was costly and lengthy. A borrower, for example, could set sail to a distant location and refuse to pay back his loan. International police cooperation was unimaginable in those days, so the fugitive would probably enjoy a prosperous life. One way or another, merchants could have easily cheated and remain unpunished, either because evidence could not be gathered against them, or because they could evade judgement.

The reputation mechanism, however, works based on two main principles. First, it establishes a merchant’s identity as a God fearing person, who would never deviate from the religious norms. These norms, among others, condemned cheating in business relations, which meant that the merchant could be safely trusted with business overseas. Second, the reputation mechanism creates a gap between the

---

expected future revenue of a merchant with good reputation as opposed to one with bad reputation. The Maghribi traders implicitly agreed to partner only with reputable peers, such that a person who lost his reputation was effectively excluded from the community. Cheating was thus deterred by economic sanctions rather than legal action: the momentary gain obtained from breaking a contract was largely offset by the future loss caused by a lost reputation. The punishment was so effective that Maghribis could strongly trust each other, and thus take full advantage of prosperous business opportunities arising from international trade. All this, with extremely rare appeal to local courts.

Khalluf's actions are suddenly easier to explain. When he sells Yeshua's pepper for a low price fearing the "*suspicious that might arise against [him]*", he explicitly reacts in order to maintain his reputation. Good reputation was a merchant's greatest asset, and nobody in their right mind would endanger it. The punishment for losing reputation was in fact so great, that Khalluf goes well beyond his contractual obligations. He shares with Yeshua the profit made by selling his own paper, hoping that his generosity will further strengthen his reputation. The fact that Khalluf wants to end his partnership with Yeshua has little importance: other community members will learn about Khalluf's behavior and will reward him by increased trust. The reputation mechanism was therefore effective even for short-time partnerships between the same two merchants, or towards the end of a merchant's life (the reputation would normally be inherited by ones offsprings).

By the late 13th century, however, the system based on inter and intra-community reputation experienced a harsh decline (Greif, 2002). As the number of communities increased and the communities grew in size, it became easier for malicious traders to change their identity, falsify their community affiliation and avoid punishments for unappropriate conduct. The cost of reputation management became simply too high (the larger the community, the longer and costlier was the propagation of negative feedback), and as a consequence, the threat of losing reputation became less effective. Starting with England and expanding all over Europe, the late Medieval period is associated with the development of state-owned contract enforcement institutions that make every individual accountable for misconduct. Slowly improving over the ages, the same system based in individual responsibility towards a central authority regulates economic transactions today. Reputation continues to play an important informal role for selecting capable business partners. Nevertheless, it is mostly lawyers and the threat of legal action, not public defamation, that keeps all parties to stick to the contractual agreements.

Why, then, should you carry on reading this thesis on reputation? The concept is old as society itself, and seems to have lost its importance over the years. The main reason is that we are looking at reputation mechanisms in a new context: that of online systems. The widespread adoption of the internet has created a global network that dramatically facilitates the communication between people and computers around the world. This technological advancement created new business opportunities, and changed the way people interact. Our day to day life has known numerous improvements ranging from basically free communication, to an explosion of choice regarding products and services.

Nevertheless, online interactions raised a whole new set of problems regarding trust and security. According to the US Federal Trade Commission approximately one half of all consumer fraud reports in 2006 were related to the Internet<sup>3</sup>. This is an area where reputation mechanisms can be useful, for the reasons I'll describe below.

First, the effective contract enforcing institutions that regulate offline trade have limited power in online settings. From this perspective, online trade resembles that of the Middle Ages: litigation claims are difficult and/or costly to verify, and legal punishments are difficult to impose on online users. Take for example eBay<sup>4</sup>, the online auctioning site where buyers and sellers from all over the world meet to exchange goods for money. Once a buyer wins an auction (by bidding the maximum amount) she pays directly to the seller the price of the good. Upon receipt of payment, the seller is supposed to deliver

---

<sup>3</sup>The full report is available at <http://www.consumer.gov/sentinel/pubs/Top10Fraud2006.pdf>

<sup>4</sup>[www.ebay.com](http://www.ebay.com)

(usually by post) the promised good directly to the buyer. If, however, the buyer does not receive the good (or receives a good of lesser quality), the legal resolution of the resulting conflict would be extremely expensive: establishing with reasonable certainty whether or not the seller shipped the right good would probably cost much more than the good itself. Moreover, even when the seller is clearly guilty, (a) establishing a link between the seller's online and physical identity is often not trivial, and (b) obtaining the cooperation of the police authority in the country of residence of the seller may be problematic (as, for example, the local legislation in that country does not recognize online fraud).

Second, a wide range of fraudulent online activities are not yet (or only starting to be) punishable by law. Phishing<sup>5</sup>, for example, has only recently been accepted as a criminal offence (the US senate voted anti-phishing laws in 2005; similar laws were introduced in UK in 2006<sup>6</sup>). Spam<sup>7</sup>, another example, is only partly punishable by law. On the other hand, many other activities are unlikely to be ever mentioned in a law. Falsifying information on Wikipedia (the online collaborative encyclopedia), deliberately providing wrong answers to a question posed on Yahoo! Answers, or tweaking web sites with commercial content in order to rank higher in search engines are just some of the examples.

Third, people like to think that online communities are natural extensions of their traditional, offline counterparts. They prefer to see the same kind of social control mechanisms functioning in online environments, where status and *reputation*, rather than abstract law, inspire people to behave cooperatively. Moreover, the underlying structure and philosophy behind the internet seem to disfavor a centralized solution. The internet itself is not subject to the control of a single authority, and functions in a distributed way. Self-organized, peer-to-peer applications designed explicitly to avoid any central control have flourished in the last years.

Forth, the cost of reputation management in online systems can be considered negligible. The communication cost for sending reputation information is significantly smaller online than offline. Moreover, information can get almost instantaneously to every corner of the world regardless the size of the community: updating reputation information in an online community of one thousand or one million members costs approximately the same.

In fact, online reputation mechanisms perform much better when they act in large communities. Higher number of reputation reports allow a better estimation of reputation information, less prone to errors in reporting and noise. For example, several Maghribi traders report having lost their reputation as a consequence of one, wrongful report. Today on eBay, on the other hand, one erroneous negative report amidst thousands of positive ones would hardly affect the business of an honest seller. Moreover, feedback can be discussed (e.g., sellers on eBay have the right to reply to negative feedback posted by clients) giving agents a powerful tool for defending themselves against mistaken or malicious reports. Therefore, the very same factors that led to the decline of Medieval reputation mechanisms (information delays and increased costs due to large communities) represent an opportunity for online reputation mechanisms.

Last but not least, reputation mechanisms can be designed individually for every community or market, in order to best take advantage of the particularities of that environment. For example, the mechanism designer can decide what information to elicit from the participants, how to aggregate feedback into reputation, and how to disseminate reputation information. While it makes perfect sense to have different reputation mechanisms acting on eBay or Amazon (governed by different rules, conceived to maximize the utility of the mechanism in the targeted community), it would certainly be strange to have different legal systems acting on the aforementioned markets. Locally designed reputation mechanisms thus have the potential to become more efficient than centralized law enforcement.

---

<sup>5</sup>The Merriam-Webster dictionary defines phishing as a scam by which an online user is duped into revealing personal or confidential information which the scammer can use illicitly.

<sup>6</sup><http://en.wikipedia.org/wiki/Phishing>

<sup>7</sup>Unsolicited usually commercial e-mail sent to a large number of addresses (the Merriam-Webster dictionary)



For all these reasons, online reputation mechanisms have known important commercial success (e.g., the success of markets like eBay or Amazon is partly attributed to their reputation mechanism) and have been the subject of intense research efforts in the last years. Today, they have been accepted as standard components of any collaborative system involving some risky interactions. We use online reputation mechanisms when searching for information, when booking trips, when purchasing electronic equipment or for choosing our service providers. Our increasing reliance on reputation information, on the other hand, creates “reputation premiums” so that products and providers with better reputation are able to gain more.

The economical value of online reputation raises questions regarding the trustworthiness of mechanisms themselves. Existing systems were conceived with the assumption that users will share feedback honestly. However, we have recently seen increasing evidence that some users strategically manipulate their reports. For some sites, the extent of *spam* feedback called for immediate solutions to eliminate suspicious reports. Known approaches include ad-hoc semi-manual, semi-automatic filters that provide satisfactory results only temporarily.

In this thesis I am concerned with a more systematic approach to making reputation mechanisms trustworthy. First, I am addressing the design of reputation mechanisms that are **incentive-compatible**, i.e., participants find it rational to share truthful feedback with the reputation mechanism. In my approach users get rewarded when submitting feedback, and the rewards are designed such that honest reporting is better than lying by at least some margin. Second, I address the problem of **collusion**. I show that some incentive-compatible rewards can also make small lying coalitions unstable, so that colluding agents find it more profitable to quit the coalition and report honestly. Third, I am interested in making the mechanism **robust** to noise and private information.

Additionally, I identify novel applications for reputation mechanisms where a reward scheme brings provable properties about the reliability of reputation information. One such example is quality of service (QoS) monitoring in decentralized markets of (web-)services. By using the feedback reported by the clients, the QoS monitoring process can be made more precise, reliable and cheaper.

Last but not least, I look at existing feedback and derive underlying factors that affect users when submitting feedback. The detailed understanding of the human feedback reporting behavior is hoped to generate new methods for aggregating reputation information that counterbalance the biases observed today.

## 1.1 Summary of the Thesis

I will start by a brief review of the state of the art related to online reputation mechanisms. Chapter 2 reviews the notions of trust, reputation and reputation mechanism, and presents some of the definitions that have been applied to these concepts in the context of online systems. For the scope of this thesis, *trust* will be understood as an agent’s (i.e., the truster) subjective decision to rely on another agent (the trustee) in a risky situation. The *reputation* of the trustee is one piece of information considered by the truster when taking the decision. Reputation information is assumed objective, and disseminated by a *reputation mechanism*. The role of reputation information is two-fold: first, to provide information about the hidden characteristics of the trustee that are relevant for the given situation (i.e., a signaling role) and second, to make future agents aware about any cheating that occurred in the past (i.e., a sanctioning role). The reputation mechanism, on the other hand, (i) provides the infrastructure for gathering feedback and for disseminating reputation information, (ii) defines the algorithms for aggregating feedback into reputation, and (iii) defines the interaction protocols and the incentives that motivate users to participate.

## Truthful Signaling Reputation Mechanisms.

The main role of signaling reputation mechanisms is to disclose hidden quality attributes of products or service providers based on the feedback of past users. Chapter 3 formally describes such mechanisms, and is mainly focused on the reporting incentives of the participants. Two factors make obtaining honest information difficult. First, feedback reporting is usually costly, and requires conscious effort to formulate and submit feedback. Many agents report only because they have alternative motives, as for example, extreme satisfaction or dissatisfaction with the service. This leads to a biased sample of reports and inaccurate reputation information.

Second, truth-telling is not always in the best interest of the reporter. In some settings for instance, false denigration decreases the reputation of a product and allows the reporter to make a future purchase for a lower price. In other contexts, providers can offer monetary compensations in exchange for favorable feedback. One way or another, external benefits can be obtained from lying and selfish agents will exploit them.

Both problems can be addressed by a payment scheme that explicitly rewards honest feedback by a sufficient amount to offset both the cost of reporting and the gains that could be obtained through lying. Fundamental results in game theory show that side payments can be designed to create the incentive for agents to report their private opinions truthfully. The best such payment schemes have been previously constructed based on *proper scoring rules* (Savage, 1971), and reward reports based on how well they update the current predictor about the observations of other agents (Miller et al., 2005). Honest reporting thus becomes a Nash equilibrium of the reputation mechanism.

I take an alternative approach using the principle of *automated mechanism design* (Conitzer and Sandholm, 2002). The mechanism is specified as an optimization problem where the payments minimize the budget required to achieve a certain margin for reporting the truth. The simplicity of a closed-form scoring rule is thus traded for a reduction of the required budget by approximately 50%. Furthermore, I find that a filtering mechanism (i.e., some reports are filtered out by the mechanism) designed automatically in conjunction with a payment mechanism can further reduce the cost of the mechanism by up to an order of magnitude.

Unfortunately, incentive compatible payments do not usually have honest reporting as a unique Nash equilibrium. Other *lying* equilibria exist, and bring forth the problem of *collusion*. Rational agents have no reason to report truthfully if they can do better by coordinating on a lying equilibrium with higher payoff.

One solution is to *fully implement* (Maskin and Sjöström, 2002) honest reporting, and have it as the unique equilibrium of the mechanism. Any lying coalition must play a non-equilibrium of the mechanism and is therefore unstable: i.e., individual colluders will find it rational to quit the coalition. A second, less strict alternative is to design a mechanism where truth-telling is the “best” equilibrium. In this case lying coalitions may be stable, but are not profitable. Assuming that rational agents will not make the effort to collude on a worse-paying lying strategy, the mechanism is collusion-resistant.

I investigate different collusion scenarios where (i) some or all agents collude, (ii) colluders can coordinate or not on asymmetric strategy profiles, and (iii) utilities are transferable or not. For each scenario I apply the principle of automated mechanism design to specify the incentive-compatible collusion-resistant payment mechanisms.

## Novel Applications of Signaling Reputation Mechanisms.

An increasing fraction of a modern economy consists of services. Services are generally provided under a contract (or Service Level Agreement) that fixes the type and quality of the service to be provided as well as penalties if these are not met. An essential requirement for such service provisioning is to be

able to monitor the quality of service that was actually delivered. As the monetary value of individual services decreases, the cost of providing accurate monitoring takes up an increasing share of the cost of providing the service itself. For example, with current technology, reliably monitoring the quality of a communication service requires constant communication with a neutral third party and would be almost as costly as providing the service itself. The cost of this monitoring remains a major obstacle to the wider adoption of a service-oriented economy.

Chapter 4 presents an alternative Quality of Service (QoS) monitoring mechanisms based on feedback provided by the clients. Clients are monitoring the quality of the service provided to them, and periodically report feedback to a reputation mechanism. The reputation mechanism aggregates the reports and estimates the QoS delivered by a provider to a group of clients. In this way, (a) several reports can be compressed into one message and communication overhead is reduced, (b) the monitoring process is as precise as possible because the monitor does not need to sample the service requests, and (c) the service provider cannot directly tamper with the monitoring process.

Accurate monitoring, however, requires clients to report the truth. The reward schemes for signaling reputation mechanisms perfectly fit to this setting, and allow the construction of efficient and robust QoS monitoring systems.

### **Sanctioning Reputation Mechanisms.**

A second important role of reputation mechanisms is to provide incentives to rational agents to behave cooperatively, even when direct benefits can be obtained from cheating. The main idea is that present feedback determines the future reputation of an agent, and implicitly affects the future revenues accessible to the agent. When carefully designed, reputation mechanisms can make it such that the momentary gain obtained by cheating is offset by the future losses caused by a bad reputation. Cheating is appropriately sanctioned by the reputation mechanism, and this encourages every participant in the market to behave cooperatively.

Chapter 5 starts by addressing the design of efficient sanctioning reputation mechanism. The internet gave us almost complete freedom over design parameters such as (i) the granularity of feedback requested from the users, (ii) the algorithms for aggregating feedback into reputation information, or (iii) the extent and form of reputation information dissemination. Dellarocas (2005) discusses binary reputation mechanisms (where seller, for example, can cooperate or cheat, and buyers can receive high or low quality) and investigates the effect these design decisions on the efficiency of the mechanism. The contribution of this chapter is to extend Dellarocas' results to general settings, where the seller can choose between several effort levels, and buyers can observe several quality levels.

I find, for example, that a reputation mechanism where the reputation has a binary value (i.e., can be either good or bad) can be equally efficient to other mechanisms where the reputation is infinitely finer grained. Moreover, I show that efficient reputation mechanism can function by considering a finite window of past reports, and finitely grained feedback values. All of these results agree with the findings of Dellarocas for the binary case.

The second part of Chapter 5 discusses a mechanism (which I have named CONFESS) for encouraging the submission of honest feedback. CONFESS works by comparing the feedback submitted by the buyer to the feedback implicitly submitted by the seller. After every transaction, the seller is allowed to acknowledge failures and reimburse the affected buyer. If, however, the seller does not reimburse the buyer, and the buyer submits negatived feedback, the reputation mechanism concludes that one of the agents is lying, and punishes them both. This simple mechanism supports and equilibrium where all sellers cooperate, and all buyers report the truth. Moreover, it allows the buyers to build a reputation for always reporting the truth, which in the end, can be proven to limit the amount of false information received by the reputation mechanism in any pareto-optimal equilibrium.

## Understanding the Human Feedback Reporting Behavior.

New reputation mechanisms could greatly benefit from the results obtained in the previous chapters. However, we will still see a great number of commercial implementations that (i) provide little or no reporting incentives, and (ii) aggregate feedback into reputation information in trivial ways. Since these naive mechanisms will continue to be important channels for *Word-of-mouth*, and users will continue using them for taking purchasing decisions, it is important to better understand the factors and interaction that drive users to participate and report feedback.

Chapter 6 extends a recent line of research by investigating the factors that (i) drive a user to submit feedback, and (ii) bias a user in the rating she provides to the reputation mechanism. For that, I consider two additional sources of information besides the numerical ratings provided by the user: first the textual comments that accompany the reviews, and second, the timely sequence of reviews submitted by previous users. All hypotheses were validated by statistical evidence from hotel reviews on the TripAdvisor website.

The first conclusion is that groups of users who amply discuss a certain feature are more likely to agree on a common rating for that feature. Moreover, the users who write lengthy textual comments are not outliers, and their reviews are generally regarded as more helpful by the other users. This observation can be used to construct feature by feature estimates of quality, where for each feature, the mechanism averages only the ratings corresponding to lengthy discussion of that feature.

The second result reveals a correlation between the effort spent in writing a review, and the risk (in terms of making a bad choice) associated to a hotel. Booking a room in a high-end hotel is argued to be a riskier transaction than booking a low-end hotel, as the traveler pays more without having the guarantee that she gets the service she wants. Human raters apparently feel motivated to decrease the decision risk of future users, hence they spend more effort rating the high-end hotels.

Third, the rating expressed by a reviewer appears to be biased by the reviews submitted by previous users. The information already available in the forum creates a prior expectation of quality, and changes the user's subjective perception of quality. The gap between the expectation and the actual quality is reflected in the rating submitted to the site.

Last but not least, the timely sequence of reviews also reveals another motivation for reporting feedback. A great proportion of users submit ratings that are significantly different from the average of previous ratings. This leads to the conclusion that human users are more likely to voice their opinion when they can bring something different to the discussion, and can contribute with new information.

## Chapter 2

# Trust, Reputation and Reputation Mechanisms

In 1970 Akerloff published an article (Akerlof, 1970) that warned about the possible collapse of markets with asymmetric information. He gives as an example a market of second-hand cars, where buyers cannot tell before the purchase if the quality of a second-hand car is good or bad. The sellers, who of course know the real quality of the car they are selling, have the incentive to exploit the lack of information of buyers, and exaggerate the quality of their cars in order to obtain higher prices. Buyers, nevertheless, anticipate such strategic incentives, and discount the price they are willing to pay such that it reflects the average quality on the market. This makes it uninteresting for sellers with good quality to stay in the market. They don't have any means to signal that their cars are better than the average, and therefore, the amount they would receive from the buyers does not make the sale worthwhile. The fact that good quality sellers will be driven away from the market further decreases the average quality expected by the buyers, and therefore the price they would be willing to pay. This, consequently drives sellers with medium quality away from the market. By induction, we end up with a "Market of Lemons" where only the worst-quality (i.e., the "lemons") is available in the market. Such a market will clearly collapse.

You might consider the above example slightly exaggerated. Some of you may have already bought a second-hand car, and never had any problems with it. You have probably test-driven the car, looked at the service history, or had it inspected by a workshop in order to verify that the actual product is as advertised. Second-hand car markets have devised mechanisms to ensure that buyers *trust* the quality description offered by the seller, and as a consequence, such markets continue to flourish.

Electronic markets, on the other hand, are much closer to the assumptions set forth by Akerloff. Take eBay, for example, the auctioning site where sellers post item descriptions and wait for buyers to bid online. Although side-communication between the seller and the buyers is possible (through e-mail), the same trust-enabling mechanisms that function in markets for second-hand cars are no longer available. A buyer from California would never fly over to the East coast in order to inspect an interesting bicycle posted by a seller from Boston. Even if the buyer and the seller were from the same region, eBay sellers can rarely accommodate visits. They often operate from the comfort of their living room, and sell thousands of items. The increased cost of having proper premises where prospective buyers can inspect the items makes the whole business unprofitable.

Electronic markets should therefore have a serious trust problem. Sellers have clear incentives to overestimate the quality of their goods, and buyers should not trust the descriptions they see online.

According to Akerloff these conditions lead to a market of lemons, and eventually, to the disappearance of trade.

Nevertheless, eBay cleared auctions worth more than 14 billion dollars in the first quarter of 2007<sup>1</sup>. There seems to be significant trust between the bidders and the seller, which encourages profitable trade. The success of eBay is mostly attributed to the feedback mechanisms that allows bidders to check the previous history of a particular seller. Following every transaction, eBay encourages both participants to give feedback about their partner. Feedback is positive, negative or neutral, but may also contains a textual comment. The history of past feedback received by a seller determines his reputation, which strongly influences buyers when deciding which seller to trust.

Successful transactions on eBay, like in most electronic markets, seem to involve three important concepts. First, the *trust* of the buyers who decide to bid for the item offered by a seller. Second, the *reputation* of the seller, which inferred from the history of previous transactions of that seller, allows buyers to decide whether and whom to trust. Third, the *reputation mechanism* which collects the information (feedback) about previous interactions and makes it available to future users.

This chapter attempts to clarify the notions of trust, reputation and reputation mechanisms. I will start with trust, and review the many definitions that have been applied to this concept. Trust has been defined as “both a noun and a verb, as both a personality trait and a belief, and as both a social structure and a behavioral intention” (McKnight and Chervany, 2001). The definition of trust and its corresponding meaning is a much disputed issue among the computer science community. Since the human understanding of the notion of trust is much too complex to be modelled within an artificial system, authors usually consider just facets of the notion of trust, and define it corresponding to their needs. I will do the same in this thesis, and regard trust as a subjective decision made by an agent regarding another agent, in a specific context.

I will continue with reputation, and consider it to be one important piece of information agents consider when taking trust decisions. While trust is subjective (e.g., the fact that agent A trusts agent B in a certain context, does not mean that another agent C will also trust B in exactly the same context) reputation is mainly assumed objective, so that all agents view it the same: e.g., the reputation of B is the same when viewed by A or C. I will also investigate the role of past information in building a reputation, and the influence of reputation on future decisions.

The last part of this chapter is reserved for an overview of existing (both commercial and academic) reputation mechanisms.

## 2.1 Modelling Trust

The Oxford English Dictionary defines trust as “confidence in or reliance on some quality or attribute of a person or thing”. Deutsch (1962) also identifies a notion of risk associated to trusting somebody and therefore, “one trusts when one has much to lose and little to gain”. The same notion of risk is mentioned by Gambetta (1988): “Trust is the subjective probability by which an individual A expects that another individual B performs a given action on which its welfare depends”.

In psychology, trust occurs when an individual (Alice) is confronted with an ambiguous path, a path that can lead to an event perceived to be beneficial, or to an event perceived to be harmful. Alice perceives that the occurrence of the beneficial, respectively harmful event is contingent on the behavior of another person, Bob. Moreover, Alice perceives that the negative impact of the harmful event is stronger than the positive impact of the beneficial event. If Alice chooses to take such an ambiguous

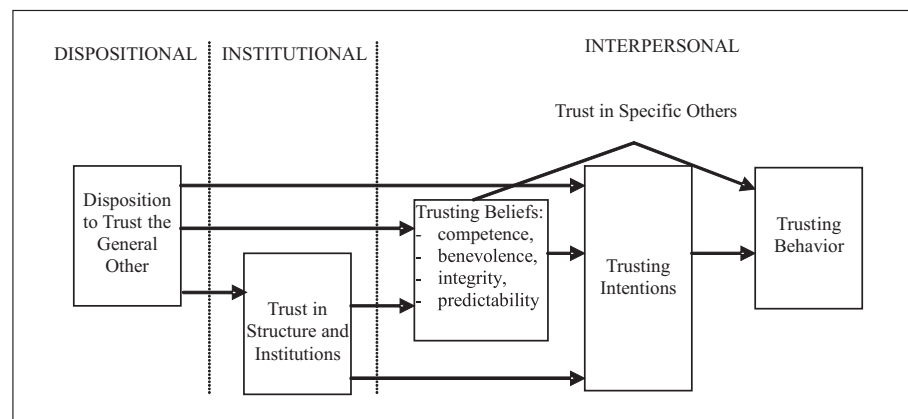
---

<sup>1</sup>Information released by eBay in its first quarter 2007 financial information report (<http://investor.ebay.com>)

path, she makes a trusting choice; otherwise she makes a distrustful choice (Marsh, 1994; Lamsal, 2001). Marsh (1994) provides an extensive formalization of this concept of trust.

Note that trust, primarily, is the belief that a person (the trustee) will act in the best interest of another (the truster) in a given situation, even when controls are unavailable and it may not be in the trustee's best interest to do so (Marsh and Dibben, 2005). However, trust also involves a decision or commitment of the truster to rely on the trustee (Sztompka, 1999).

McKnight et al. (2002) and McKnight and Chervany (2001) develop a conceptual topology of the factors that contribute towards trust and distrust decisions and define as subsets of the high level concepts measurable constructs for empirical research (see Figure 2.1). In their model, the *trust-related behavior* refers to a person's voluntary reliance on another person with a feeling of relative security, even though negative consequences are possible. The truster thus gives the trustee some measure of power over himself/herself.



**Figure 2.1:** The McKnight et al. interdisciplinary model of trust.

Trusting behavior is influenced by *trusting intentions* and *trusting beliefs*. *Trusting intentions* refer to one's subjective willingness to depend on the other party with a feeling of relative security, in spite of lack of control over that party, and even though negative consequences are possible. Intentions may depend on the general state of mind of the truster, but also on the specific context and behavior in mind: e.g., the intention to trust somebody with keeping information private is generally more likely than the intention to trust somebody with ones life.

*Trusting beliefs*, on the other hand, refer to the extent to which one believes, with feeling of relative security, that the trustee has characteristics beneficial to one. Qualities of interest are: *competence* (the truster securely believes that the trustee has the necessary abilities to successfully do what the truster asked it to do), *benevolence* (the truster securely believes that the trustee cares about and is motivated to act in the truster's interest), *integrity* (one securely believes that the other party makes good faith agreements) and *predictability* (one securely believes that the other person's actions, good or bad, are consistent enough that one can foresee them).

*Institutions* can also encourage trusting behavior. For example, one may believe that proper laws and authorities are in place in order to increase the probability of successful outcomes in risky situations. Such structures and institutions encourage trusting behavior generally, across all trustees and contexts. Last, but not least, trusting behavior also depends on one's general *disposition to trust*: i.e., the extent to which one displays a consistent tendency to be willing to depend on general others across a broad

spectrum of situations and persons. Since disposition to trust is a generalized tendency across situations and persons, it has a major influence on the trust-related behavior only in novel, unfamiliar situations when we cannot rely on the other components. Parts of disposition to trust are faith in humanity (assumption that people in general are usually honest, benevolent, competent and predictable) and trusting stance (belief that regardless of what one assumes about other people generally, one assumes that one will achieve better outcomes by dealing with people as they were well-meaning and reliable).

Several works extend the model of McKnight, Choudhury, Kacmar and Chervany by explicitly modeling the risk involved by trusting decisions. Povey (1999), Dimitrakos (2003), Manchala (1998), and Jøsang and Lo Presti (2004) all argue in favor of trust policies, or algorithms, that based on the available information quantifying risk and opportunities (i.e., transaction value, probability of failure, contingency outcomes, etc.) output the decision to trust or not.

Kramer (2001) focuses on explaining how humans decide whether to trust or not. He proposes the Social Auditor Model (SAM) where the decision process is characterized by a set of rules. These rules can be classified as:

- *interpretation rules*, i.e., rules that help us categorize a given trust dilemma and prescribe the sort of evidence we should look for when trying to assess the another decision maker's trustworthiness, and
- *action rules*, i.e., rules about what behaviors we should engage in when responding to those interpretations.

Rules are applied based on direct, momentary perceptions from the environment, but also using a personal mental model of the world. The Social Auditor Model also includes a post-decision *auditing* process, when the decision maker evaluates the results of his application of the rules, and updates accordingly both the set of rules, and their mental model of the world. Kramer also studies the efficiency of different rules and strategies that can be used within an artificial society.

Bacharach (2002), Falcone and Castelfranchi (2001) and Castelfranchi and Falcone (2000) look at the dynamics associated with the notion of trust. Trust and distrust responsiveness (trust from the trustee increases the probability of cooperative behavior from the trustor, while distrust from the trustee increases the probability of defective behavior) are presented as facts of human behavior. They also address the dialectic link between trust and degree of control.

From all facets of trust briefly described above, the one I'll be using most in this thesis refers to the *interpersonal* trust between the truster and the trustee. I will abstract away both subjective (e.g., disposition or intention to trust) and contextual (e.g., institutions, situations, etc) considerations involved in trusting decisions, and assume that decisions rely entirely on the beliefs of the truster regarding the trustee. Moreover, I will generally unite competence, benevolence, integrity and predictability under one single measure, referred to as *reputation*.

## 2.2 Reputation

Similar disputes exist around the notion or reputation. In economics (and game theory) reputation measures an agent's ability to commit for a certain strategy. In computer science, reputation is usually an ad-hoc measure of trustworthiness. Sociologists, anthropologists and biologists also use reputation to characterize the position of an individual within the society.

Nevertheless, all these different aspects of reputation have two things in common. First, reputation



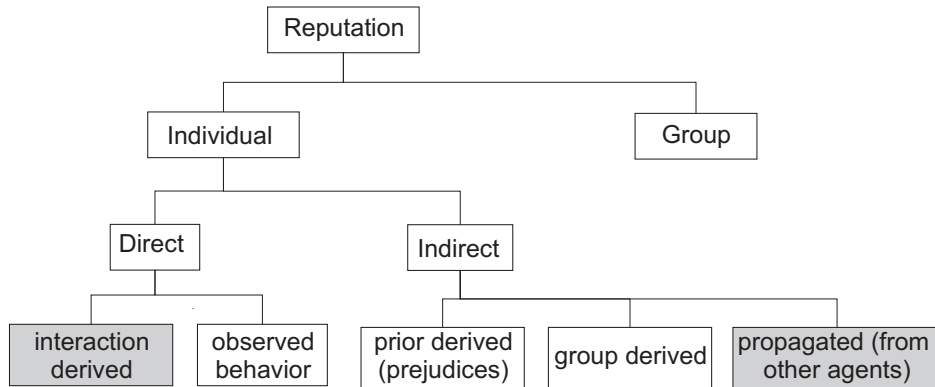


Figure 2.2: Reputation Topology (presented by Mui et al.)

is *distilled* about information available from an agents past behavior. Second, reputation information seems to be important when taking trust decisions.

### 2.2.1 The nature of reputation information

Mui, Halberstadt, and Mohtashemi (2002a), Mui, Mohtashemi, and Halberstadt (2002b) and Mui (2003) provide an extensive classification of reputation (Figure 2.2) by the means of collecting it. At the topmost level, reputation is derived from the individual characteristics of the agent, or inherited from the group the agent belongs to. The group component of reputation is usually important in our society: for example, the quality of job applicants is often judged by the prestige of the companies they previously worked for. Most online systems, however, only consider the individual component of reputation. Notable exceptions are the work of Sabater and Sierra (2001) and Halberstadt and Mui (2002). The former studied the social dimension of reputation which is inherited from a group reputation; the latter proposed a hierarchical group model and have studied group reputation based on simulations.

Individual reputation can be further split as *direct* or *indirect* reputation. Direct reputation comes from the first-hand information (either from direct encounters or from direct observations of others' encounters) one has about the agent whose reputation is in question. Indirect reputation, on the other hand, is based on second-hand information provided by others. Prior-derived reputation refers to the prior beliefs (prejudices) agents bring with them about strangers, group-derived indirect reputation refers to the prior beliefs agents have about groups, while propagated reputation is formed based on the feedback from other agents.

Of all reputation categories above, online reputation mainly consists of direct, interaction derived individual reputation, and indirect individual reputation propagated from other agents (shaded in Figure 2.2). The two categories naturally correspond to:

- information an agent directly obtains from her interactions with other agents, and,
- feedback reported to a reputation mechanism that can be used by other agents.

Two more aspects worth mentioning are that reputation is context-dependent (Mui et al., 2002a), and usually has a personal scope (Conte and Paolucci, 2002). For example, someone's reputation as a

cook, should be separated from that person's reputation as a singer. Most online reputation mechanisms, however, maintain only one value of reputation, which characterizes all activities of an agent within the system. As online reputation mechanisms are usually used in a narrow context (e.g., on eBay, reputation characterizes the trading behavior of a seller, on Slashdot reputation refers to the quality of the articles submitted by a user) this simplification is not over-restrictive. There are, however, cases where finer-grained reputation would improve the overall performance of the system (e.g., empirical studies of eBay reputation show that human users take into account different dimensions of reputation (Pavlou and Dimoka, 2006), like speed of delivery, accuracy of item description, etc.).

Likewise, agent A may perceive agent B's reputation differently from agent C. Nevertheless, since most online reputation mechanisms are centralized, the probability that individual agents will have private information that greatly differs from the information maintained by the reputation mechanism is insignificant. Therefore, the assumption that reputation information is global (viewed uniformly by all agents in the system) is usually correct. One notable exception are reputation mechanisms in peer-to-peer systems (Despotovic, 2005). Here, different agents may judge reputation differently because the underlying P2P network gives them different views on the set of feedback reports recorded by the reputation mechanism.

## 2.2.2 The role of reputation information

If reputation is entirely constructed from information from the past, why and how is it helpful in future decisions? Intuitively, there are two answers to this question. First, by looking at the history of an agent's behavior, one can hope to discover qualities and habits that are likely to remain valid in the near future. On eBay, for example, a seller with thousands of positive feedback reports on his record is probably committed to behave honestly. There is no guarantee whatsoever that the seller will also honor the next transaction with the same honesty and care as the previous ones. Nevertheless, given that the seller behaved honestly so many times in the past, a buyer is likely to believe that honesty is an inherent characteristic of the seller, unlikely to change from one day to the next.

As another example, consider a hotel who has been subject to a number of reviews from travellers who used to spend a night there. The fact that the hotel has a good reputation for having offered great service in the past, creates a belief that a traveller going there next week will also be satisfied. The management of the hotel might, of course, have radically changed the quality of service (e.g., by cutting costs and firing some of the staff); nevertheless, the probability of this happening under normal circumstances is so small, that as long as further reviews do not say otherwise, one continues to believe the hotel will offer great service in the immediate future.

Therefore, the first role of reputation information is to *signal* lasting attributes (or *types*) characterizing the behavior of an agent (Kuwabara, 2003). This role is mostly descriptive; reputation provides information to decision makers, without really specifying what the decision should be.

The second role of reputation is normative in nature, and encourages decision makers to punish agents with a bad reputation. This *sanctioning* role of reputation has as a main objective to encourage long-term honest behavior by eliminating momentary incentives to cheat. The whole idea is simple: when trusting decisions are contingent on reputation information, the present actions reflect on future reputation and therefore on future revenues. Cheating may be profitable for the moment, but when rational agents factor the future losses provoked by the resulting bad reputation, honest behavior becomes the optimal strategy. An eBay seller, for example, may refuse to ship the product to the buyer after receiving the payment. This gives the seller the momentary gain of not losing the good, but also decreases his future reputation. Future buyers may account for the risk of being cheated and bid less for the goods auctioned by the seller. The overall future loss may be greater than the momentary gain

		Player 2	
		Confess	Don't Confess
Player 1	Confess	-3, -3	0, -4
	Don't Confess	-4, 0	-1, -1

**Figure 2.3:** The Prisoners' Dilemma.

obtained from cheating, and therefore, shipping the good in the first place becomes the optimal strategy.

The two roles of reputation are complementary, and solve the two important problems associated with online markets (Dellarocas, 2006a). The signaling role acts against information asymmetries, and allows agents to accurately identify capable partners. The sanctioning role, on the other hand, acts against cheating incentives and encourages honest behavior.

There are settings, however, where one role of reputation is predominant. Online reviews such as those posted by Amazon, ePinions or BizRate act primarily as signaling devices, and convey information about products whose quality remains unchanged over a period of time. Buyers who read the reviews are mainly interested in assessing the overall quality of the product, and do not generally worry that someone out there will explicitly decide to ship them a product of significantly lower quality. Defective items may of course reach end customers, but this unfortunate event is the result of chance (a random process independently distributed across all buyers), not malice.

On eBay, on the other hand, reputation can be viewed primarily as a sanctioning mechanism. Ratings reflect a seller's ability to deliver the product as described in the auction, and do not reflect the quality of the product itself. Assuming that all eBay sellers are equally capable of honest behavior, the role of the reputation mechanism is to promote honest trade, without trying to identify the sellers that auction better quality products.

The distinction between the two roles of reputation is crucial for the design of effective markets. Before going into more details for each of these roles, I will briefly introduce some basic notions of game theory.

### Basic notions of Game Theory

Game Theory provides a set of tools that allow us to analyze and predict how self-interested decision makers interact. Game Theory works on simplified, abstract models of reality (also known as *games*) that define the actions that may be taken by the agents (or the *players*), the information available to them, and the possible outcomes of the interaction, contingent on the actions (or the *play*) of the agents. One important assumption is that players are *rational* in the sense that they wish to maximize their welfare (or *utility*) and therefore reason *strategically* about their actions and about the actions of their opponents.

One famous example of a game is the *Prisoners' Dilemma*, and corresponds to the following situation. Two suspects in a crime are arrested and placed in separate cells. If they both confess, enough evidence can be gathered against both of them and they will be sentenced to 3 years in prison. If only one confesses, he will be freed and used as a witness against the other. The prisoner who didn't confess will be sentenced to 4 years in prison. Nevertheless, if neither of the prisoners confesses, they will both be convicted for minor offenses and get sentenced to one year.

A schematic representation of the game is shown in Figure 2.3. Player 1 (the first prisoner) has a choice between two actions (Confess or Don't Confess) which label the rows of the table. The actions of Player 2 (the second prisoner) are in this case the same, and label the columns of the table. Every cell in the table contains the outcome of the game when the players choose the corresponding actions. The first number in every cell is the payoff (values are negative because spending time in jail is usually associated to a loss, not a gain) to player 1, the second number is the payoff to player 2.

In the Prisoners' Dilemma both players choose their actions once and for all, independently from each other. Such games are called **strategic games**, and formally consist of:

- a finite set  $N$  of players; e.g.,  $N = \{\text{Player 1, Player 2}\}$ ;
- for each player  $i \in N$ , a nonempty set  $A_i$  (the set of actions available to player  $i$ ); e.g.,  $A_1 = A_2 = \{\text{Confess, Don't Confess}\}$ ;
- for each player  $i \in N$  a preference relation  $\succsim_i$  on  $A = \times_{j \in N} A_j$  (the preference relation of player  $i$ ).

A **pure strategy**  $s_i$  of player  $i$  is an action  $a_i \in A_i$ . In the Prisoners' Dilemma, a pure strategy of the first player is, for example, to confess. A **mixed strategy** of player  $i$  is a probabilistic combination of actions,  $s_i = \sum_{a_k \in A_i} \alpha_k \cdot a_k$ , where the action  $a_k$  of player  $i$  is chosen with probability  $\alpha_k$ . Of course,  $\sum_{a_k \in A_i} \alpha_k = 1$ . A mixed strategy in the Prisoners' Dilemma is to randomly choose whether to confess (with some probability  $\alpha$ ) or not (with the remaining probability  $1 - \alpha$ ).

An action (or a pure strategy) **profile** is a collection of action values  $a = (a_j)_{j \in N}$ , specifying one action  $a_j \in A_j$  for every player  $j$ . (Confess, Don't Confess), for example, is an action profile of the Prisoners' Dilemma, where the first player confesses and the second doesn't. I will also use the notation  $a_{-i}$  to define an action profile for all players except  $i$ : i.e.,  $a_{-i} = (a_j)_{j \in N \setminus \{i\}}$ , and  $a = (a_{-i}, a_i) = (a_i)_{i \in N}$  for any player  $i$ . The set  $A_{-i} = \times_{j \in N \setminus \{i\}} A_j$  contains all possible action profiles  $a_{-i}$  of players other than  $i$ . When there are only two players, as in the Prisoners' Dilemma,  $a_{-1}$  for example is the action taken by player 2.

An action profile  $a = (a_j)_{j \in N}$  is also an **outcome** of the game. The preference relation  $\succsim_i$  provides an ordering (as viewed by player  $i$ ) of the possible outcomes of the game. For example, player 1 orders the outcomes of the Prisoners' Dilemma in the following way: (Confess, Don't Confess)  $\succsim_1$  (Don't Confess, Don't Confess)  $\succsim_1$  (Confess, Confess)  $\succsim_1$  (Don't Confess, Confess). Under a wide range of circumstances the preference relation  $\succsim_i$  of player  $i$  can be represented by a **payoff function**  $u_i : A \rightarrow \mathbb{R}$  (also called *utility function*), in the sense that  $u_i(a) \geq u_i(b)$  whenever outcome  $a$  is preferred to outcome  $b$  (i.e.,  $a \succsim_i b$ ). Values of such a function are called **payoffs** (or utilities), and they are the ones we have put in the cells of the table in Figure 2.3. We will denote a game by  $\langle N, (A_i), (\succsim_i) \rangle$  or by  $\langle N, (A_i), (u_i) \rangle$ .

Game theory predicts the class of "reasonable" outcomes that may occur in a game. Let us reexamine the payoff and the actions of the players in the Prisoner's Dilemma. It is immediately obvious that for player 1, the best thing to do is to confess:

- if player 2 confesses, confession gives  $-3$  while not confessing give  $-4$ ;
- if player 2 doesn't confess, confession gives  $0$  while not confessing gives  $-1$ .

The game being symmetric, player 2 can make exactly the same reasoning, and therefore, both will end up confessing. The dilemma in this game is that rational reasoning determines both players to confess, while both would benefit if neither confesses. The outcome where neither of the players confesses is,

		Player 2	
		h	l
Player 1	H	2, 3	0, 2
	L	3, 0	1, 1

**Figure 2.4:** The Product-choice game.

**pareto-optimal**<sup>2</sup>, but unfortunately not “stable”, since neither of the players can trust the other not to confess.

Confessing is a **dominant strategy** in the Prisoners’ Dilemma, since no matter what the other player does, confessing gives the highest payoff. A profile of dominant strategies, one for each player, defines a **dominant equilibrium** of the game. Formally, a dominant equilibrium is defined as a strategy profile  $s = (s_i)_{i \in N}$ , where every player  $i$  plays a dominant strategy  $s_i$ : i.e.  $u_i(s_i, s_{-i}^*) > u_i(s'_i, s_{-i}^*)$  for all strategies  $s'_i \in A_i, s'_i \neq s_i$  of player  $i$ , and all strategies  $s_{-i}^* \in A_{-i}$  of the other players.

While dominant equilibrium is a very convenient (and strong) equilibrium concept, it does not exist in all games. Consider the product-choice game in Figure 2.4, where a provider can exert high effort (H) or low effort (L) to produce a good. The consumer on the other hand (player 2) can choose to buy either a high priced (h) or a low priced (l) product from the provider. Player 1, for example, can be considered a financial investor who offers two investment products, one requiring a high upfront investment, the second requiring only a small investment. Another example is a restaurant who offers on its menu both fast-food and fancy cuisine (Mailath and Samuelson, 2006).

The client most prefers the expensive product if the provider exerts high effort; nevertheless, if the provider exerts low effort, the buyer is better off with the cheap product. One can gain more by entrusting a high sum of money to a diligent financial investor, and one gets more pleasure from eating a fine dinner prepared by a careful chef. The same person, however, would prefer the low investment or the fast food if the investor, respectively the restaurant chef is doing a hasty job. The optimal action of the client therefore depends on the action chosen by the provider, and there is no strategy for the client that dominates all others.

The provider, however, does have a dominant strategy which prescribes low effort. The client anticipates this, and will therefore choose the low priced product. The resulting outcome is a **Nash Equilibrium**, one of the most commonly used equilibrium notions in game theory. The Nash Equilibrium defines a “stable state”, where unilateral deviations are not profitable: no player can gain by choosing another strategy, as long as all other players continue playing the equilibrium. Formally, the strategy profile  $s = (s_i)_{i \in N}$  is a Nash Equilibrium if and only if  $u_i(s_i, s_{-i}) \geq u_i(s'_i, s_{-i})$  for all players  $i$  and all strategies  $s'_i \neq s_i$ . A mixed strategy Nash equilibrium always exists (Nash, 1950).

Not all situations can be modeled through games where players simultaneously choose their actions. In online auctions, for example, the seller ships the good only after the buyer has paid the required amount. In chess, another example, black and white alternate their moves, and each player is informed about the move made by the opponent before taking the next action. Games where players move in turns are called **extensive form game**, and are usually represented as in Figure 2.5.

Here, a chain-store (player  $CS$ ) has branches in several cities. In each city there is a single potential competitor, player  $k$ . The potential competitor decides first whether or not she wants to compete with

<sup>2</sup>A payoff profile is pareto-optimal when none of the players can obtain a better payoff without decreasing the payoff of at least one of the other players.




Figure 2.5: The Chain-Store Game.

$CS$  by opening another store. If the competitor decides to enter the market (action  $In$ ), the chain store can either fight (action  $F$ ), or cooperate with (action  $C$ ) the competitor. If the competitor chooses not to compete (action  $Out$ ) the chain-store does not have to take a decision.

The payoffs of the players are represented as numbers associated to terminal nodes: the first number is the payoff of the chain-store, the second number is the payoff to the competitor. When challenged, the chain-store prefers to cooperate rather than fight (i.e., fighting is expensive). However, the chain-store best prefers not to have the competition in the first place. The competitor is better off staying out, rather than enter and being fought by the chain-store; nevertheless, the competitor obtains the highest payoff when she enters and the chain-store cooperates.

Formally, an extensive form game is defined by the following components:

- a set  $N$  of players, e.g.,  $N = \{CS, k\}$  for the Chain-Store game above;
- a set  $H$  of all possible **histories** of the game. A history is a sequence of actions taken by the players that correspond to a valid path of moves in the game. For the Chain-Store game the set of possible histories is  $H = \{\emptyset, In, (In, F), (In, C), Out\}$ , where  $\emptyset$  is the initial history (no move has occurred yet) and the subset  $Z = \{(In, F), (In, C), Out\}$  contains all **terminal histories** (the game ends after any of these histories). Note that  $(Out, C)$  is not a history, since the chain-store cannot cooperate once player  $k$  chooses  $Out$ .
- a function  $P$  (called the **player function**) that assigns to any nonterminal history a subset of  $N$ . If  $h \in H \setminus Z$  is a non-terminal history,  $P(h)$  is the player who takes an action after the history  $h$ ; e.g.  $P(In) = CS$
- for each player  $i \in N$  a preference relation  $\succsim_i$  on  $Z$  (the *preference relation* of player  $i$ ). As we have seen for the strategic games, the preference relation is often expressed as a utility function over the set of terminal histories.

Extensive form games are said to have **perfect information** when all players perfectly see all the actions taken by the previous players. If there is uncertainty about the actions previously taken, the game has **imperfect information**. A **strategy** in an extensive form game is a sequence of mappings, one for every node where the player has to take an action, describing the action to be taken by the player for every possible history that precedes that node.

In the Chain-Store game, both players have two strategies: player  $k$  moves only once (after the empty history,  $\emptyset$ ), and can choose between  $In$  or  $Out$ , player  $CS$  also moves only once (after the history

*In*) and can choose between  $F$  and  $C$ . The game has two Nash equilibria:  $(Out, F)$  and  $(In, C)$ . It is easy to verify that none of the players can gain by changing their strategy as long as the opponent plays the equilibrium. While you probably regard the second equilibrium as natural, the first one,  $(Out, F)$  may look strange. The chain-store is supposed to fight the competitor who never enters the market, something clearly forbidden by the rules of the game. The strategy  $(Out, F)$  is nevertheless a Nash equilibrium. The strategy  $F$  of the chain-store can be regarded as a commitment (or threat): given that  $CS$  will ever be allowed to move, it will fight. This makes it optimal for the competitor to stay out of the market.

There is, however, something wrong with the equilibrium  $(Out, F)$ : if the chain-store is actually allowed to move, fighting is not rational. With the competitor already on the market, the chain-store is clearly better by cooperating. The commitment of the chain-store to play  $F$  is not rational, and anticipating this, the competitor will enter the market.

The above intuition is captured by the notion of **subgame perfect equilibrium**, a refinement of the Nash equilibrium that requires strategies to be in equilibrium in *every* subgame of the original extensive form game. This requirement applies even for subgames that are never reached by the equilibrium strategy. Thus,  $(Out, F)$  is not a Nash equilibrium in the subgame that starts with the chain-store moving, so the only subgame perfect equilibrium of the Chain-Store game is  $(In, C)$ . Intuitively, subgame perfect equilibrium strategies do not contain threats that are not credible.

Extensive form games are important because they allow the representation of **repeated games**: i.e., players repeatedly engaging in a game  $G$ , called the **constituent game** or the **stage game**. There may be no limit to the number of times  $G$  is played, in which case, we have an **infinitely repeated game**. The players play the game  $G$  sequentially, such that the  $t^{th}$  repetition of the game (or the  $t^{th}$  round) starts after the end of the repetition  $t - 1$ . Players carry information from the outcome of the previous rounds, and can condition their strategy on such information.

The Chain-Store game can easily be viewed as a repeated game. Consider that there are  $K$  branches (in  $K$  different cities) of the chain-store, and that competitors decide sequentially whether or not to compete the local branch of the chain-store. Competitor  $k$  is assumed to see the outcome of the games played between all previous competitors and the chain-store. Although individual competitors play only once, the fact that there is information transmitted from previous competitors makes the whole interaction a  $K$  time repetition of the game in Figure 2.5, where the same chain-store (the *long-run* player) plays against a sequence of competitors (the *one-shot* players).

As we will see in the next sections, the transmission of information to the later rounds of a repeated game will have a major impact on the equilibrium strategies the players can adopt. Players may condition their strategies on the outcome of earlier games (and consequently on the past behavior of their opponent) which leads to new equilibria. What I have intuitively described in the beginning of this chapter as the influence of reputation information, will be formalized and characterized in the remaining of this section.

### The sanctioning role of reputation

As we have seen in the beginning of the section, reputation information can foster cooperation by allowing agents to effectively punish their partners who have cheated in the past. This mechanism works when agents care enough about the future, such that punishments inflicted at a later time can cause a big enough loss to offset the momentary gain one can obtain from cheating.

To give an example, let us first give a new interpretation of the Prisoners' Dilemma game from Figure 2.3. Instead of two prisoners, we have two agents facing a business opportunity. If both partners

		Player 2	
		<i>C</i>	<i>D</i>
Player 1	<i>C</i>	3, 3	0, 4
	<i>D</i>	4, 0	1, 1

**Figure 2.6:** The Prisoners' Dilemma game, without prisoners.

cooperate (denoted as *C*) each receives a return of 3. However, if one agent cheats (or defects, denoted as *D*) on his partner the cheater gets 4, while the cheated agent doesn't receive anything. When both agents cheat, each gets a default payment of 1. The outcome preferred by both agents is when they both cooperate (the prisoners similarly obtain the best deal when they cooperate with each other and do not confess). Nevertheless, agents are always better off by defecting, which makes (*D*, *D*) the (unique) dominant strategy of this game. The modified game is represented in Figure 2.6

Assume now that the same two agents repeatedly play this version of the Prisoners' Dilemma. In fact, they intend to play this game forever, although death, or other unforeseen events will eventually force them to stop playing. The interaction between them can be best modeled as an infinitely repeated game. The game, of course, will not be played an infinite number of times; however, as none of the agents deliberately foresees the end of the interaction, from their perspective, the game could go on forever.

The uncertainty about the future (which eventually will lead to the termination of the game) can be modeled in an infinitely repeated game by a discount factor applied to future revenues. The decisions made now by an agents takes into account the future revenues, but gains from the near future count more than gains from the distant future (exactly because gains from the distant future are less certain).

Formally, assume that both players expect with probability  $(1 - \delta)$  that the current game they are playing is the last one. For example,  $1 - \delta$  is the probability that one of the players will suffer a serious accident before the time the next game is scheduled, which prevents her from ever playing the game again. Therefore, following round  $t$ , the repeated game carries on to round  $t + 1$  with probability  $\delta$ , and stops with probability  $1 - \delta$ . If  $u_i(s_t)$  is the payoff to player  $i$  from the game in round  $t$  (where the strategy profile played by the players is  $s_t$ ), the lifetime revenue of player  $i$  is:

$$V_i = \sum_{t=0}^{\infty} \delta^t u_i(s_t);$$

The revenue from the current game ( $t = 0$ ) is certain, therefore is not discounted ( $\delta^0 = 1$ ); the next game will happen with probability  $\delta$ , so the revenue of the next game counts towards the *current time* estimation of life-time revenues with a discount ( $\delta^1$ ); The tenth repetition of the game will happen with probability  $\delta^9$ , and therefore, the corresponding revenue is discounted by  $\delta^9$ .

If every game gives player  $i$  the same payoff  $u_i(s_t) = u$ , the player's life time revenue is:

$$V_i = \sum_{t=0}^{\infty} \delta^t u = \frac{u}{1 - \delta};$$

For convenience, we will normalize life time revenues by multiplying them with  $(1 - \delta)$ ; in this way the life-time revenue is expressed in the same units as the stage-game payoffs, and can be easily compared to the latter.

Putting this background notation and formalism aside, we can now see how repeated games support new kind of equilibrium strategies, and new outcomes. Assume for example the following strategy: the



agent cooperates in the first round. In every subsequent game, the agent cooperates as long as the opponent cooperated in all previous games. Following a game where the opponent defected, the agent defects forever. This strategy is commonly referred to as the Tit-For-Tat (TFT) strategy because it involves immediate and definitive retaliation: once the opponent defects, a player retaliates by defecting in every subsequent round.

When both agents play TFT, the outcome of the repeated game consists of a sequence of cooperative outcomes, and both agents obtain the pareto-optimal payoff. Moreover, the pair of strategies (TFT, TFT) is an equilibrium, since none of the agents has an interest to deviate from the equilibrium strategy. To see this, imagine that one of the players, (player 1) considers deviating from TFT, by defecting in the current round. Player 1 can make the following reasoning: by defecting in the current round he gains 4 (because player 2 unsuspectingly cooperates), but switches the strategy of the opponent for the rest of the games. Following deviation, player 1 will obtain a payoff equal to 1 for every game starting with the next one. His life-time revenue is therefore:

$$V_1(D) = (1 - \delta) \left( 4 + \sum_{t=1}^{\infty} 1 \cdot \delta^t \right) = 4(1 - \delta) + \delta;$$

If player 1, on the other hand, sticks to TFT and cooperates, he obtains a payoff equal to 3 in the current round, and for every other round that follows. His life-time revenue is thus:

$$V_1(C) = (1 - \delta) \left( 3 + \sum_{t=1}^{\infty} 3 \cdot \delta^t \right) = 3;$$

TFT is an equilibrium strategy only when the deviation is not profitable: i.e.,  $V_1(D) < V_1(C)$ . Simple arithmetics shows that as long as  $\delta > 1/3$ , the payoff obtained for cooperation is always greater than the payoff obtained by defecting. In other words, as long as player 1 cares enough about future earnings (the discount factor is high enough) the momentary gain obtained by cheating is offset by the future losses.

Viewed from the perspective of reputation information, the Tit-For-Tat strategy uses binary reputation: good or bad. An agent has good reputation only if she had always cooperated in the past; otherwise the agent has bad reputation. The equilibrium rewards good reputation with cooperative play, and sanctions bad reputation with the uncooperative, lower payoff. When the future is important enough agents care to maintain their good reputation, and the equilibrium outcome is efficient. The TFT strategy is not the only one that leads to an efficient equilibrium of the Prisoners' Dilemma game. Axelrod (1984) analyses a series of strategies and evaluates them both empirically (based on simulation tournaments) and analytically.

More generally, the intuition that infinitely repeated games support equilibrium outcomes that are not possible in the stage game is formally characterized by the Folk Theorems. The good news is that when future payoffs are important enough (the discount factor  $\delta$  is close enough to 1), efficient outcomes can be sustained in an equilibrium of infinitely repeated games. The bad news, on the other hand, is that the set of equilibrium outcomes supported by equilibria of the repeated games is huge. As a matter of fact, almost all feasible, individually-rational outcomes<sup>3</sup> can be the result of an equilibrium strategy profile.

The core idea behind equilibrium strategies of the repeated game are future punishments triggered by present deviations from the equilibrium. As we have seen with Tit-For-Tat, players change their behavior when they detect deviations in the past behavior of their opponents. The sanctions for past

<sup>3</sup>An outcome is feasible if there is a pair of strategies (non necessarily equilibrium strategies) that generate it in the stage game. An outcome is individually rational if all players receive higher payoffs than what they can secure for themselves in the stage game. The minimum payoff a player can secure in the stage game is called the minimax payoff and corresponds to the worst-case payoff the other players can inflict on player  $i$ .

deviations is scheduled such that it offsets any gains an agent can obtain from deviating. Notable folk theorems for different settings have been proven by Fudenberg and Maskin (1989), Abreu, Pearce, and Stacchetti (1990) or Fudenberg, Levine, and Maskin (1994). For a detailed overview of the various Folk Theorems the reader is redirected to the work of Osborne and Rubinstein (1997) or Mailath and Samuelson (2006).

### The signaling role of reputation

While the infinite repetition of the Prisoners' Dilemma supports an equilibrium strategy profile that gives the players the efficient payoff, the same is not true for a finite repetition of the Prisoners' Dilemma. Assume two agents know they will be playing the Prisoners' Dilemma  $T$  times, where  $T$  is a finite number, no matter how large. At time  $t = 0$ , when the agents evaluate what strategy they should adopt, they will realize that at some point, they will end up playing the last game, at time  $t = T$ . Regardless of what has happened in the past, the best strategy in the last game is to defect: nothing follows after the last game, so agents maximize their payoff in the one-shot game. Since the outcome of the last round is fixed (both agents defect), the round before the last round will also be regarded as an isolated game: both agents defect in the last round anyway, so the outcome of the game  $T - 1$  does not have any influence on the future play. The only equilibrium, therefore, asks both agents to defect in round  $T - 1$  as well. The same reasoning would apply to rounds  $T - 2$ ,  $T - 3$ , etc., so by backward induction the only equilibrium asks both agents to defect in every round of the finitely repeated game (Nowak and Sigmund, 1994).

Although repeated defection is the only rational equilibrium, when  $T$  is large we would intuitively expect some cooperation to take place. Not surprisingly, controlled experiments actually demonstrate that humans do cooperate in finite repetitions of the Prisoners' Dilemma game (Lave, 1962; Roth, 1988; Wedekind and Milinski, 1996).

Another example where theoretical equilibrium predictions contradict intuitive considerations is the finitely repeated Chain-Store game (Figure 2.5). Here, again, the last repetition of the game must be played as an isolated game and generates the outcome  $(In, C)$  (i.e., to competitor enters the market and the chain-store cooperates), the only subgame perfect equilibrium outcome of the game. By backward induction all competitors will enter the market, and the chain-store cooperates with every one of them. The only subgame equilibrium of the finitely repeated Chain-Store games prescribes the outcome  $(In, C)$  in every game. This outcome, however, is not at all what you would expect from real chain-stores. In real markets, monopolists fight strongly with all competitors. They hope in this way to develop a reputation for being aggressive, and thus scare away future competition.

The apparent failure of game-theoretic equilibrium predictions have been elegantly solved by the three seminal papers of Kreps and Wilson (1982), Milgrom and Roberts (1982) and Kreps et al. (1982). They explain the apparent *irrational* behavior as a *signal* intended to alter the beliefs of the opponents. In particular, a player, let's call her Alice, would like her opponents to believe that she is a *crazy type* who does not play according to the equilibrium. When such beliefs become strong enough, opponents might as well give up playing the equilibrium, and start responding to what they believe Alice will play in the future. This change in the opponents' strategy might be beneficial for Alice, which justifies the signaling effort.

But why would such a change in beliefs be rational for the opponents? In the end, always cooperating with competitors is the *only* rational action a chain-store can take. The chain-store that does otherwise, is clearly intending to modify the beliefs of the competitors in order to deter their entry. Rational competitors should therefore refuse to modify their beliefs, knowing that it is just a trick of the chain-store.




Figure 2.7: The Chain-Store Game (modified).

Well, not necessarily. In reality, competitors might have real uncertainty regarding the game that is being played. Imagine, for example, that competitors are not exactly sure about the payoffs obtained by the chain-store. It is almost certain that fighting is costly, but maybe the cost is so small that it is largely compensated by the manager's thrill to engage the competitors. Competitors, therefore, might be in one of two possible situations:

1. they are playing against a normal chain-store that loses by fighting and obtains payoffs as described in Figure 2.5 (apriorily most likely), or,
2. they are playing against a chain-store headed by a strange manager who enjoys fighting and who obtains the payoffs as described in Figure 2.7 (unlikely, but possible).

In the former case, a chain-store fights competitors because it wants to alter the beliefs of future players, deter their entry, and thus gain more. In the latter case, however, the chain-store fights competitors because it is the optimal strategy following entrance, regardless of what happens in the future. Clearly competitors should commit to always enter the market when playing against the first type (the *normal* type) of chain-store, in order to determine the chain-store to cooperate with every entrant. Against the second type (the *crazy* type), however, competitors are better off by staying out.

Such a game where players are not perfectly informed about the structure of the payoffs is said to have **incomplete information**. Competitors do not know exactly against which type they are playing, so they form beliefs. The initial belief would normally assign high probability to the normal type, but following a series of fights, future competitors will update their beliefs in favor of the crazy type. When the probability of the crazy type becomes high enough, competitors cannot risk entering the market anymore, and they change their behavior.

It is exactly this uncertainty about the real type that allows chain-stores to obtain another equilibrium even in the finitely repeated Chain-Store game. Normal chain-stores (the ones that really lose by fighting) can start by fighting competitors, and thus falsely signal their belonging to the crazy type. They build a *reputation* for being crazy, in the hope that future competitors, at some point, will stop entering the market. Building this reputation is expensive, but if the repeated game is long enough (though still finite) the future revenues accessible when the reputation becomes credible compensates the cost of building the reputation.

This *reputation effect* can be extended to all games where a player ( $A$ ) could benefit from committing to a certain strategy  $\sigma$  that is not credible in a complete information game. In an incomplete information game where the commitment type has positive probability,  $A$ 's opponent ( $B$ ) can at some point become

convinced that  $A$  is playing as if she were the commitment type. At that point,  $B$  will play a best response against  $\sigma$ , which gives  $A$  the desired payoff. Establishing a reputation for the commitment strategy requires time and cost. When the higher future payoffs offset the cost of building reputation, the reputation effect prescribes minimum payoffs any equilibrium strategy should give to player  $A$  (otherwise,  $A$  can profitably deviate by playing as if she were a commitment type).

Fudenberg and Levine (1989) study the class of all repeated games in which a long-run player faces a sequence of single-shot opponents who can observe all previous games. If the long-run player is sufficiently patient and the single-shot players have a positive prior belief that the long-run player might be a commitment type, the authors derive a lower bound on the payoff received by the long-run player in any Nash equilibrium of the repeated game. This result holds for both finitely and infinitely repeated games, and is robust against further perturbations of the information structure (i.e., it is independent of what other types have positive probability).

Schmidt (1993) provides a generalization of the above result for the two long-run player case in a special class of games called of “conflicting interests”, when one of the players is sufficiently more patient than the opponent. A game is of conflicting interests when the commitment strategy of one player ( $A$ ) holds the opponent ( $B$ ) to his minimax payoff. The author derives an upper limit on the number of rounds  $B$  will not play a best response to  $A$ ’s commitment type, which in turn generates a lower bound on  $A$ ’s equilibrium payoff. This result holds even when players can observe only noisy signals of each other’s actions (Fudenberg and Levine, 1992).

Chan (2000), however, points out that reputation effects cannot appear when the two long-run players are equally (or approximately equally) patient. Unless the commitment strategy of player one gives player two her best feasible and individually rational payoff in the stage game<sup>4</sup> (i.e., player two *likes* the commitment strategy of player one), player one can build a reputation only too slowly to make it worth it.

Online settings, however, are often described by games of common interest, where the cooperation of the two players brings benefits to both of them. For such games, however, the reputation effect cannot be generalized. Schmidt (1993) and Cripps and Thomas (1997) show that in general, it is impossible to eliminate strategies that enforce smaller payoffs than the cooperative one. Nevertheless, Aumann and Sorin (1989) show that by restricting the set of possible types to those of bounded recall<sup>5</sup>, all pure strategy equilibria will be close to the cooperative outcome.

One last result we would like to mention in this section addresses the stability of reputation in settings with imperfect observations (i.e., the perception of players regarding the actions of their opponents are noisy). Cripps et al. (2004) surprisingly prove that in such noisy environments reputations are not permanent: in the long run, the players will learn each other’s type and play will revert to a static equilibrium.

### The Sanctioning vs. the Signaling Role

As mentioned in the beginning of this section, there is a fundamental difference between the two roles of reputation information. The sanctioning role allows agents and markets to specify strategies and behavior norms that threaten present misbehavior by future punishments. These threats make a priory commitments on cooperative behavior possible (even in settings where otherwise, cooperation is a dominated action) and thus contribute to efficient exchanges. As there are numerous ways of specifying

<sup>4</sup>Such games are called of *strictly conflicting interest* and are also treated by Cripps et al. (2005)

<sup>5</sup>A type is of bounded recall when the corresponding strategy prescribes actions that depend on the outcome of a finite number of past games. The strategy requires bounded memory since it does not need to store infinite information about the past.

punishments contingent of past behavior (or reputation), the set of possible equilibrium outcomes is large, as proven by the Folk Theorems.

The signaling role, on the other hand, allows agents to “learn” from past feedback the characteristics (or type) of other agents. Types can be innate, or acquired: e.g., the *normal* type chain-store has the incentive to play as if it were committed to fight all competitors. Every single report about an agent’s behavior triggers an update of beliefs regarding the agent’s type, and consequently an update of beliefs regarding how that agent will play in the future. The updated beliefs are later used to make more informative decisions or choices about future interactions.

## 2.3 Reputation Mechanisms for Online Systems

We define the reputation mechanism as the information system that allows to create and disseminate reputation information. Its main role is to gather feedback from the agents, to aggregate it into reputation information, and to disseminate this information afterwards.

Existing reputation mechanisms differ on several dimensions. The first is the format of solicited feedback. Some mechanisms require binary feedback, and agents may “vote” positively or negatively a product, a service, or the behavior of another agent. Other mechanisms allow agents to submit one out of a finite number of possible feedback values. Hotels, for example, are rated with stars (one star for the worst hotel, five stars for the best one), and movies are rated with points (one point is the minimum rating, ten points is the maximum one). Yet other mechanisms allow (almost) an infinity of possible feedback values, usually expressed as free textual comments<sup>6</sup>. Combinations of the above are also quite common: e.g., users may give both numerical ratings, and textual comments.

The second dimension which differentiates reputation mechanisms is the way they aggregate feedback reports into reputation information. Some mechanisms accumulate all reports into a reputation “score” that may potentially grow forever. Other mechanisms use a Bayesian approach where every report is used to update a belief regarding the reputation of an agent. Yet other mechanisms use discrete values for reputation information (e.g., an agent may be *very trustworthy*, *trustworthy*, *untrustworthy* or *very untrustworthy*) with clear rules describing how sets of feedback are mapped to reputation values.

Finally, the third dimension relates to how reputation information is disseminated in the community. Some mechanisms rely on a central entity to gather all reports and compute reputation information. The central component is later responsible for disseminating reputation information (e.g., by publishing on a web site). Other mechanisms rely on decentralized solutions where every peer stores information about the other agents with which he interacted. Reputation information is disseminated on demand, through the social network capturing the trust relations between peers.

In the rest of this section I will review existing reputation mechanisms and explain what is the context they are used in, what kind of feedback they collect, how they aggregate reputation information, and how that information is later disseminated to the other agents. I will first review commercial mechanisms that are in use by real communities, followed by a discussion of the mechanisms proposed by the academic community.

---

<sup>6</sup>Strictly speaking, the set of possible feedback values is not infinite. Most mechanisms restrict the length of the textual comment, which limits the number of possible feedback values.

### 2.3.1 Online Implementations

Perhaps the most famous implementation of a reputation mechanism is the Feedback Forum of eBay. The main goal of the reputation mechanism is to promote trust between sellers and buyers that are geographically far apart, and to encourage trading in the market. The reputation mechanism is often cited as one of the important factors that contributed to the incredible growth of the eBay market.

For every transaction, eBay encourages both participants (the seller and the buyer) to rate each other. Feedback can be positive (1), negative (-1) or neutral (0), followed by a comment that contains free text. As of May 2007 buyers may also rate:

- the degree to which the description offered by the seller matched the item;
- the quality of the communication with the seller;
- the shipping time;
- the shipping and handling charges offered by the seller.

For these supplementary dimensions feedback is from one to five stars, one star being the lowest, and five stars being the highest rating. Both users may reply with a comment to the feedback they received, and may also add a follow-up note (free text) to the feedback they have previously submitted.

eBay aggregates all feedback reports into a *Feedback Score*. A positive rating increases the Feedback Score by one point, a negative rating decreases the Feedback Score by one point, while a neutral rating leaves the Feedback Score unmodified. Each member, however, can only affect another member's Feedback Score by no more than one point. The Feedback Score is displayed as part of a *Feedback Profile*, together with (Figure 2.8):

- the percentage of positive feedback received by the user,
- the number of unique users who left positive feedback,
- the number of unique users who left negative feedback,
- the total number of positive feedback received for all transactions,
- a profile of most recent reports (summarized for the last month, the last six months and for the last year),
- the average of the ratings for the additional dimensions,
- a list of the textual comments accompanying every feedback received by the user. The list can be sorted to display the comments received as a seller and the ones received as a buyer,
- a list of the feedback left by that user.

Similar mechanisms are used by most of the marketplaces where complete strangers can interact. In Amazon's secondary market, users may rate other users who sell second-hand books. Ratings are from 1 to 5, (where 4 and 5 are considered positive ratings, 3 is a neutral rating, 1 and 2 are negative ratings) followed by one line of comments. The "reputation" of a seller is displayed as the percentage of positive, neutral and negative ratings, followed by the number of raters, and the list of comments.

eLance.com is a marketplace for professional services. Contractors rate their satisfaction with subcontractors by a rating from 1 to 5, and a short comment. The reputation of a user is displayed as

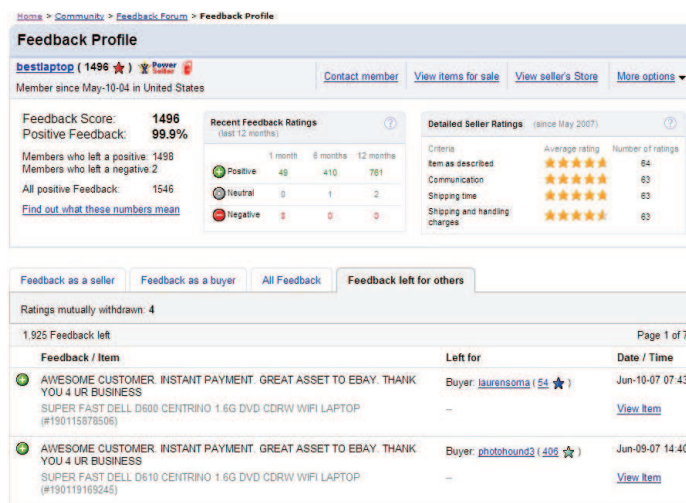


Figure 2.8: The eBay Feedback Profile.

the percentage of positive reviews, the total earnings made by the user, followed by the project history of the user. RentACoder.com, a marketplace for software programming services employs a similar mechanism.

A second family of reputation mechanisms control the quality of content (mostly information) contributed by a community of users. Slashdot is an online discussion board where postings (short news articles) are prioritized or filtered according to ratings received from the readers. The site uses a 2-layer semi-automatic moderation scheme. Level 1 moderators rate posts, while level 2 moderators (or the meta-moderators) moderate the level 1 moderators. Moderators are selected randomly according to their “karma”. Karma may be *terrible*, *bad*, *neutral*, *positive*, *good*, and *excellent* depending on how well the contribution of a user has been moderated in the past. Moderators are given a number of points (proportional to their karma) which limits the number of comments they may rate. Every comment may be rated on a scale from -1 to 5 by choosing labels that best fit the content. Example of labels are *insightful*, *funny*, *troll*. Meta-moderators judge the fairness of the level 1 moderators, which influences the karma of level 1 moderators and hence the ability to become moderators in the future. One particularity of Slashdot is that karma is not explicitly visible. Nevertheless, users with better karma get more power in the system and receive higher visibility for the content they submit.

Digg.com is a news site where the best articles are promoted on the top of the list by the community of readers. Unlike in Slashdot, any reader can vote for an article and thus “digg it” up. Users may also “burry” (i.e. remove spam) and comment on stories. All activities made by a user are recorded in a profile visible to the user’s friends.

Yahoo! Answers, and Live QnA allow users to post questions and receive answers from other users. The entire community may vote for the answers already posted, and thus contributes to selecting the best answer to a specific question. Users are rewarded for their effort (answering questions or voting) by points that can later be used to ask more questions. The number of points gained by a user measures the user’s reputation within the community. However, this reputation is not visible to the outside world, and is only used internally by the site owners.

A third family of mechanisms keeps reputation information about products or services. ePinions.com allows users to rate goods products and services on a scale from 1 to 5 for various product-dependant

aspects (e.g., ease of use, reliability, durability, etc). Ratings are also accompanied by a detailed textual review. The reputation of a product is displayed by aggregating the numerical ratings across the quality dimensions for that product. Moreover, readers have easy access to the list of detailed textual reviews, ordered by the date of submission.

Interesting enough, Epinions also keeps an internal reputation mechanism for the reviewers. Members can rate the reviews as helpful or not. Reviewers with a lot of activity over a period of time may become *advisors, top reviewers, or category leads*. Advisors and Top Reviewers are automatically chosen. Category leads are chosen by the company based on member nominations. Epinions also creates a web of trust where members can decide to “trust” or “block” other members. The fact of being trusted or blocked by other members impacts a member’s qualification to become a Top Reviewer.

Similar product review systems may be seen on Bizrate.com, Amazon.com, IMDB.com or TripAdvisor.com. As a matter of fact, all major retail e-commerce sites allow users to post reviews about the products they have bought. The reputation of a product is usually displayed as the average of all numerical ratings, plus a list of the textual comments submitted by the reviewers.

Most reputation mechanisms implemented today support with preponderance the signaling role of reputation. Feedback from the past is aggregated into “different scores” that are displayed as such, and called *reputation*. Users are not offered any support for interpreting reputation information, and they generally regard it as *estimates* about future behavior.

Empirical studies, nevertheless, have shown the existence of *reputation premiums*, where users (or products) with higher reputation obtain higher revenues. This suggests an implicit sanctioning role of reputation information, as human users naturally discount the price (or the amount of business) offered to peers with lesser reputation. This sanctioning, however, is mostly ad-hoc, without any support from the mechanism. The absence of clear guidelines for interpreting reputation information, is, I believe, one of the major drawbacks of existing reputation mechanisms.

### 2.3.2 Academic Models

The reputation mechanisms modeled in research papers are also mostly ad-hoc in the way they interpret reputation information. They differ, nevertheless, from the existing commercial implementations in the fact that most of them do not rely on the existence of a central place to gather feedback and aggregate reputation information.

The view these works assume is that of a social network, where two agents that have interacted in the past share an edge in the network. Every edges  $(A, B)$  is assigned a weight representing the “trust” of agent  $A$  towards agent  $B$  aggregated across all interactions between them in which agent  $A$  happened to have relied on agent  $B$ . Having the local interactions among the agents encoded this way, the challenge is how to merge these local information to enable the agents to compute the reputation of non-neighboring agents, whom they never met before. The main distinguishing points among the numerous works belonging to this class are:

1. the strategy to aggregate individual experiences to give the mentioned weights: i.e., the trust an agent  $A$  has in agent  $B$ ;
2. the strategy to aggregate the weights along a path of an arbitrary length to give a path wide gossip between two non-connected agents;
3. the strategy to aggregate this gossip across multiple paths between two non-connected agents.



Beth et al. (1994) presents an early example in which a clear distinction between direct experiences and recommendations has been made, which is reflected in the strategy for path wide gossip aggregation. However, this separation of the two contexts led to an exponential complexity of the trust derivation algorithm. Clearly, this is unacceptable for large scale networks.

Yu and Singh (2000) do not treat recommendations and direct service provisions separately. The authors use a variation of the delta learning method to aggregate “positive” and “negative” experiences of the agents into the weights assigned to the corresponding branches and simple multiplication as the strategy to compute the path wide gossips. As for the strategy to aggregate the gossips of different paths the authors use a variation of the simple maximum function. All this results in a polynomial time algorithm for the overall trust aggregation.

Richardson et al. (2003) offers important theoretical insights on how the computational complexity of the trust derivation algorithms relates to the mentioned aggregation strategies by characterizing the combinations of path and across-path aggregation strategies that may lead to a non-exponential trust computation algorithm (we note that many other works use such combinations: e.g., (Page et al., 1998) and (Kamvar et al., 2003)). The authors also offer such an algorithm which is, however, based on a synchronous participation of all agents in the network. As such it is not quite appropriate for usage in P2P networks due to their inherent high dynamicity. With respect to this problem Xiong and Liu (2004) offer a considerable improvement in terms of an appropriate caching scheme that enables asynchronous computation while retaining good performance.

Srivatsa et al. (2005) describe techniques for minimizing the impact of malicious nodes on the reputation mechanism. The framework guards against strategic oscillation of behavior, detects fake transaction reports, and filters out dishonest reports.

A common denominator of all these works is that the computed values have unclear semantics and are hard to interpret on an absolute scale, without ranking them. In many applications this imposes certain problems. On the other hand, as shown by many simulations, they are very robust to a wide range of misbehavior.

Probabilistic estimation techniques present certain improvement with respect to the meaningfulness of the computed values. Namely, they output probability distributions (or at least the most likely outcome) over the set of possible behaviors of the trusted agents enabling thus the trusting agents to evaluate explicitly their utilities from the decision to trust or not. Mui et al. (2002b) present the well-known method of Bayesian estimation as the right probabilistic tool for assessing the future trusting performance based on past interactions. Only direct interactions were studied - the question of including recommendations was not considered. Buchegger and Le Boudec (2003) go a step further by taking into account the “second-hand” opinions also. However, the strategy for merging own experiences with those of other witnesses is intuitive (giving more weight to own experiences, though plausible, is still intuitive) rather than theoretically founded. Aberer and Despotovic (2001) and Despotovic and Aberer (2004) is another work belonging to this group.

As we said previously, many works can be regarded as belonging to the class of “social networks” and using different strategies to aggregate the gossip available. Thus Dellarocas (2000), Dellarocas (2004), Zacharia et al. (1999) and Abdul-Rahman and Hailes (2000) use collaborative filtering techniques to calculate personalized reputation estimates of as weighted averages of past ratings in which weights are proportional to the similarity between the agent who computes the estimate and the raters. Birk (2000), Birk (2001), Biswas et al. (2000) and Witkowski et al. (2001) use either well-known machine learning techniques or heuristic methods to increase the global performance of the system by recognizing and isolating defective agents. Common to these works is that they consider only direct reputation. Similar techniques, extended to take into account indirect reputation, are used by Barber and Kim (2001), Schillo et al. (2000), Jurca and Faltings (2003), Yu and Singh (2002) and Sen and Sajja (2002).

### 2.3.3 Empirical studies of Reputation Mechanisms

Empirical studies on reputation were conducted both in labs and in real settings. The effects of reputation can better be isolated in lab experiments. Keser (2003) studies the *investment game* where player one may entrust the opponent with an *investment* that brings sure returns (e.g., player two receives three times the amount invested by player one). Player two then decides how much to return to player one. Trust increases the total revenues (the higher the investment, the higher the returns) but leaves player one vulnerable to player two taking all the profit. When players receive information about each other's play both trust (investment) and trustworthiness (return of profits to the trustor) were higher. Similar studies of the investment game have been conducted by Berg et al. (1995), McCabe et al. (2003), Cox and Deck (2005) Fehr and Gächter (2000) and Buskens and Barrera (2005).

Bolton et al. (2004) investigate a similar two-stage game where buyers decide whether to pay, and sellers then decide whether to ship the item. They also consider two settings: one where players are matched randomly and have information regarding the history of play of their partners, one where players are matched randomly without reputation information, and one where the same two partners interact repeatedly. Availability of reputation information (history of past play) increases both trust and trustworthiness with respect to the no reputation setting. Nevertheless, the repetition of the game between the same players yields the best results.

Chen et al. (2004) investigate the same setting for a wider set of choices. For example, buyers can choose the sellers they want to trade with, both buyers and sellers may choose to cheat on their partners, and players may choose to misreport feedback. Despite a more general setting, the results are consistent with those of Bolton et al. (2004): reputation information significantly increases the efficiency of the market. Interestingly, the fact that players may misreport feedback did not have a major impact on the performance of the market.

An impressive body of lab experiments conducted on the Prisoners Dilemma are discussed by Rapoport and Chammah (1965). The authors asked pairs of students to play the Prisoners Dilemma several hundred times. The statistics yielded by the experiments give valuable insights to how human users act and reason in conflict scenarios.

With few notable exceptions (e.g., Diekmann and Wyder, 2002), the vast majority of field experiments on trust and reputation were conducted on eBay's reputation mechanism. An overview of the different studies is described by Dellarocas (2005), Resnick et al. (2006) and Bajari and Hortacsu (2004) with the following main conclusions:

- reputable sellers seem to have their items purchased with higher probability and for higher prices. The precise effects are, however, ambiguous and studies on different categories of products often lead to contradictory results;
- reputation is more important for riskier (high value) transactions;
- from all the information displayed in the feedback profile, buyers mostly look the total number of positive (respectively negative) reports, followed by the number of most recent negative comments.

Resnick and Zeckhauser (2002) identify main patterns of human behavior with respect to trust. The authors argue that, despite clear incentives to free ride (not leave feedback) and leave only positive feedback, trust among eBay traders emerges due to its reputation system.

Resnick et al. (2006) conducted the first randomized controlled field experiment of the eBay reputation mechanism. The same goods (vintage postcards) were sold under two identities: that of a new

seller and that of a highly reputable seller. As predicted, the seller with a good reputation did significantly better, and obtained, on the average, 8.1% higher prices than the new seller. The same authors conducted another experiment and conclude that few negative ratings do not drive away buyers.

The same effect of reputation on eBay prices was observed by:

- Houser and Wooders (2006) for the sale of Pentium processors,
- Ba and Pavlou (2002) for the sale of music, software and electronics,
- Kalyanam and McIntyre (2001) for the sale of Palm Pilot PDAs,
- McDonald and Slawson Jr. (2002) for the sale of dolls,
- Melnik and Alm (2002) for the sale of circulated coins,
- Dewan and Hsu (2004) for the sale of collectible stamps,
- Dewally and Ederington (2006) for the sale of collectible comic books.

Other studies only find a correlation between reputation and the probability of sale. Livingston (2002), for example, studies the auction of golf clubs and concludes that more positive feedback increases the probability of sale. Jin and Kato (2004) reach the same conclusion for the sale of sports trading cards. Similar conclusions hold for the German eBay market, as pointed out by Wehrli (2005).

Yet other studies (Lucking-Reiley et al., 2000; Kauffman, R. J. and Wood, C., 2000; Eaton, 2002) don't find an influence of positive feedback on the price or probability of sale, however, do acknowledge that negative feedback drives away buyers and consequently reduces the probability of sale.

One surprising fact about eBay's reputation mechanism is the overwhelming percentage of positive reports: only 1% of all reports submitted to eBay are negative. This might lead to the naive interpretation that 99% of the transactions on eBay result in satisfactory outcomes. Dellarocas and Wood (2006) argue that this interpretation is biased and does not reflect the reality. They propose a method to account for the *sound of silence* and factor in the transactions on which no feedback was submitted. The results are surprising as they suggest that approximately 20% of the eBay users are not entirely satisfied with the outcome of the transactions. Deriving information from missing feedback enables buyers to more accurately assess the risks involved by a transaction. The bias of eBay feedback (mainly due to retaliation effects) was also mentioned by Reichling (2004) and Klein, Lambertz, Spagnolo, and Stahl (2006).

Several recent results investigate the role of reputation in the context of product review forums. Chevalier and Mayzlin (2006) examine the impact of reviews on the sale of books on Amazon and BarnesandNoble.com. They find a positive correlation between an improvement in a book's review and volume of sales of that book. They also note that book reviews are overwhelmingly positive, and therefore, the few negative reviews have a larger impact.

Hennig-Thurau et al. (2004) conduct a survey to determine what motivates users to submit online reviews. They identify four motives: social benefits (fun and pleasure resulting from the participation in an online forum), economic incentives (rewards offered in exchange of the review), concern for other consumers (help others take more informative decisions) and extraversion (pleasure for sharing one's opinion with the rest of the community).

### 2.3.4 Other Aspects related to Reputation Mechanisms

Economic theory predicts that voluntary feedback will be underprovisioned. Since feedback constitutes a public good, users might be tempted to free-ride. Avery, Resnick, and Zeckhauser (1999) analyze mechanisms where early evaluators are paid to provide information, and later clients pay in order to balance the budget. They conclude that voluntary participation, no price discrimination and budget balance cannot be simultaneously achieved; however, any combination of two such properties is feasible.

Another problem associated to online reputation mechanism comes from the ease with which participants can change their identity. Friedman and Resnick (2001) discuss the risks associated with cheap *online pseudonyms*, where users can exploit their initial reputation and then start new under a new identity. Using a game theoretic model, the authors conclude that newcomers must start with the lowest possible reputation. The cost of building reputation exceeds the benefits obtained from “milking” it afterwards, which gives incentives to the users to keep their identity. This property is later used by Dellarocas (2005) to design moral hazard reputation mechanisms that are robust to identity changes.

Dellarocas (2000) identifies a number of possible attacks on reputation reporting systems: e.g., *ballot stuffing* (i.e., artificially inflating someone’s reputation), *bad mouthing* (i.e., artificially denigrating an agent), *negative discrimination* (i.e., providers deliver good service to all except a few clients) and *positive discrimination* (i.e., providers deliver good service only to a few selected agents). Dellarocas discusses several solutions to reduce the effects of those attacks.

The strategic manipulation of opinion forums has been more formally studied by Dellarocas (2006b). Firms whose products are being discussed in such forums can try to manipulate consumer perceptions by posting costly anonymous messages that praise their products. The striking result of this result is that manipulation can both increase and decrease the information value of online forums. A theoretical analysis backed up by examples describes the settings in which such phenomena occur.

Reputation mechanisms would not be needed in the first place if the transactions in the electronic market can be designed to be *safe*. Sandholm and Lesser (1995) propose the idea of *safe exchanges* where the contract between the buyer and the seller is fulfilled without enforcement. The exchange is managed such that at any point in time, the future gains from carrying out the rest of the exchange are larger than the gains from cheating for both the seller and the buyer. The underlying principle is to split the exchange into smaller chunks that are intelligently managed to avoid defection. Sandholm and Ferrandon (2000) operationalize this idea into a safe exchange planner that can automatically carry out the transaction on behalf of a player. Buttyan and Hubaux (1999) theoretically model the safe exchange by a dynamic game, and Sandholm and Wang (2002) characterize what is inherently possible and impossible to achieve in safe exchange. Safe exchanges, however, can only be designed for divisible goods.

## Chapter 3

# Truthful Signaling Reputation Mechanisms

As discussed in Section 2.2, one of the two important roles of reputation information is to bridge the asymmetry of information that usually exists in online trading situations, by allowing the buyers to differentiate between good and bad providers. The reputation mechanisms that primarily support this role of reputation information are called *signaling* reputation mechanisms, as they signal to the buyers some a priori unknown characteristics of the seller.

Most online feedback mechanisms available today on the internet are primarily signaling reputation mechanisms. For example, product review sites like, Amazon.com, ePinions.com or BizRate.com fit into this category; their main goal is to disclose to future buyers hidden, *experience-related* (Parasuraman et al., 1985) product attributes (e.g., quality, reliability, ease of use, etc.) gathered from the testimonies of previous buyers. Having access to this previously unavailable information allows buyers to take better decisions and buy the products that best satisfy their needs.

From this perspective, signaling reputation mechanisms do not just place white or black tags on the products in a certain market. Rather they help buyers identify product “*types*”, where every type is characterized by a set of properties relevant to the market. Note that most of the times there is no complete order relation on the set of possible types a product could have: Type *A* is not necessarily better than type *B*, it is just different. It is therefore up to the buyer to select the product type that is best in a particular context.

As a very simple example, assume that the relevant properties defining the product type are (1) a one dimensional measure of quality, and (2) the speed of delivery. A good quality product that is delivered fast is definitely better than a bad quality product delivered after a long time. Nevertheless, the choice between a very good product with slow delivery (type *A*), and another slightly worse product with fast delivery (type *B*) is harder to make. Depending on the context and the subjective preferences, one buyer will choose the product type *A*, while another buyer will choose the product type *B*. When more attributes are taken into account for defining a type, it becomes more likely that different types will offer different tradeoffs to the buyers.

Signaling reputation mechanisms are therefore valuable for the sellers as well. First, they allow the existence of specialized products in the market that will be recognized by the corresponding niche. This can greatly increase the value of the market, as niche products cover the long tail distribution of preferences and account for a significant volume of trading (Anderson, 2006). Second, sellers can

optimally plan the investment in their products, knowing that buyers will be able to differentiate the tradeoffs between different product attributes. The reputation mechanism thus contributes to the overall efficiency of the market.

From a technical point of view, the reputation of a product indicates the product's type, as inferred by the reputation mechanism from the feedback reported by previous users. In most cases, however, the reputation mechanism can only make an approximate prediction of a product's type, and therefore, reputation information is represented through a probability distribution over the possible types. As more and more information is gathered about the product, the probability distribution (i.e., the reputation information) will concentrate on the real type.

The signaling reputation mechanism therefore implements a learning algorithm. Given a description of all possible types, the problem of the reputation mechanism becomes to continuously update reputation information (i.e., the probability distribution over the types) such that the feedback submitted by the users is optimally taken into account. As we will see in the next section, one of the most natural solutions to this problem is given by the Bayesian learning theory.

Before proceeding to a more formal treatment of signaling reputation mechanisms, let us make the following remark. So far, I have mainly talked about the reputation and the type of a product. The same discussion can be easily extended to sellers, service providers, algorithms, information or other entities on which reputation information may be useful. Nevertheless, signaling reputation mechanisms work under several important assumptions. First, the set of possible types, and their exact description, is assumed known. Second, entities are assumed to have only one type, which does not change in time. Finally, the entity is assumed to act consistently according to its type, such that all agents interacting with that entity see (a noisy representation) of exactly the same type.

### 3.1 A Formal Model

I model the market and a signaling reputation mechanism in the following way. From all products and buyers, I isolate all buyers (or "agents", assumed rational) who buy and experience one particular product. The quality of the product remains fixed, and defines the product's unknown *type*. As discussed in the previous section, types may be characterized by different quality attributes, and every type is defined by a specific value for each quality attribute.  $\Theta$  is the finite set of possible types, assumed known by all agents, and  $\theta$  denotes a member of this set.

A central reputation mechanism publishes the reputation of the product as a probability distribution over the set of types. I denote by  $Pr[\theta]$  the probability that the product has the type  $\theta$ , hence the reputation information published by the reputation mechanism is a vector  $(Pr[\theta])_{\theta \in \Theta}$ , where  $\sum_{\theta \in \Theta} Pr[\theta] = 1$ . All buyers are assumed to know the reputation of the product, so they share a common belief regarding the prior probability  $Pr[\theta]$ .

After the purchase, the buyer inspects the product and forms an opinion about its quality and true type. I model the opinion of the buyer  $i$  as a noisy random signal  $O_i$ ; the value of signal  $O_i$  is denoted  $o_i$ , belonging to the set  $Q = \{q_0, q_1, \dots, q_{M-1}\}$  of possible signal values. Every signal  $q_j$  may actually represent a set of different attribute values characterizing the quality of the product. For example, the signal  $q_0$  could mean that the product is unreliable and has a long delivery time. The set  $Q$  and the semantics of the signals  $q_j$  is assumed known by all agents.

Different buyers may perceive different signals about the same product, either because they have different preferences, or because of the inherent noise in the assessment process. However, I assume that the observations of different buyers are conditionally independent, given the type of the product.

Let  $Pr[q_j|\theta] = Pr[O_i = q_j|\theta]$  be the probability that a buyer observes the signal  $q_j$  when the true type of the product is  $\theta$ .  $Pr[\cdot]$  is assumed common knowledge, and  $\sum_{q_j \in Q} Pr[q_j|\theta] = 1$  for all  $\theta \in \Theta$ . A summary of the formal notation is given in the Appendix 3.A.

For the sake of simplicity, some results in this chapter are obtained on a binary setting where there are only two possible quality signals. Here, buyers are assumed to form and express a binary opinion about the quality of the product, so that they will regard the product as either *good* or *bad*. The resulting feedback they can submit to the reputation mechanism is binary, where  $q_1$  (or simply 1) stands the good quality signal, while  $q_0$  (or simply 0) stands for bad quality. When referring to the submitted report, I will also call the two feedback values as the *positive*, respectively the *negative* report. To distinguish between the general and the binary feedback set, I will also use the notation  $Q_2 = \{q_0, q_1\} = \{0, 1\}$ .

The reputation mechanism asks every buyer to submit feedback. Assuming that buyers report truthfully, every feedback report can be used to update the reputation of the product, such that the updated reputation becomes the posterior distribution over the possible types, given the prior distribution and the observed signal. Therefore, the updated reputation can be computed by Bayes's Law, where the posterior probability that the product is of type  $\theta$  given the feedback  $q_j \in Q$  is:

$$Pr[\theta|q_j] = \frac{Pr[q_j|\theta]Pr[\theta]}{Pr[q_j]} \quad (3.1)$$

where  $Pr[q_j]$  is the overall probability that a buyer observes the signal  $q_j \in Q$ :

$$Pr[q_j] = \sum_{\theta \in \Theta} Pr[\theta]Pr[q_j|\theta];$$

Note that a signaling reputation mechanism represents a very simple instance of an *online* mechanism<sup>1</sup>. Feedback is received sequentially, and after every submitted report, the mechanism updates the probability distribution over the possible types. However, since all buyers are assumed to experience a signal from the same product (whose type remains fixed) the optimal solution to this online process is governed by Bayes' Law (expressed in Eq. 3.1).

The assumption that the type does not change in time can be relaxed such that the type changes sufficiently slowly. If the model of type changes is known (e.g., after every interaction, the type of the product changes randomly to some other type with a very small probability) it can be factored into the Bayesian updating process to reflect the true posterior beliefs. If, on the other hand, the model of change is unknown, we can have a mechanism that works on batches of reports: the type is assumed to change only once every  $T$  transactions. During every window of  $T$  feedback reports, the mechanism functions as described by Eq. 3.1. Nevertheless, from one batch to the next, the mechanism assumes that the type may have changed, and relaxes the prior beliefs in order to reflect this possible change.

The computation of reputation information is therefore extremely simple as long as the mechanism receives truthful feedback. The remaining of this chapter will therefore be dedicated to this problem: how can the reputation mechanism provide incentives for rational agents to report honestly? Section 3.2.1 reviews the work of Miller et al. (2005) and explains that payment mechanisms (i.e., the reputation mechanism pays agents for reporting feedback) can in general create an equilibrium where rational agents find it optimal to report the truth, a property called incentive-compatibility. The key idea is to score a report against the report submitted by another agent, called the reference report. The payment to the first reporter depends on this score, and maximizes the expected revenue of the reporter when submitting true feedback. Section 3.3 uses the idea of *automated mechanism design* (Conitzer and

<sup>1</sup>Online mechanisms function in dynamic environments where agents may arrive and depart dynamically, and information can change over time. Online mechanism design (Friedman and Parkes, 2003) addresses the design of mechanisms for such dynamic environments. A survey of the main results is offered by Parkes (2007).

Sandholm, 2002; Sandholm, 2003) to construct incentive-compatible payment schemes that are also optimal: i.e., the expected cost to the reputation mechanism is minimized, while offsetting both the cost of reporting and the external gains an agent could obtain from lying. The average saving to the reputation mechanism is around 50% when compared to the traditional payments based on proper scoring rules (Miller et al., 2005).

Section 3.4 investigates two methods that can further decrease the incentive-compatible feedback payments. The first requires the use of several reference reports and guarantees that the cost to the reputation mechanism decreases as payments are computed by considering and increasing number of reference reports. The second method involves probabilistic filtering mechanisms that discard some feedback reports. If the payments and the filtering mechanism are designed together, significant reductions of cost can be achieved without affecting the quality of reputation information.

Section 3.5 extends the design of incentive-compatible payments to settings with uncertain information. The improved payments make honesty the optimal strategy even for agents who poses private information unknown to the reputation mechanism.

Finally, 3.6 addresses the collusion between reporters. First, I show that general incentive-compatible payments have several equilibria besides the truthful one. Some of the lying equilibria generate higher payoffs than the honest equilibrium, which might motivate selfish agents to coordinate their reports and game the mechanism. Fortunately, supplementary constraints can be added to the design problem such that honest reporting becomes the unique, or the pareto-optimal equilibrium. For different collusion scenarios, I describe algorithms for computing such incentive-compatible, collusion-resistant payments.

## 3.2 Incentives for Honestly Reporting Feedback

Recent studies raise important questions regarding the quality of reputation information as reflected by contemporary online feedback mechanisms. First, the absence of clear incentives drives only some of the users to voice their opinions and report feedback. For example, Hu, Pavlou, and Zhang (2006) and Admati and Pfleiderer (2000) show that Amazon<sup>2</sup> ratings of books or CDs follow with great probability bi-modal, U-shaped distributions where most of the ratings are either very good, or very bad. As controlled experiments on the same items reveal normally distributed opinions, the authors conclude that users with a moderate outlook are unlikely to report. Talwar, Jurca, and Faltings (2007) identify another factor that promotes rating, namely the desire to contribute with something new to the previously submitted reports. In both cases, the reputation mechanism collects an unrepresentative sample of reviews, that is not necessarily informative for the average user.

Second, and even more distressful, some users intentionally lie in order to gain external benefits from the distorted reputation. Harmon (2004) reports that some authors write fake reviews on Amazon in order to boost the sale of their own books, or to trash the reputation of competing titles. White (1999) describes manipulation techniques for pushing songs up the charts, and Elliott (2006) and Keates (2007) identify problems associated to fake hotel reviews on the travel reputation site TripAdvisor.com. Although we still see high levels of altruistic (i.e., honest) reporting, the increasing awareness that gains can be made by manipulating online reputation will likely attract more strategic reporting in the future.

Both problems can be solved by explicitly rewarding users for reporting feedback. These *payments*<sup>3</sup> made by the reputation mechanism to the reporters have two roles. First, they must cover the cost of reporting feedback so that more users report, and give the reputation mechanism a more representative

---

<sup>2</sup><http://www.amazon.com>

<sup>3</sup>Using the term payments does not exclude non-monetary rewards such as preferential access to resources, social status or bonus “points”.



collection of feedback. Second, the payments must be designed such that self-interested agents find it in their best interest to report the truth.

Fundamental results in the mechanism design literature (d'Aspremont and Grand-Varet, 1979; Crémer and McLean, 1985) show that side payments can be designed to create the incentive for agents to reveal their private opinions truthfully. Such payment schemes have been constructed based on *proper scoring rules* (Kandori and Matsushima, 1998; Johnson et al., 1990; Clemen, 2002), and exploit the correlation between the observations of different buyers about the same good. The first adaptation of these results to online feedback mechanisms is due to Miller, Resnick, and Zeckhauser (2005), and is described below.

### 3.2.1 Incentive-compatible Payment Mechanisms

The model used by Miller et al. (2005) is similar to the one presented in Section 3.1: a set of users experience the same product or service, and later report the privately perceived quality signal to a central reputation mechanism.

The reputation mechanism *scores* every submitted feedback by comparing it with another report (called the *reference* report) submitted by a different user about the same good. The payment received by a reporter is then directly proportional to the computed score. Let  $r_i \in Q$  be the report submitted by the buyer  $i$ , let  $ref(i)$  be the reference reporter of  $i$ , and let  $r_{ref(i)} \in Q$  be the report submitted by the reference reporter of buyer  $i$ . The payment received by the buyer  $i$  is denoted by  $\tau(r_i, r_{ref(i)})$ , computed in the following way:

$$\tau(r_i, r_{ref(i)}) = ScR(r_{ref(i)}|r_i); \quad (3.2)$$

where  $ScR(\cdot|\cdot)$  is a proper scoring rule. The three best studied proper scoring rules are:

- the logarithmic scoring rule:

$$ScR_{log}(q_k|q_j) = \ln(Pr[q_k|q_j]); \quad (3.3)$$

- the spherical scoring rule:

$$ScR_{sph}(q_k|q_j) = \frac{Pr[q_k|q_j]}{\sqrt{\sum_{q_h \in Q} Pr[q_h|q_j]^2}}; \quad (3.4)$$

- the quadratic scoring rule:

$$ScR_{quad}(q_k|q_j) = 2Pr[q_k|q_j] - \sum_{q_h \in Q} Pr[q_h|q_j]^2; \quad (3.5)$$

and depend on the conditional probability of the reference report given the report of buyer  $i$ .  $Pr[q_k|q_j]$  is the probability that the reference reporter observes the signal  $O_{ref(i)} = q_k$  given that buyer  $i$  observed the signal  $O_i = q_j$ :

$$Pr[q_k|q_j] = \sum_{\theta \in \Theta} Pr[q_k|\theta]Pr[\theta|q_j]; \quad (3.6)$$

where  $Pr[\theta|q_j]$  is the posterior probability of the type  $\theta$  given the observation  $q_j$ , computed from Bayes' Law as in Eq. (3.1).

Assume that the true signal observed by the buyer  $i$  is  $o_i \in Q$ . If the reference report is truthful, the payment expected by buyer  $i$  from the reputation mechanism is the weighted sum of the payments received for all possible values of the reference report:

$$\sum_{r_{ref(i)} \in Q} Pr[r_{ref(i)}|o_i] \tau(r_i, r_{ref(i)});$$

Miller et al. prove that for any proper scoring rule  $ScR(r_{ref(i)}|r_i)$ , the expected payment of a reporter is maximized by reporting the truth: i.e.,

$$\sum_{r_{ref(i)} \in Q} Pr[r_{ref(i)}|o_i] \left( \tau(o_i, r_{ref(i)}) - \tau(r_i, r_{ref(i)}) \right) > 0;$$

for all signals  $o_i \in Q$ ,  $r_i \neq o_i \in Q$ . This makes honest reporting a Nash equilibrium of the mechanisms. Moreover, they show that by scaling the payments appropriately, i.e.,

$$\tau(r_i, r_{ref(i)}) = aScR(r_{ref(i)}|r_i) + b; \quad (3.7)$$

where  $a$  and  $b$  are constants, the expected reward when reporting truthfully is large enough to cover the effort of reporting.

Intuitively, incentive-compatible payments exploit the correlation between the private signal observed by an agent, and the agent's beliefs regarding the reference report. Every quality signal privately observed by the agent will generate a different posterior belief regarding the true type of the product<sup>4</sup>, and consequently a different expected distribution for the value of the reference report. By paying reporters according to how well the public posterior belief (updated with the submitted report) predicts the actual reference report (assumed honest), agents have the incentive to "align" the public belief to their private beliefs, and thus report the truth.


### 3.3 Automated Design of Incentive-compatible Payment Mechanisms

The payments based on scoring rules proposed by Miller et al. (Equations (3.2) and (3.7)) make honest reporting optimal for all possible signals privately observed by a buyer. Practical mechanisms, however, require certain margins for truth-telling. Honest reporting must be better than lying by at least some margin  $\Lambda$ , chosen by the mechanism designer to offset the external benefits an agent might obtain by lying. The payments based on proper scoring rules can be scaled to account for this margin. Nevertheless, this scaling can lead to arbitrarily high feedback payments. This can be a problem because the payments cause a loss to the reputation mechanism that must be made up in some way, either by sponsorship or by charges levied on the users of the reputation information.

In this section I will investigate incentive compatible payments derived by *automated mechanism design* (Conitzer and Sandholm, 2002; Sandholm, 2003). The basic idea is to define the payments through an optimization problem that, for example, minimizes the budget required by the reputation mechanism to achieve a certain truth-telling margin  $\Lambda$ . We thus lose the simplicity of a closed-form scoring rule, but gain in efficiency of the mechanism by having the guarantee that the resulting payments are the lowest we can have in a particular context.

Formally, let  $s_i = (s_i(0), \dots, s_i(M-1))$  be the reporting strategy of buyer  $i$ , such that the buyer announces  $s_i(j) \in Q$  whenever she observes the signal  $q_j$ . The honest reporting strategy is  $\bar{s}$ , such that

<sup>4</sup>The fact that different observations trigger different posterior beliefs is grounded in Bayesian theory, but was also experimentally proven in psychological studies (Prelec, 2004)



**Figure 3.1:** Reporting feedback. Choices and Payoffs.

the buyer always reports the truth:  $\bar{s}(j) = q_j$  for all signals  $q_j \in Q$ . A summary of the notation is presented in Appendix 3.A.

Assuming that buyer  $i$  actually observed the quality signal  $q_j$ , her report  $r_i = s_i(j)$  is determined by the reporting strategy,  $s_i$ . Likewise, when the reference reporter actually observes the signal  $q_k$ , the report she submits to the reputation mechanism is  $r_{ref(i)} = s_{ref(i)}(k)$ , according to the reporting strategy  $s_{ref(i)}$ . The payment expected by  $i$  is therefore:

$$\begin{aligned} V(s_i, s_{ref(i)}|q_j) &= E_{q_k \in Q} \left[ \tau(s_i(j), s_{ref(i)}(k)) \right] \\ &= \sum_{q_k \in Q} Pr[O_{ref(i)} = q_k | O_i = q_j] \tau(s_i(j), s_{ref(i)}(k)); \end{aligned} \quad (3.8)$$

where the expectation is taken with respect to the signal  $q_k \in Q$  that is observed by the reference reporter, and  $Pr[q_k|q_j]$ , computed according to Eq. (3.6), is the conditional probability that the reference reporter observes  $q_k$  given that buyer  $i$  observed  $q_j$ .

Reporting feedback is usually costly, and buyers may obtain external benefits from lying. Let  $C_r \geq 0$  be an upper bound for the feedback reporting cost of one buyer, and let  $\Lambda(q_j, q_h)$  be an upper bound on the external benefit a buyer can obtain from falsely reporting the signal  $q_h$  instead of  $q_j$ . The cost of reporting  $C_r$  is assumed independent of the beliefs and observations of the buyer; moreover, for all signals  $q_j \neq q_k \in Q$ ,  $\Lambda(q_j, q_j) = 0$  and  $\Lambda(q_j, q_k) \geq 0$ .

Let us now consider the buyer  $i$  who purchases the product and observes the quality signal  $O_i = q_j$ . When asked by the reputation mechanism to submit feedback, the buyer can choose: (a) to honestly report  $q_j$ , (b) to report another signal  $s_i(j) \neq q_j \in Q$  or (c) not to report at all. Figure 3.1 presents the buyer's expected payoff for each of these cases, given the payment scheme  $\tau(\cdot, \cdot)$  and the reporting strategy  $s_{ref(i)}$  of the reference reporter.

Truthful reporting is a Nash equilibrium (NE) if the buyer finds it optimal to announce the true signal, whenever the reference reporter also reports the truth. Formally, the honest reporting strategy  $\bar{s}$  is a NE if and only if for all signals  $q_j \in Q$ , and all reporting strategies  $s^* \neq \bar{s}$ :

$$\begin{aligned} V(\bar{s}, \bar{s}|q_j) &\geq V(s^*, \bar{s}|q_j) + \Lambda(q_j, s^*(j)); \\ V(\bar{s}, \bar{s}|q_j) &\geq C_r; \end{aligned}$$

When the inequalities are strict, honest reporting is a strict NE, and the corresponding payment mechanism  $\tau(\cdot, \cdot)$  is incentive-compatible.

For any observed signal  $O_i = q_j \in Q$ , there are  $M - 1$  different dishonest reporting strategies  $s^* \neq \bar{s}$  the buyer can use: i.e., report  $s^*(j) = q_h \in Q \setminus \{q_j\}$  instead of  $q_j$ . Using Eq. (3.8) to expand the

expected payment of a buyer, the NE conditions become:

$$\begin{aligned} \sum_{q_k \in Q} Pr[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) &> \Lambda(q_j, q_h); \\ \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) &> C_r; \end{aligned} \quad (3.9)$$

for all  $q_j, q_h \in Q$ ,  $q_j \neq q_h$ .

Given the incentive-compatible payment scheme  $\tau(\cdot, \cdot)$ , the expected amount paid by the reputation mechanism to an honest buyer is:

$$E_{q_j \in Q} [V(\bar{s}, \bar{s}|q_j)] = \sum_{q_j \in Q} Pr[q_j] \left( \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \right);$$

The optimal payment scheme minimizes the budget required by the reputation mechanism, and therefore solves the following linear program (i.e., linear optimization problem):

### LP 3.3.1

$$\begin{aligned} \min \quad & E_{q_j \in Q} [V(\bar{s}, \bar{s}|q_j)] = \sum_{q_j \in Q} Pr[q_j] \left( \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \right) \\ \text{s.t.} \quad & \sum_{q_k \in Q} Pr[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) \geq \Lambda(q_j, q_h); \quad \forall q_j, q_h \in Q, q_j \neq q_h; \\ & \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \geq C_r; \quad \forall q_j \in Q \\ & \tau(q_j, q_k) \geq 0; \quad \forall q_j, q_k \in Q \end{aligned}$$

The payment scheme  $\tau(\cdot, \cdot)$  solving LP 3.3.1 depends on the cost of reporting, on the external benefits from lying, and on the prior belief about the type of the product. To illustrate these payments the next subsection introduce a very simple example.

### 3.3.1 Example

Alice, the owner of a new house, needs some plumbing work done. She knows there are *good* (type  $\theta_G$ ) and *bad* (type  $\theta_B$ ) plumbers, i.e.,  $\Theta = \{\theta_G, \theta_B\}$ . Alice picks the plumber from the Yellow Pages, and given the reputation of the source, she believes that the plumber, Bob, is likely to be good: e.g.,  $Pr[\theta_G] = 0.8$  and  $Pr[\theta_B] = 0.2$ . However, even a good plumber can sometimes make mistakes and provide low quality service. Similarly, a bad plumber gets lucky from time to time and provides satisfactory service. Alice does not have the necessary expertise to judge the particular problem she is facing; she therefore perceives the result of the plumber's work as a random signal conditioned on Bob's true type. Let us assume that the probability of a successful service (i.e., *high* quality,  $q_1$ ) is 0.9 if the plumber is good, and 0.15 if the plumber is bad (the probabilities of a *low* quality service,  $q_0$ , are 0.1 and 0.85 respectively). Following the notation in Section 3.1, we have:  $Pr[q_1|\theta_G] = 1 - Pr[q_0|\theta_G] = 0.9$  and  $Pr[q_1|\theta_B] = 1 - Pr[q_0|\theta_B] = 0.15$ .

Considering the prior belief, and the conditional distribution of quality signals, Alice expects to receive high quality with probability:  $Pr[q_1] = 1 - Pr[q_0] = Pr[q_1|\theta_G]Pr[\theta_G] + Pr[q_1|\theta_B]Pr[\theta_B] = 0.75$ .

Once Bob gets the work done, Alice observes the result and learns something new about Bob's type. If Alice sees good work, her posterior belief regarding the type of Bob will be  $Pr[\theta_G|q_1] =$

	$\tau(q_1, q_1)$	$\tau(q_1, q_0)$	$\tau(q_0, q_1)$	$\tau(q_0, q_0)$	
min	0.6525	0.0976	0.0975	0.1525	
s.t.	0.87	0.13	-0.87	-0.13	$\geq 0.06$
	0.87	0.13			$\geq 0.01$
	-0.39	-0.61	0.39	0.61	$\geq 0.02$
			0.39	0.61	$\geq 0.01$
	$\geq 0$	$\geq 0$	$\geq 0$	$\geq 0$	

**Table 3.1:** The optimization problem LP 3.3.1 for the example in Section 3.3.1.

$1 - Pr[\theta_B|q_1] = 0.96$  (computed by Bayes' Law), and therefore, Alice will believe that some other client will get good service from Bob with probability:  $Pr[q_1|q_1] = Pr[q_1|\theta_G]Pr[\theta_G|q_1] + Pr[q_1|\theta_B]Pr[\theta_B|q_1] = 0.87$ . On the other hand, if Alice is not happy with the work done by Bob, her posterior belief will be:  $Pr[\theta_G|q_0] = 1 - Pr[\theta_B|q_0] = 0.32$ , and she will expect another client to receive good service from Bob with probability:  $Pr[q_1|q_0] = Pr[q_1|\theta_G]Pr[\theta_G|q_0] + Pr[q_1|\theta_B]Pr[\theta_B|q_0] = 0.39$ .

Alice can submit one binary feedback (i.e.,  $q_0$  or  $q_1$ ) to an online reputation mechanism. Let the price of the plumber's work be fixed and normalized to 1, and the cost of formatting and submitting feedback be  $C_r = 0.01$ . Alice has clear incentives to misreport:

- by reporting low quality when she actually received high quality, Alice can hope to both decrease the price and increase the future availability of this (good) plumber. Assume that the external benefits of lying can be approximated as  $\Lambda(q_1, q_0) = 0.06$
- by reporting high quality when she actually received low quality, Alice can hope to decrease the relative reputation of other plumbers and thus obtain a faster (or cheaper) service from a better plumber in the future. Assume the lying incentive can be approximated as  $\Lambda(q_0, q_1) = 0.02$

The optimal feedback payments solve the optimization problem presented in Table 3.1, and have the following structure:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$
$q_0$	0.082	0
$q_1$	0	0.085

The expected payment to an honest buyer is 0.066, i.e., 6.6% of the price of the service.

Note that the payment mechanism described in Table 3.1 can only be used for the first report submitted by Alice after the prior distribution  $Pr[\theta_G] = 1 - Pr[\theta_B] = 0.8$ . Once the reputation mechanism recorded the feedback reported by Alice and updated the reputation of the Bob, new payments must be computed to provide honest reporting incentives for the next client employing Bob's services. In an online setting where feedback reports arrive sequentially, the payment mechanism must be computed anew after every report. It is therefore essential that the payments be computed efficiently (further details in Section 3.3.3).

Another remark is that the design process relies heavily on the prior distribution over types. The reputation mechanism may collect these priors by considering all information available in the environ-

ment about a specific product. From a practical perspective, the priors need not be precise: as long as agents believe in the same priors as the reputation mechanism (this will typically be the case as individual agents believe that they possess less information than the mechanism) they will subjectively regard honest reporting as the optimal strategy. Moreover, Section 3.5 shows how to design robust incentive compatible payments even when the beliefs of the agents differ slightly from the beliefs of the reputation mechanism.

### 3.3.2 Unknown Lying Incentives

LP 3.3.1 reveals a strong correlation between the minimum expected cost and the external benefits obtained from lying: low lying incentives generate lower expected payments. When finding accurate approximations for the lying incentives is difficult, the mechanism designer might want to compute the payment scheme that satisfies certain budget constraints, and maximizes the tolerated misreporting incentives. The algorithm for computing these payments follows directly from LP 3.3.1: the objective function becomes a constraint (e.g., expected budget is bounded by some amount,  $B$ ) and the new objective is to maximize the worst case (i.e., minimum) expected payment loss caused by misreporting:

#### LP 3.3.2

$$\begin{aligned}
 & \max \quad \Lambda \\
 & \text{s.t.} \quad \sum_{q_j \in Q} Pr[q_j] \left( \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \right) \leq B; \\
 & \quad \sum_{q_k \in Q} Pr[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) \geq \Lambda; \quad \forall q_j, q_h \in Q, q_j \neq q_h; \\
 & \quad \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \geq \Lambda; \quad \forall q_j \in Q \\
 & \quad \tau(q_j, q_k) \geq 0; \forall q_j, q_k \in Q
 \end{aligned}$$

The resulting scheme guarantees that any buyer will report honestly when the reporting costs and external lying benefits are smaller than  $\Lambda$ .

Coming back to the example in Section 3.3.1, let us assume that the benefits obtained from lying cannot be estimated accurately. Given the same limit on the expected budget (i.e.,  $B = 0.066$ ), we want to compute the payment scheme that maximizes the tolerance to lying. Solving LP 3.3.2 gives  $\Lambda = 0.047$  and the payments:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$
$q_0$	0.073	0
$q_1$	0	0.124

### 3.3.3 Computational Complexity and Possible Approximations

The linear optimization problems LP 3.3.1 and LP 3.3.2 are similar in terms of size and complexity: LP 3.3.1 has  $M^2$  variables and  $M^2$  inequality constraints, LP 3.3.2 has  $M^2 + 1$  variables and  $M^2 + 1$  inequality constraints. This section will therefore analyze the complexity (and runtime) of LP 3.3.1, knowing that the same conclusions extend to LP 3.3.2 as well.

M	CPU time [ms]	M	CPU time [ms]
2	11.16 (std = 3.5)	10	92.79 (std = 7.5)
4	19.24 (std = 3.7)	12	174.81 (std = 11.1)
6	29.22 (std = 4.4)	14	316.63 (std = 18.4)
8	55.62 (std = 6.7)	16	521.47 (std = 25.4)

**Table 3.2:** Average CPU time (and standard deviation) for computing the optimal payment scheme, depending on the number  $M$  of quality signals.

The worst case complexity of linear optimization problems is  $O(n^4L)$ , where  $n = M^2$  is the number of variables, and  $L$  is the size of the problem (approximately equal to the total number of bits required to represent the problem). The average time required to solve LP 3.3.1 has been experimentally evaluated using the standard linear solver in the Optimization Toolbox of Matlab 7.0.4. For different sizes of the feedback set (i.e., different values of  $M$ ) 2000 problems were randomly generated as described in Appendix 3.B. Table 3.2 presents the average CPU time required to find the optimal payment scheme on an average laptop: e.g., 1.6 GHz Centrino processor, 1Gb RAM, WinXP operating system. Up to  $M = 16$  possible quality signals, general purpose hardware and software can find the optimal payment scheme in less than half a second.

The optimal payment scheme depends on the prior belief regarding the type of the product, and therefore, must be recomputed after every submitted feedback. Although linear optimization algorithms are generally fast, frequent feedback reports could place unacceptable workloads on the reputation mechanism. Two solutions can be envisaged to ease the computational burden:


- publish batches of reports instead of individual ones. The beliefs of the buyers thus change only once for every batch, and new payments must be computed less frequently. The right size for the batch should be determined by considering the frequency of submitted reports and the tradeoff between computational cost, and the efficiency losses due to delayed information.
- approximate the optimal payments, either by closed form functions (e.g., scoring rules) or by partial solutions of the optimization problem. The rest of this section develops on these latter techniques.

The first approximation for the optimal incentive compatible payment scheme is provided Miller et al. (2005), where  $\tau(q_j, q_k) = ScR(q_k|q_j)$  is defined by a proper scoring rule. The three best known proper scoring rules are the logarithmic, the spherical and the quadratic scoring rules defined by Eq. (3.3), (3.5) and (3.4) respectively.

The constraints from LP 3.3.1 can be satisfied by: (a) adding a constant to all payments such that they become positive: i.e.,  $\tau(q_j, q_k) = \tau(q_j, q_k) - \min_{q_h, q_l \in Q} \tau(q_h, q_l)$ , and (b) multiplying all payments with a constant such that the expected payment loss when lying outweighs external benefits: i.e.,  $\tau(q_j, q_k) = const \cdot \tau(q_j, q_k)$  where:

$$const = \max_{\substack{s^*(j), q_j \in Q \\ s^*(j) \neq q_j}} \frac{\Lambda(q_j, s^*(j))}{\left( V(\bar{s}, \bar{s}|q_j) - V(s^*, \bar{s}|q_j) \right)}; \quad (3.10)$$

For the example in Section 3.3.1, the payments computed based on scoring rules (properly scaled according to Eq. (3.10)) are the following:



**Figure 3.2:** Incentive-compatible payments based on proper scoring rules.

- $\tau_{log}(q_1, q_1) = 0.27$ ,  $\tau_{log}(q_1, q_0) = 0$ ,  $\tau_{log}(q_0, q_1) = 0.17$ ,  $\tau_{log}(q_0, q_0) = 0.21$  and an expected cost of 0.22 for the logarithmic scoring rule;
- $\tau_{sph}(q_1, q_1) = 0.2$ ,  $\tau_{sph}(q_1, q_0) = 0$ ,  $\tau_{sph}(q_0, q_1) = 0.11$ ,  $\tau_{sph}(q_0, q_0) = 0.15$  and an expected cost of 0.17 for the spherical scoring rule;
- $\tau_{quad}(q_1, q_1) = 0.23$ ,  $\tau_{quad}(q_1, q_0) = 0$ ,  $\tau_{quad}(q_0, q_1) = 0.13$ ,  $\tau_{quad}(q_0, q_0) = 0.18$  and an expected cost of 0.19 for the quadratical scoring rule;

The payments based on scoring rules are two to three times more expensive than the optimal ones. The same ratio remains valid for more general settings. For 2000 randomly generated settings (see Appendix 3.B) and different number of quality signals, Figure 3.2 plots the average expected payment to one buyer when payments are computed using scoring rules.

Computational methods can also be used to obtain faster approximations of the optimal payment scheme. Most linear programming algorithms find the optimal solution by iterating through a set of feasible points that monotonically converge to the optimal one. Such algorithms are *anytime* algorithms as they can be stopped at any time, and provide a feasible solution (i.e., a payment scheme that is incentive-compatible, but maybe not optimal). The more time available, the better the feasible solution. The reputation mechanism can thus set a deadline for the optimization algorithm, and the resulting payment scheme makes it optimal for the buyers to report the truth.

Figure 3.3 plots the convergence of the Matlab linear programming algorithm for large problems (i.e., large number of possible quality signals) where approximations are likely to be needed. For 500 randomly generated settings, we plot (on a logarithmic scale) the average relative cost (relative to the optimal one) of the partial solution available after  $t$  iteration steps of the algorithm. As it can be seen, most of the computation time is spent making marginal improvements to the partial solution. For  $M = 50$  possible quality signals, the full optimization takes 20 steps on average. However, the partial solution after 6 steps generates expected costs that are only 40% higher on the average than the optimal ones.

Finally, the two techniques can be combined to obtain fast accurate approximations. As many linear programming algorithms accept initial solutions, one could use an approximation computed using the scoring rules to specify a starting point for an iterative optimization algorithm.



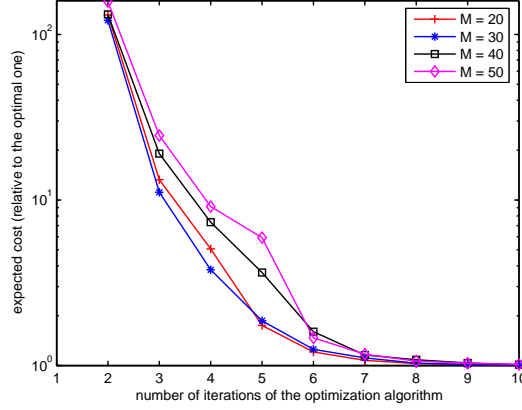


Figure 3.3: Incentive-compatible payments based on partial solutions.

### 3.4 Further Decreasing the Feedback Payments

The payment scheme computed through LP 3.3.1 generate, by definition, the lowest expected budget required by an incentive-compatible reputation mechanism, for a given setting. In this section I investigate two modifications of the mechanism itself, in order to lower the cost of reputation management even further. The first, proposes the use of several reference reporters for scoring one feedback. I formally prove that the higher the number of reference raters, the lower becomes the expected cost of reputation management.

The second idea is to reduce the potential lying incentives by filtering out false reports. Intuitively, the false feedback reports that bring important external benefits must significantly differ from the average reports submitted by honest reporters. Using a probabilistic filter that detects and ignores *abnormal* reports, lying benefits can be substantially reduced. The constraints on the optimal payments thus become more relaxed, and the optimal expected cost decreases.

#### 3.4.1 Using Several Reference Reports

The setting described in Section 3.1 is modified in the following way: we consider  $N_{ref}$  (instead of only one) reference reports when computing the feedback payment due to an agent. By an abuse of notation we use the same  $ref(i)$  to denote the set of  $N_{ref}$  reference reporters of agent  $i$ . The signals observed and reported by the reference reporters are now sets containing  $N_{ref}$  elements.

The order in which reference reports are submitted to the reputation mechanism is not important. We therefore represent a set of  $N_{ref}$  reports or observed signals by a vector  $(n_0, n_1, \dots, n_{M-1})$  where  $n_j$  is the number of times the signal  $q_j \in Q$  is present in the set. The set containing all possible unordered sets of  $N_{ref}$  signals is  $Q(N_{ref})$ , formally defined as:

$$Q(N_{ref}) = \left\{ (n_0, \dots, n_{M-1}) \in \mathbb{N}^M \mid \sum_{j=0}^{M-1} n_j = N_{ref} \right\};$$

Let  $\bar{q}_k \in Q(N_{ref})$  be a set of signals observed by the reference reporters  $ref(i)$ . If they report honestly, (i.e., every reference reporter reports according to the truthful strategy  $\bar{s}$ ), the expected

payment received by buyer  $i$ , given that she observed the signal  $q_j$  and reports according to the strategy  $s_i$  is:

$$V(s_i, \bar{s}|q_j) = \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k|q_j] \tau(s_i(j), \bar{q}_k);$$

where the payment function  $\tau(\cdot, \cdot)$  now depends on the report of buyer  $i$ , and the set of reports submitted by the reference reporters. The optimal payment function solves the following optimization problem:

### LP 3.4.1

$$\begin{aligned} \min \quad & E_{q_j \in Q} [V(\bar{s}, \bar{s}|q_j)] = \sum_{q_j \in Q} Pr[q_j] \left( \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k|q_j] \tau(q_j, \bar{q}_k) \right); \\ \text{s.t.} \quad & \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k|q_j] \left( \tau(q_j, \bar{q}_k) - \tau(q_h, \bar{q}_k) \right) \geq \Lambda(q_j, q_h); \quad \forall q_j, q_h \in Q, q_j \neq q_h; \\ & \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k|q_j] \tau(q_j, \bar{q}_k) \geq C_r; \quad \forall q_j \in Q \\ & \tau(q_j, \bar{q}_k) \geq 0; \forall q_j \in Q, \bar{q}_k \in Q(N_{ref}) \end{aligned}$$

The optimization problem LP 3.4.1 has  $M^2$  constraints and  $M \cdot |Q(N_{ref})|$  variables, where  $|Q(N_{ref})| = \binom{N_{ref} + M - 1}{M - 1}$  is the cardinality of  $Q(N_{ref})$  (proof in Appendix 3.C).

**Proposition 3.4.1** *The minimum expected budget required by an incentive compatible reputation mechanism decreases as the number of reference reporters increases.*

PROOF. I will prove that any incentive-compatible payment scheme designed for a  $N$  reference reports can generate an incentive-compatible payment scheme for  $N + 1$  reference reports, for any value of  $N$ . Moreover, the expected payment to an honest reporter is equal under the two payment schemes. It follows that the minimum budget required by a reputation mechanism that uses  $N + 1$  reference reports is less or equal to the minimum budget required by a reputation mechanism that uses only  $N$  reference reports. Hence the conclusion of the proposition.

Formally, let  $\tau_N(\cdot, \cdot)$  be the payment scheme that solves the linear program LP 3.4.1 for  $N$  reference reports. Let  $\bar{q}_k^N \in Q(N)$  be a set of  $N$  reference reports, and let  $seq_k^N \in Q^N$  be an ordered sequence of  $N$  reference reports. The notation  $seq_k^N \sim \bar{q}_k^N$  means that the unordered set of reports present in the sequence  $seq_k^N$  is equal to  $\bar{q}_k^N$ . Let  $seq_k^N \oplus q_j$  also denote the sequence of  $N + 1$  signals obtained by appending the signal  $q_j$  at the end of the sequence  $seq_k^N$ .

Slightly abusing the notation, I will also use  $\tau_N(q_j, seq_k^N) = \tau_N(q_j, \bar{q}_k^N)$  to denote the payment when the reference reports are ordered according to the sequences  $seq_k^N \sim \bar{q}_k^N$ . The payment scheme  $\tau_{N+1}(\cdot, \cdot)$  corresponding to a setting with  $N + 1$  reference reports, can now be defined as follows: Take all possible orderings of the signals in the set  $\bar{q}_k^{N+1}$ . For each sequence  $seq_k^{N+1} \sim \bar{q}_k^{N+1}$ , discard the last signal (e.g., if  $seq_k^{N+1} = seq_k^N \oplus q_h$  then discard  $q_h$ ) and consider the payment  $\tau(q_j, seq_k^N)$  based

on the remaining  $N$  reference reports. These payments are then averaged:

$$\tau_{N+1}(q_j, \bar{q}_k^{N+1}) = \frac{\sum_{\substack{seq_k^{N+1} = seq_k^N \oplus q_h \\ seq_k^{N+1} \sim \bar{q}_k^{N+1}}} \tau_N(q_j, seq_k^N)}{\sum_{seq_k^{N+1} \sim \bar{q}_k^{N+1}} 1};$$

Let us now prove that the payment scheme  $\tau_{N+1}(\cdot, \cdot)$  satisfies the constraints of the optimization problem LP 3.4.1 corresponding to  $N + 1$  reference reports:

$$\begin{aligned} & \sum_{\bar{q}_k \in Q(N+1)} Pr[\bar{q}_k | q_j] \left( \tau_{N+1}(q_j, \bar{q}_k) - \tau_{N+1}(q_h, \bar{q}_k) \right) = \\ &= \sum_{\bar{q}_k \in Q(N+1)} \sum_{seq_k' \sim \bar{q}_k} Pr[seq_k' | q_j] \frac{\sum_{\substack{seq_k^{N+1} = seq_k^N \oplus q_l \\ seq_k^{N+1} \sim \bar{q}_k^{N+1}}} \tau_N(q_j, seq_k^N) - \tau_N(q_h, seq_k^N)}{\sum_{seq_k^{N+1} \sim \bar{q}_k^{N+1}} 1} \\ &= \sum_{\bar{q}_k \in Q(N+1)} \sum_{\substack{seq_k^{N+1} = seq_k^N \oplus q_l \\ seq_k^{N+1} \sim \bar{q}_k^{N+1}}} Pr[seq_k' | q_j] \left( \tau_N(q_j, seq_k^N) - \tau_N(q_h, seq_k^N) \right) \\ &= \sum_{\substack{seq_k^{N+1} = seq_k^N \oplus q_l \\ seq_k^{N+1} \in Q^{N+1}}} Pr[seq_k^{N+1} | q_j] \left( \tau_N(q_j, seq_k^N) - \tau_N(q_h, seq_k^N) \right) \\ &= \sum_{seq_k^N \in Q^N} Pr[seq_k^N | q_j] \left( \tau_N(q_j, seq_k^N) - \tau_N(q_h, seq_k^N) \right) \\ &= \sum_{\bar{q}_k \in Q(N)} Pr[\bar{q}_k | q_j] \left( \tau_N(q_j, \bar{q}_k) - \tau_N(q_h, \bar{q}_k) \right) \\ &> \Lambda(q_j, q_h); \forall q_j, q_h \in Q, q_j \neq q_h; \end{aligned}$$


The first transformation was obtained because  $Pr[seq_k' | q_j]$  takes the same value for all sequences  $seq_k' \sim \bar{q}_k$ . The second transformation is possible because a nested iteration over all  $\bar{q}_k \in Q(N)$  and all  $seq_k \sim \bar{q}_k$  is equivalent to one iteration over all possible sequences  $seq_k \in Q^{N+1}$ . The third transformation takes into account that  $Pr[seq_k^N \oplus q_l | q_j] = Pr[seq_k^N | q_j] \cdot Pr[q_l | q_j]$  and that  $\sum_{q_l \in Q} Pr[q_l | q_j] = 1$ . Finally, the last transformation is possible because  $\tau_N(\cdot, seq_k^N) = \tau_N(\cdot, \bar{q}_k)$  for all sequences  $seq_k^N \sim \bar{q}_k$ .

Similarly, it is easy to prove that:

$$\sum_{\bar{q}_k \in Q(N+1)} Pr[\bar{q}_k | q_j] \tau_{N+1}(q_j, \bar{q}_k) = \sum_{\bar{q}_k \in Q(N)} Pr[\bar{q}_k | q_j] \tau_N(q_j, \bar{q}_k) > C_r; \forall q_j \in Q;$$

and that the expected payment made to an honest reporter is the same under the payment schemes  $\tau_N(\cdot, \cdot)$  and  $\tau_{N+1}(\cdot, \cdot)$ .  $\blacksquare$

Using several reference reports decreases the cost of reputation management, but also increases the complexity of the algorithm defining the optimal payment scheme. To visualize the tradeoff between the two, I generated 2000 random settings (details available in Appendix 3.B) and plotted in Figure 3.4 the average cost of the incentive-compatible payment mechanism as the number of reference reports



**Figure 3.4:** Average expected payment to one agent when several reference reports are used.

increases from 1 to 5. Significant savings (approx. 25% for a setting with  $M = 2$  quality signals, and 4% for a setting with  $M = 8$  quality signals) are mainly obtained from the second and third reference reports. Practical systems can therefore function in this sweetspot where 2 to 4 reference reports are used to decrease the budget required by the reputation mechanism without increasing too much the complexity of the design problem.

For the example in Section 3.3.1, using 2 reference reports reduces the expected payment to one reporter from 0.066 to 0.054, and generates the following payment mechanism:

$\tau(\cdot, \cdot)$	$(2 \times q_0)$	$(q_0, q_1)$	$(2 \times q_1)$
$q_0$	0.087	0	0
$q_1$	0	0	0.081

Using 3 reference reports further reduces the expected cost of the reputation mechanism to 0.051 and generates the following payments:

$\tau(\cdot, \cdot)$	$(3 \times q_0)$	$(2 \times q_0, q_1)$	$(q_0, 2 \times q_1)$	$(3 \times q_1)$
$q_0$	0.112	0	0	0
$q_1$	0	0	0	0.091

### 3.4.2 Filtering out False Reports

The feedback payments naturally decrease when the reporting and honesty costs become smaller. The cost of reporting can be decreased by software tools that help automate as much as possible the process of formatting and submitting feedback. On the other hand, the external incentives for lying can be reduced by filtering out the reports that are likely to be false.

*Truth filters* can be constructed based on statistical analysis. When all agents report truthfully, their reports follow a common distribution given by the product's true type. Reports that stand out

from the common distribution are either particularly unlikely, or dishonest. Either way, by filtering them out with high probability, the reputation information does not usually suffer significant degradation.

Probabilistic filters of false reports have been widely used in decentralized and multi-agent systems. Vu, Hauswirth, and Aberer (2005) use clustering techniques to isolate lying agents in a market of web-services. Their mechanism is based on a small number of “trusted” reports that provide the baseline for truthful information. The technique shows very good experimental results when lying agents use probabilistic strategies and submit several reports. Buchegger and Le Boudec (2005) use Bayesian methods to detect free riders in a wireless ad-hoc network. Nodes consider both direct and second-hand information, however, second-hand information is taken into account only when it does not conflict with direct observations (i.e., second hand reports do not trigger significant deviations in the agent’s beliefs). In peer-to-peer reputation mechanisms (e.g., TRAVOS (Teacy et al., 2005), CREDENCE (Walsh and Siner, 2005) and (Yu and Singh, 2003; Ismail and Jøsang, 2002; Whitby et al., 2004)) agents weigh the evidence from peers by the distance from the agent’s direct experience.

However, all of the results cited above rely on two important assumptions: a) every agent submits several reports, b) according to some *probabilistic* lying strategy. Self-interested agents can strategically manipulate their reports to circumvent the filtering mechanisms and take profit from dishonest reporting<sup>5</sup>. When all buyers are self-interested and submit only one feedback report, filtering methods based entirely on similarity metrics can never be accurate enough to filter out effectively *all* lying strategies without important losses of information.

In this section, I present an alternative filtering method that also exploits the information available to the agents. The intuition behind this method is simple: the probability of filtering out the report  $r_i$  submitted by agent  $i$  should not only depend on how well  $r_i$  fits the distribution of peer reports, but also on the benefits that  $r_i$  could bring to the reporter if it were false. When  $\Lambda(q_j, r_i)$  is big (i.e. the agent has strong incentives to report  $r_i$  whenever her true observation was  $q_j$ ), the filtering mechanism should be more strict in accepting  $r_i$  given that peer reports make the observation of  $q_j$  probable. On the other hand, when  $\Lambda(q_j, r_i)$  is small, filtering rules can be more relaxed, such that the mechanism does not lose too much information. In this way, the filter adapts to the particular context and allows an optimal tradeoff between diminished costs and loss of information.

Concretely, let  $Pr[\theta]$ ,  $\theta \in \Theta$  describe the current common belief regarding the true type of the product, let  $o_i, r_i \in Q$  denote the signals observed, respectively reported by agent  $i$ , and let  $\bar{q}_k \in Q(N_{ref})$  describe the set of  $N_{ref}$  reference reports. The publishing of the report submitted by agent  $i$  is delayed until the next  $N_{ftr}$  reports (i.e., the filtering reports) are also available. A filtering mechanism is formally defined by the table of probabilities  $\pi(r_i, \hat{q}_k)$  of accepting the report  $r_i \in Q$  when the filtering reports take the value  $\hat{q}_k \in Q(N_{ftr})$ . With probability  $1 - \pi(r_i, \hat{q}_k)$  the report  $r_i$  will not be published by the reputation mechanism, and therefore will not be reflected in the reputation information. Note, however, that all reports (including the dropped ones) are paid for as described in the previous sections.

The payment scheme  $\tau(\cdot, \cdot)$  and the filtering mechanism  $\pi(\cdot, \cdot)$  are incentive compatible if and only if for all signals  $q_j, q_h \in Q$ ,  $q_j \neq q_h$ , the expected payment loss offsets the expected gain obtained from lying:

$$\begin{aligned} \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k | q_j] \left( \tau(q_j, \bar{q}_k) - \tau(q_h, \bar{q}_k) \right) &> \hat{\Lambda}(q_j, q_h) \\ \hat{\Lambda}(q_j, q_h) &= \sum_{\hat{q}_k \in Q(N_{ftr})} Pr[\hat{q}_k | q_j] \cdot \pi(q_h, \hat{q}_k) \cdot \Lambda(q_j, q_h); \end{aligned} \tag{3.11}$$

where  $\hat{\Lambda}(\cdot, \cdot)$  is obtained by discounting  $\Lambda(\cdot, \cdot)$  with the expected probability that a false report is

<sup>5</sup>It is true, however, that some mechanisms exhibit high degrees of robustness towards such lying strategies: individual agents can profit from lying, but as long as the big majority of agents reports honestly, the liars do not break the properties of the reputation mechanism

recorded by the reputation mechanism.

Naturally, the feedback payments decrease with decreasing probabilities of accepting reports. However, a useful reputation mechanism must also limit the information lost by discarding feedback reports. As a metric for information loss I chose the number (or percentage) of *useful* reports that are dropped by the mechanism. A feedback report is *useful*, when given the true type of the product and a prior belief on the set of possible types, the posterior belief updated with the report is closer to the true type than the prior belief.

For the example in Section 3.3.1, when the plumber is actually good, recording a *high* quality report is *useful* (because the posterior belief is closer to reality than the prior belief), while recording a *low* quality report is not. Conversely, when the plumber is bad, recording a *low* quality report is useful, while recording a *high* quality report is not. The notion of usefulness captures the intuition that some reports can be filtered out in some contexts without any loss of information for the buyers (on the contrary, the community has more accurate information without the report).

Formally, information loss can be quantified in the following way. Given the true type  $\theta^* \in \Theta$  and the prior belief  $Pr[\cdot]$  on the set of possible types, the report  $q_j$  is useful if and only if  $Pr[\theta^*] < Pr[\theta^*|q_j]$ : i.e. the posterior belief updated with the signal  $q_j$  is closer to the true type than the prior belief. Given the filtering mechanism  $\pi(\cdot, \cdot)$ , and the true type  $\theta^*$ , the expected probability of dropping  $q_j$  is:

$$Pr[\text{drop } q_j | \theta^*] = 1 - \sum_{\hat{q}_k \in Q(N_{ftr})} Pr[\hat{q}_k | \theta^*] \pi(q_j, \hat{q}_k); \quad (3.12)$$

where  $Pr[\hat{q}_k | \theta^*]$  is the probability that the filtering reports take the value  $\hat{q}_k$ , when the true type of the product is  $\theta^*$ . To limit the loss of information, the reputation mechanism must insure that given the current belief, whatever the true type of the product, no useful report is dropped with a probability greater than a given threshold,  $\gamma_{infLoss}$ :

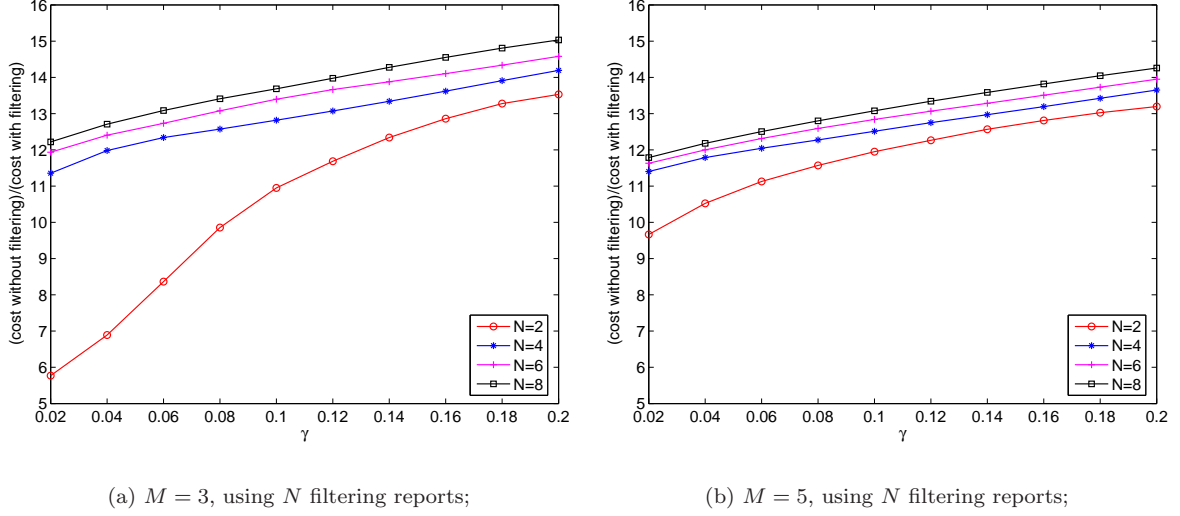
$$\forall q_j \in Q, \theta \in \Theta, \quad Pr[\theta] < Pr[\theta|q_j] \Rightarrow Pr[\text{drop } q_j | \theta] < \gamma_{infLoss}; \quad (3.13)$$

I can now define the incentive-compatible payment mechanism (using  $N_{ref}$  reference reports) and the filtering mechanism (using  $N_{ftr}$  filtering reports) that minimize the expected payment to an honest reporter:

### LP 3.4.2

$$\begin{aligned} \min \quad & E_{q_j \in Q} [V(\bar{s}, \bar{s}|q_j)] = \sum_{q_j \in Q} Pr[q_j] \left( \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k | q_j] \tau(q_j, \bar{q}_k) \right); \\ \text{s.t.} \quad & \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k | q_j] \left( \tau(q_j, \bar{q}_k) - \tau(q_h, \bar{q}_k) \right) \geq \hat{\Lambda}(q_j, q_h); \quad \forall q_j, q_h \in Q, q_j \neq q_h, \\ & \hat{\Lambda}(q_j, q_h) \text{ is defined in (3.11);} \\ & \sum_{\bar{q}_k \in Q(N_{ref})} Pr[\bar{q}_k | q_j] \tau(q_j, \bar{q}_k) \geq C_r; \quad \forall q_j \in Q \\ & Pr[\theta] < Pr[\theta|q_j] \Rightarrow Pr[\text{drop } q_j | \theta] \leq \gamma_{infLoss}; \quad \forall \theta \in \Theta, \forall q_j \in Q \\ & \tau(q_j, \bar{q}_k) \geq 0, \pi(q_j, \hat{q}_k) \in [0, 1] \quad \forall q_j \in Q, \forall \bar{q}_k \in Q(N_{ref}), \forall \hat{q}_k \in Q(N_{ftr}); \end{aligned}$$

The effect of using probabilistic filtering of reports was experimentally studied on 500 randomly generated settings, for different number of filtering reports (i.e.,  $N_{ftr}$ ), different number of quality signals (i.e.,  $M$ ) and different values of the parameter  $\gamma_{infLoss}$ . Figure 3.5 plots the tradeoff between the reduction of cost (i.e. the ratio between the optimal cost without probabilistic filtering and the



**Figure 3.5:** Tradeoff between cost decrease and information loss.

optimal cost with probabilistic filtering) and the information loss for  $M = 3$  and  $M = 5$  quality signals. When  $M = 3$ , and we accept losing 2% of the useful reports, the cost decreases 6 times by using  $N_{ftr} = 2$  filtering reports, and 12 times by using  $N_{ftr} = 8$  filtering reports. As intuitively expected, the cost decreases when we can use more filtering reports, and accept higher probabilities of losing useful feedback.

As a next experiment, I studied the accuracy of the reputation information published by a mechanism that filters out reports. For each of the random settings generated above, I also generate 200 random sequences of 20 feedback reports corresponding to a randomly chosen type. For different parameters (i.e., number of signals,  $M$ , number of filtering reports,  $N_{ftr}$ , and threshold probability  $\gamma_{infLoss}$ ), Figure 3.6 plots the mean square error of the reputation information<sup>6</sup> published by a mechanism that filters, respectively doesn't filter submitted reports. As expected, filtering out reports does not significantly alter the convergence of beliefs; on the contrary, filtering out reports may sometimes help to focus the beliefs on the true type of the product.

Finally, for the example in Section 3.3.1 using one reference report and 3 filtering reports with a threshold for dropping useful reports of 2%, the expected payment to one agent is 0.02, and the payments  $\tau(\cdot, \cdot)$ , respectively the filtering probabilities  $\pi(\cdot, \cdot)$  are the following:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$	$\pi(\cdot, \cdot)$	$(3 \times q_0)$	$(2 \times q_0, q_1)$	$(q_0, 2 \times q_1)$	$(3 \times q_1)$
$q_0$	0.033	0	$q_0$	1	1	0.718	0
$q_1$	0	0.023	$q_1$	0.04	0.42	0.995	1

<sup>6</sup>The mean square error after  $i$  submitted reports is defined as:  $mse_i = \sum_{\theta \in \Theta} (Pr[\theta|i] - I(\theta))^2$ , where  $Pr[\cdot|i]$  describes the belief of the agents regarding the type of the product after  $i$  submitted reports,  $I(\theta) = 1$  for  $\theta = \theta^*$  (the true type of the product), and  $I(\theta) = 0$  for  $\theta \neq \theta^*$ .

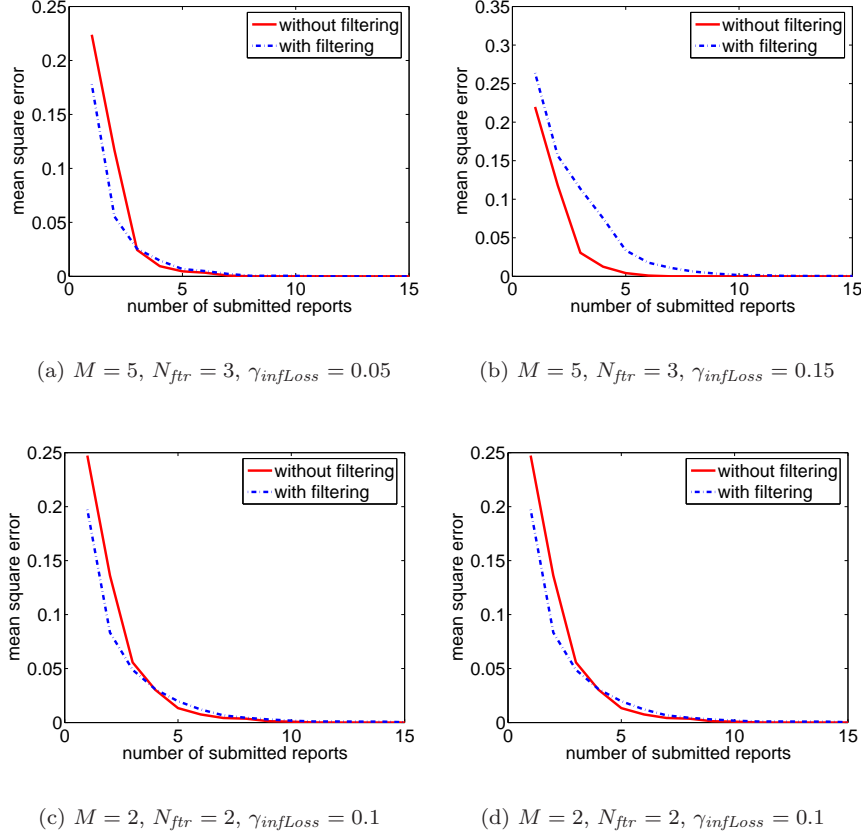


Figure 3.6: Convergence of reputation information.

## 3.5 Robust Incentive-Compatible Payment Mechanisms

One key assumption behind the payment mechanisms described in the previous section is that the reputation mechanism and the reporters share the same prior information regarding the reputation of the rated product. Only in this case the honest report aligns the posterior reputation (as computed by the reputation mechanism) with the private posterior beliefs of the agent. When reporters have private prior information unknown to the reputation mechanism, it may be possible that some lying report is better than truth telling. In this section, I investigate feedback payment mechanisms that are incentive compatible even when reporters have some private information that is unknown to the reputation mechanism.

### 3.5.1 Dishonest Reporting with Unknown Beliefs

The optimal incentive-compatible payments computed in Sections 3.2.1 and 3.3 rely on the posterior beliefs (i.e., the probabilities  $Pr[q_k|q_j]$ ) of the reporters regarding the value of the reference reports. These can be computed by the reputation mechanism from:



- the prior belief,  $Pr[\theta]$ , that the product is of type  $\theta$ ,
- the conditional probabilities,  $Pr[q_j|\theta]$ , that a product of type  $\theta$  generates the signal  $q_j$ ,

using Bayes' Law as shown in the Eq. (3.6).

However, when the reporters have different beliefs regarding the values of the reference reports, the constraints in LP 3.3.1 do not accurately reflect the decision problem of the agents, and therefore, do not always guarantee an honest equilibrium.

Let us reconsider the example from Section 3.3.1, and the corresponding payments from Table 3.1. Assume that the prior belief of Alice differs only slightly from that of the reputation mechanism. For example, Alice might have talked on the phone with the plumber Bob, so she's slightly more likely to believe that Bob is a good plumber: e.g., for Alice,  $Pr^*[\theta_G] = 0.83$ , while for the reputation mechanism,  $Pr[\theta_G] = 0.8$ . All other values remain the same. If Alice is dissatisfied with Bob's work, her private posterior belief regarding Bob's type becomes  $Pr^*[\theta_G|q_0] = 0.365$  instead of  $Pr[\theta_G|q_0] = 0.32$ , so Alice expects another client to get good service from Bob with the probability:  $Pr^*[q_1|q_0] = 1 - Pr^*[q_0|q_0] = 0.424$  (instead of  $Pr[q_1|q_1] = 0.39$  considered by the reputation mechanism). Simple arithmetics reveals that Alice is better off by reporting a positive instead of the truthful negative report:

- the expected gain from falsely submitting positive feedback is:  $0.424 \cdot 0.085 + 0.576 \cdot 0 + \Lambda(q_0, q_1) = 0.056$ ,
- while honest reporting brings only:  $0.424 \cdot 0 + 0.576 \cdot 0.082 = 0.047$ .

### 3.5.2 Declaration of Private Information

To eliminate lying incentives, Miller et al. (2005) suggest that reporters should also declare their prior beliefs before submitting feedback. The reputation mechanism could then use the extra information to compute the payments that makes truth-telling rational for every agent.

Unfortunately, such a mechanism can be exploited by self-interested agents when external benefits from lying are positive. Consider the same example as above: Alice has a prior belief which assigns probability 0.83 to the plumber being *good*.

If Alice truthfully declares her prior belief, the reputation mechanism computes the optimal payments by solving LP 3.3.1:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$
$q_0$	0.095	0
$q_1$	0	0.082

so that by truthfully reporting negative feedback, Alice expects a payment equal to:  $0.424 \cdot 0 + 0.576 \cdot 0.095 = 0.055$ .

Alice can, however, declare the prior belief:  $\overline{Pr}[\theta_G] = 1 - \overline{Pr}[\theta_B] = 0.1$ . In this case, the payment scheme computed by the reputation mechanism will be:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$
$q_0$	0.064	0
$q_1$	0	0.212

and the optimal strategy for Alice following the negative experience is to submit a positive report. Her expected payment is:  $0.424 \cdot 0.212 + 0.55 \cdot 0 + \Lambda(q_0, q_1) = 0.11 > 0.055$ , where  $\Lambda(q_0, q_1) = 0.02$  is the external benefit Alice can obtain by falsely declaring  $q_1$  instead of  $q_0$ .

The example provided above is, unfortunately, not unique. Profitable lying is possible because agents can find false prior beliefs that determine the reputation mechanism to compute feedback payments that make lying optimal. Thus, the agent obtains both the optimal feedback payment, and the external benefit from lying.

The false prior beliefs that make lying profitable can be easily computed based on the following intuition. The payment scheme defined by LP 3.3.1 makes it optimal for the agents to reveal their true *posterior* belief regarding the type of the product. When the prior belief is known, only the truly observed quality signal “aligns” the posterior belief of the reputation mechanism with that of the agent. However, when the prior belief must also be revealed, several combinations of prior belief and reported signal, can lead to the same posterior belief. Hence, the agent is free to chose the combination that brings the best external reward.

The false prior belief  $(\overline{Pr}[\theta])_{\theta \in \Theta}$  and the false signal  $q_h$  that lead to the same posterior belief  $(Pr[\theta|q_j])_{\theta \in \Theta}$ , can be computed by solving the following system of linear equations:

$$\overline{Pr}[\theta|q_h] = \frac{Pr[q_h|\theta]\overline{Pr}[\theta]}{\sum_{\theta' \in \Theta} Pr[q_h|\theta']\overline{Pr}[\theta']} = \frac{Pr[q_j|\theta]Pr[\theta]}{\sum_{\theta' \in \Theta} Pr[q_j|\theta']Pr[\theta']} = Pr[\theta|q_j]; \quad \forall \theta \in \Theta; \quad (3.14)$$

The system has  $|\Theta|$  equations and  $|\Theta| + 1$  variables (i.e., the probabilities  $\overline{Pr}[\theta]$  and the signal  $q_h$ ); therefore, there will generally be several solutions that make lying profitable. The agent may choose the one that maximizes her expected payment by solving the following nested linear problem:

$$\begin{aligned} \max \quad & \Lambda(q_j, q_h) + \sum_{q_k \in Q} Pr[q_k|q_j]\tau(q_h, q_k) \\ \text{s.t.} \quad & \overline{Pr}[\theta|q_h] = Pr[\theta|q_j]; \quad \forall \theta \in \Theta; \\ & \tau(\cdot, \cdot) \text{ solves LP 3.3.1 for the prior beliefs } \overline{Pr}[\theta] \end{aligned}$$

To enforce truth-telling, Prelec (2004) suggests payments that also depend on the declared priors. Agents are required to declare both the observed signal, and a prediction of the signals observed by the other agents (which indirectly reflects the agent’s private information). The proposed “*truth serum*” consists of two additive payments: an *information* payment that rewards the submitted report, and a *prediction* payment that rewards the declared private information. Prelec shows that honesty is the highest paying Nash equilibrium. Nonetheless, his results rely on the assumption that a prior probability distribution over all possible private beliefs (not the belief itself) is common knowledge.

Another solution has been suggested by Miller et al. (2005). Misreporting incentives can be eliminated if agents declare their prior beliefs before the actual interaction takes place. As posteriors are not available yet, the agent cannot manipulate the declared prior belief in order to avoid the penalty from lying. However, such a process has several practical limitations. First, the enforcement of prior belief declaration before the interaction can only be done if a central authority acts as an intermediary between the buyer and the seller. The central proxy may become a bottleneck and adds to the transaction cost. Second, the declaration of prior beliefs could significantly delay the access to the desired good. Finally, the reporting of priors adds to the reporting cost (reporting probability distributions is much more costly than reporting observed signals) and greatly increases the budget required by an incentive-compatible mechanism.

### 3.5.3 Computing Robust Incentive-Compatible Payments

In this section I pursue an alternative solution for dealing with unknown beliefs. I start from the assumption that the private beliefs of most rational agents will not differ significantly from those of the reputation mechanism. The beliefs of the reputation mechanism, as reflected in the publicly available reputation information, have been constructed by aggregating all feedback reports submitted by all previous users. Assuming that agents trust the reputation mechanism to publish truthful information, their private information will trigger only marginal changes to the beliefs. Thus, rather than build a system that can accommodate all private beliefs, I focus on mechanisms that are incentive-compatible for most priors, i.e., the priors within certain bounds from those of the reputation mechanism.

Let  $(Pr[\theta])_{\theta \in \Theta}$  characterize the prior belief of the reputation mechanism and let  $(Pr^*[\theta] = Pr[\theta] + \varepsilon_\theta)_{\theta \in \Theta}$  be the range of private beliefs the clients might have, where:  $\sum_{\theta \in \Theta} \varepsilon_\theta = 0$ , and

$$\max(-\gamma_{maxPI}, -Pr[\theta]) \leq \varepsilon_\theta \leq \min(\gamma_{maxPI}, 1 - Pr[\theta]), \gamma_{maxPI} > 0.$$

Replacing the private beliefs in Eq. (3.6), the conditional probabilities for the reference reporter's observed signals become:

$$Pr^*[q_k|q_j] = \frac{\sum_{\theta \in \Theta} Pr[q_k|\theta]Pr[q_j|\theta](Pr[\theta] + \varepsilon_\theta)}{\sum_{\theta \in \Theta} Pr[q_j|\theta](Pr[\theta] + \varepsilon_\theta)};$$

Let  $Pr_m^*[q_k|q_j]$  and  $Pr_M^*[q_k|q_j]$  be the minimum, respectively the maximum values of  $Pr^*[q_k|q_j]$  as the variables  $(\varepsilon_\theta)_{\theta \in \Theta}$  take values within the acceptable bounds. If the constraints of LP 3.3.1 are modified such that honest reporting is optimal for all possible values of the probabilities  $Pr^*[q_k|q_j]$ , the resulting payment mechanism is incentive compatible for all private beliefs that are not *too far* from the belief of the reputation mechanism.

Representing linear constraints for a continuous range of parameters is not accepted by linear solvers. The constraint:

$$\sum_{q_k \in Q} Pr^*[q_k|q_j] (\tau(q_j, q_k) - \tau(q_h, q_k)) > \Lambda(q_j, q_h); \quad (3.15)$$

is satisfied for all possible values of  $Pr^*[q_k|q_j] \in [Pr_m^*[q_k|q_j], Pr_M^*[q_k|q_j]]$ , only when:

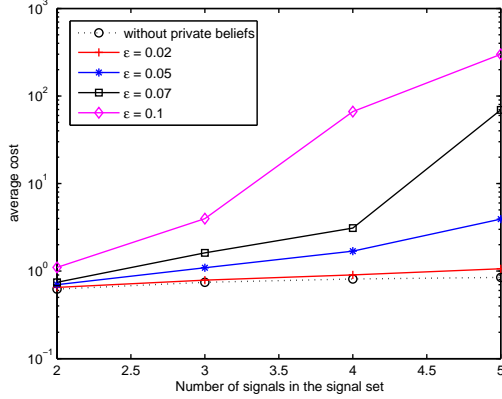
$$\min_{Pr^*[q_k|q_j]} \left( \sum_{q_k \in Q} Pr^*[q_k|q_j] (\tau(q_j, q_k) - \tau(q_h, q_k)) \right) > \Lambda(q_j, q_h); \quad (3.16)$$

If the probabilities  $Pr^*[q_k|q_j]$  were independent<sup>7</sup>, the minimum would be given by one of the combinations of extreme values: i.e.,  $Pr^*[q_k|q_j] = Pr_m^*[q_k|q_j]$  or  $Pr^*[q_k|q_j] = Pr_M^*[q_k|q_j]$ . Therefore, by replacing every constraint (3.15), with  $2^M$  linear constraints, one for every combination of extreme values of  $Pr^*[q_k|q_j]$ , we impose stricter condition than in (3.16). The optimization problem defining the payment scheme is similar to LP 3.3.1, where every constraint has been replaced by  $2^M$  linear constraints, one for every combination of extreme values of  $Pr^*[q_k|q_j]$ .

The effect of private beliefs on the expected cost of the incentive-compatible mechanism has been experimentally evaluated by generating 2000 random problems as described in Appendix 3.B. For each problem, I considered different tolerance levels to private beliefs (i.e.,  $\gamma_{maxPI} = \{0, 0.02, 0.05, 0.07, 0.1\}$ ) and solved the linear optimization problem that defines the robust, incentive compatible payments. I used average hardware (e.g., Pentium Centrino 1.6MHz, 1Gb RAM, WinXP) and the CPLEX<sup>8</sup> linear

<sup>7</sup>The probabilities  $Pr^*[q_k|q_j]$  are not independent because they depend on the same variables  $(\varepsilon_\theta)$

<sup>8</sup>www.ilog.com



**Figure 3.7:** Average expected payment to one agent for different tolerance levels of private beliefs.

M	CPU time [ms]	Std. dev. [ms]
2	14.117	4.9307
3	38.386	4.1765
4	485.33	50.546
5	798.28	722.5

**Table 3.3:** Average CPU time (and standard deviation) for computing the optimal payment scheme with private beliefs.

solver. Table 3.3 presents the average CPU time required for computing the payments. Due to the exponential number of constraints, the time required to compute the optimal payments increases exponentially with the number of signals,  $M$ . For  $M = 6$  signals, the computation already takes more than one second.

Figure 3.7 presents the average cost of an incentive-compatible payment scheme that tolerates private beliefs. It plots the average expected payment to one agent for different number of signals, and different tolerance levels for private beliefs. The cost of the mechanism increases quickly with  $\gamma_{maxPI}$ , the tolerated range of beliefs. For beliefs within  $\pm 10\%$  of those of the reputation mechanism, the cost of the mechanism increases one order of magnitude. Note, however, that the constraints defining the payment scheme are stricter than necessary. As future research, one can define non-linear algorithms that can approach the truly optimal payments.

### 3.5.4 General Tolerance Intervals for Private Information

Instead of modeling private information as small perturbations to the prior belief regarding the true type of the product, we consider in this section a more general case, where the conditional probabilities  $Pr[q_k|q_j]$  that parameterize LP 3.3.1 are allowed to vary within certain limits. Such variations can account for various sources of private information: e.g., private beliefs regarding the true type of the product, private information regarding the conditional distribution of signals, or even small changes to the true type of the product. This approach is similar to the work of Zohar and Rosenschein (2006).

Without modeling the real source of private information, I assume that the conditional probability distributions  $Pr^*[\cdot|q_j]$  (for all  $q_j$ ) are not too far from the probability distributions  $Pr[\cdot|q_j]$  computed by the reputation mechanism. I will use the  $L_2$  norm for computing the distance, and assume that:

$$\sum_{q_k \in Q} \left( Pr^*[q_k|q_j] - Pr[q_k|q_j] \right)^2 \leq \varepsilon^2; \forall q_j \in Q; \quad (3.17)$$

for some positive bound  $\varepsilon$ . The incentive-compatibility constraints must enforce that for any value of

the probabilities  $Pr^*[\cdot|\cdot]$ , honesty gives the highest payoff. Formally,

$$\begin{aligned} \min_{Pr^*[\cdot|\cdot]} & \left( \sum_{q_k \in Q} Pr^*[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) \right) > \Lambda(q_j, q_h); \forall q_j \neq q_h \in Q; \\ \text{s.t.} & \sum_{q_k \in Q} \left( Pr^*[q_k|q_j] - Pr[q_k|q_j] \right)^2 \leq \varepsilon^2; \end{aligned}$$

This optimization problem can be solved analytically by writing the Lagrangian and enforcing the first order optimality conditions. Therefore:

$$\begin{aligned} \min & \left( \sum_{q_k \in Q} Pr^*[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) \right) = \\ & \sum_{q_k \in Q} Pr[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) - \varepsilon \sqrt{\sum_{q_k \in Q} \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right)^2}; \end{aligned}$$

and the best (i.e., cheapest) incentive compatible payments that are robust to private information (i.e., have robustness level  $\varepsilon^2$ ) are obtained by solving the conic optimization problem:

#### CP 3.5.1

$$\begin{aligned} \min & E[V(\bar{s}, \bar{s}|q_j)] = \sum_{q_j \in Q} Pr[q_j] \left( \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) \right) \\ \text{s.t.} & \sum_{q_k \in Q} Pr[q_k|q_j] \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right) - \varepsilon \sqrt{\sum_{q_k \in Q} \left( \tau(q_j, q_k) - \tau(q_h, q_k) \right)^2} \geq \Lambda(q_j, q_h); \\ & \sum_{q_k \in Q} Pr[q_k|q_j] \tau(q_j, q_k) - \varepsilon \sqrt{\sum_{q_k \in Q} \tau(q_j, q_k)^2} \geq C_r; \\ & \forall q_j \neq q_h \in Q; \quad \tau(q_j, q_k) \geq 0; \forall q_j, q_k \in Q \end{aligned}$$


where  $Pr[\cdot|\cdot]$  are the probabilities computed by the reputation mechanism. Such problems can be solved by polynomial time algorithms.

I evaluated experimentally the cost of general private information as reflected on the expected payment to one reporter. For 2000 randomly generated problems (details in Appendix 3.B) and for different levels of robustness, I solved CP 3.5.1 to obtain the robust incentive compatible payments. Table 3.4 presents the average CPU time required to compute the payments. As expected, the values are much smaller than those of Table 3.3.

Figure 3.8 plots the average expected payment to one agent for different number of signals, and different tolerance levels to private information. Like in Figure 3.7, the cost of the mechanism increases exponentially with the robustness level,  $\varepsilon^2$ .

One important remark about the results of this section is that agents trust the reputation mechanism to publish truthful information. Only in this case agents are likely to adopt (with marginal changes) the beliefs of the reputation mechanism, and have incentives to report honestly. While the trustworthiness of the reputation mechanism is an interesting topic on its own, let us note that agents can verify<sup>9</sup> whether or not the payments advertised by the reputation mechanism actually “match” the beliefs and the robustness bound published by the reputation mechanism. On the other hand, the payments that

<sup>9</sup>by checking that the payments  $\tau(\cdot, \cdot)$  solve the optimization problems LP 3.3.1 or CP 3.5.1



**Figure 3.8:** Average expected payment to one agent for different levels of robustness to general private information.

M	CPU time [ms]	Std. dev. [ms]
2	3.46	24.40
3	7.31	9.04
4	16.05	8.07
5	44.89	15.74

**Table 3.4:** Average CPU time (and standard deviation) for computing the optimal payment scheme with general private information.

match the true beliefs of the reputation mechanism are the smallest possible, as guaranteed by the corresponding optimization problems.

However, understanding exactly what the reputation mechanism can do in order to manipulate reputation information without being detected, while still providing decent honest reporting incentives requires further work.

### 3.6 Collusion-resistant, Incentive-compatible Rewards

Sections 3.2.1 and 3.3 describe several ways of designing incentive-compatible payment mechanisms where honest reporting is a Nash equilibrium. Unfortunately, honest reporting is not the only Nash Equilibrium (NE) of the mechanism. Consider the plumber example from Section 3.3.1 and the payments computed by solving the optimization problem LP 3.3.1:

$\tau(\cdot, \cdot)$	$q_0$	$q_1$
$q_0$	0.082	0
$q_1$	0	0.085

Besides the honest reporting equilibrium, this payment mechanism has two other equilibria where all agents report the same thing: either  $q_0$  or  $q_1$ . Moreover, if the expected payment in the honest equilibrium is 0.066 (computed in Section 3.3.1), the payment given by the equilibrium where all agents report positively, respectively negatively, is 0.085, respectively 0.082. While the multiplicity of equilibria is generally seen a problem for other reasons (e.g., equilibrium selection), in this context it also brings forth the problem of collusion: rational agents could potentially coordinate on a lying equilibrium to extract the maximum payments from the reputation mechanism.

In this section, I will address the design of incentive-compatible reputation mechanisms that are also resistant to collusion. First, I will motivate the need for collusion-resistance measures, and show that all binary reputation mechanisms have multiple equilibria and are thus vulnerable to lying coalitions.

tions. Second, I will describe a method for designing incentive-compatible, collusion-resistant payment mechanisms by using several reference reports.

The basic idea behind deterring lying coalitions is to design incentive compatible rewards that make honest reporting the unique, or at least the “best” equilibrium. Ideally, honest reporting would be the dominant strategy for the colluders, such that no matter what other colluders do, honest reporting is optimal for every individual colluder. Clearly no rational agent should be expected to join a lying coalition under such circumstances.

When dominant truth-telling is not possible, the second most preferable option is to have honest reporting as the unique NE. Any lying coalition would imply a non-equilibrium situation, where individual colluders would rather report differently than specified by the colluding strategy. Assuming non transferable utilities and the absence of effective punishments that a coalition can enforce on its members, any lying coalition would be unstable, and therefore, unlikely to form.

Finally, when honest reporting is neither the dominant strategy, nor the unique NE, collusion resistance can emerge if honesty is the Pareto-optimal Nash equilibrium. The intuition here is that any stable (i.e., equilibrium) lying coalition would make at least one of the colluders worse off than in the honest equilibrium, which hopefully prevents that agent for joining the coalition in the first place (assuming, again, non-transferable utilities).

An interesting scenario is when a single strategic entity controls a number of fake online identities, or *sybils* (Cheng and Friedman, 2005). As colluders can now transfer payments among themselves, the reward mechanism must ensure that the cumulative revenue of the coalition is maximized by the honest reports. The different options discussed above are not relevant for this case, as individual reporting incentives are not important. Instead, the appropriate payments must elicit truthful information from the coalition as a whole.

The results I will describe in this section relate to the literature on (computational) mechanism design, implementation theory, and incentive contracts for principal-(multi)agent settings. The literature on mechanism design (see Jackson, 2003, for a survey) and implementation theory (see Jackson, 2001, for a survey) addresses the design of mechanisms and institutions that satisfy certain properties, given that the agents using the mechanism behave strategically. The main difference between mechanism design and implementation theory is that of multiple equilibria. In the *mechanism design* literature, the goal of the designer is to find the mechanism that has the desired outcome as *an* equilibrium. In the *implementation* literature, on the other hand, the mechanism is required to have *only* the desired equilibria. From this perspective, our results are closer to the implementation theory. More details on the relation between the results of this section and the related work are given in Section 3.7.

For the simplicity of the presentation, I will limit my results from this section to *binary* feedback mechanisms where the set of signals privately observed by the buyers, and the set of feedback values submitted to the reputation mechanism contains only two elements:  $q_0$  and  $q_1$ , or simply 0 and 1. The extension to general mechanisms where the feedback set contains  $M > 2$  elements is conceptually simple, but technically non-trivial. Whenever possible I will give directions for extending the binary results to the general case.

### 3.6.1 Collusion Opportunities in Binary Payment Mechanisms

The proposition below shows that all binary incentive-compatible payment mechanisms using a single reference report have several Nash equilibria (NE). Moreover, at least one lying equilibrium gives the agents higher expected payoffs than the truthful NE.

**Proposition 3.6.1** *Always reporting positive feedback or always reporting negative feedback are both Nash equilibria of any binary incentive-compatible payment mechanisms using a single reference report. One of these equilibria generates a higher payoff than the truthful equilibrium.*

PROOF. The proof follows two steps. First, I prove that in all settings,  $Pr[1|1] \geq Pr[1|0]$ , and  $Pr[0|0] \geq Pr[0|1]$ . Second, I derive the results of the proposition using the definition of payoffs and the incentive-compatibility constraints.

$$\begin{aligned}
Pr[1|1] - Pr[1|0] &= \sum_{\theta \in \Theta} Pr[1|\theta]Pr[\theta|1] - \sum_{\theta \in \Theta} Pr[1|\theta]Pr[\theta|0] \\
&= \sum_{\theta \in \Theta} \frac{Pr[1|\theta]^2 Pr[\theta]}{Pr[1]} - \sum_{\theta \in \Theta} \frac{Pr[1|\theta]Pr[0|\theta]Pr[\theta]}{1 - Pr[1]} \\
&= \sum_{\theta \in \Theta} Pr[1|\theta]Pr[\theta] \frac{Pr[1|\theta] - Pr[1]}{Pr[1](1 - Pr[1])} \\
&= \frac{\sum_{\theta \in \Theta} Pr[1|\theta]Pr[\theta](Pr[1|\theta] - \sum_{t \in \Theta} Pr[1|t]Pr[t])}{Pr[1](1 - Pr[1])};
\end{aligned}$$

However,

$$\begin{aligned}
\sum_{\theta \in \Theta} Pr[1|\theta]Pr[\theta] \left( Pr[1|\theta] - \sum_{t \in \Theta} Pr[1|t]Pr[t] \right) &= \\
&= \sum_{\theta, t \in \Theta, \theta \neq t} (Pr[1|\theta] - Pr[1|t])Pr[\theta]Pr[t](Pr[1|\theta] - Pr[1|t]) \\
&= \sum_{\theta, t \in \Theta, \theta \neq t} (Pr[1|\theta] - Pr[1|t])^2 Pr[\theta]Pr[t] \geq 0
\end{aligned}$$

Similarly,  $Pr[0|0] - Pr[0|1] \geq 0$ .

From the incentive-compatibility constraints:

$$\begin{aligned}
Pr[0|0] \left( \tau(0, 0) - \tau(1, 0) \right) + Pr[1|0] \left( \tau(0, 1) - \tau(1, 1) \right) &> \Lambda(0, 1) > 0 \\
Pr[0|1] \left( \tau(1, 0) - \tau(0, 0) \right) + Pr[1|1] \left( \tau(1, 1) - \tau(0, 1) \right) &> \Lambda(1, 0) > 0
\end{aligned}$$

Because  $Pr[1|1] \geq Pr[1|0]$  and  $Pr[0|0] \geq Pr[0|1]$ , we have  $\tau(0, 0) - \tau(1, 0) > \Lambda(0, 1)$  and  $\tau(1, 1) - \tau(0, 1) > \Lambda(1, 0)$ . Let  $s^{neg} = (0, 0)$  and  $s^{pos} = (1, 1)$  be the constant reporting strategies where an agent always reports 0, respectively 1. It is easy to verify that  $V(s^{neg}, s^{neg}|\cdot) > V(a^{pos}, s^{neg}|\cdot)$  and  $V(s^{pos}, s^{pos}|\cdot) > V(s^{neg}, s^{pos}|\cdot)$ , which makes both  $s^{neg}$  and  $s^{pos}$  be Nash equilibria.

Moreover,

$$\begin{aligned}
V(\bar{s}, \bar{s}|\cdot) &= Pr[0|\cdot]\tau(\cdot, 0) + Pr[1|\cdot]\tau(\cdot, 1) \\
&< \max(\tau(0, 0), \tau(1, 1)) \\
&< \max(V(s^{pos}, s^{pos}|\cdot), V(s^{neg}, s^{neg}|\cdot));
\end{aligned}$$

so at least one of the constant reporting NE strategies generates a higher expected payoff than the truthful equilibrium.  $\blacksquare$



Using several reference reports does not, by default, eliminate the conformity rating lying equilibria, where every agent reports the same thing. In the following result I will show that the optimal binary payment mechanisms based on the feedback submitted by  $N$  agents (i.e.,  $N_{ref} = N - 1$  reference reports) also supports the lying equilibria mentioned in Proposition 3.6.1. However, before doing that, let us introduce some simplifying notation (a summary of the notation in this chapter is presented in the Appendix 3.A).

First, as mentioned above, I assume the reputation mechanism can dispose of  $N$  feedback reports (besides the prior belief  $Pr[\theta]$  and the conditional probabilities  $Pr[1|\theta]$  when designing the payment mechanism. This means that any report can be compared against  $N_{ref} = N - 1$  others for being paid. Second, as the order of reports is not important, the set of  $N_{ref}$  reports can be accurately described by the integer  $n \in \{0, \dots, N - 1\}$  denoting the number of positive reports the other  $N - 1$  agents submitted to the reputation mechanism. This leads to a simplified representation of the payment mechanism by the function  $\tau : Q_2 \times \{0, \dots, N - 1\} \rightarrow \mathbb{R}^+$ , where  $\tau(r_i, n)$  is the payment to agent  $i$  when she reports  $r_i \in Q_2$  and  $n$  out of the other  $N - 1$  agents report positively. Third, let  $S = \{(s(0), s(1)) | s(0), s(1) \in Q_2\}$ , be the set of pure reporting strategies of one agent, where  $s = (s(0), s(1))$  denotes the strategy according to which the buyer announces  $s(0) \in Q_2$  when she observes low quality (i.e., 0), and  $s(1) \in Q_2$  when she observes high quality (i.e., 1). To ease the notation, we name the four members of the set  $S$  as:

- the honest strategy,  $\bar{s} = (0, 1)$ ,
- the lying strategy,  $s^{lie} = (1, 0)$ ,
- the always reporting positive feedback strategy,  $s^{pos} = (1, 1)$ , and
- the always reporting negative feedback strategy,  $s^{neg} = (0, 0)$ .

Forth, a strategy profile  $s$  is a vector  $(s_i)_{i=1, \dots, N}$ , prescribing the reporting strategy  $s_i \in S$  for each of the  $N$  agents. We will sometimes use the notation  $s = (s_i, s_{-i})$ , where  $s_{-i}$  is the strategy profile for all agents except  $i$ ; i.e.,  $s_{-i} = (s_j)$ , for  $j = 1, \dots, i - 1, i + 1, \dots, N$ .

Fifth, given the profile of reporting strategies  $(s_i, s_{-i})$ , let  $\mu[n, s_{-i}]$  describe the belief of agent  $i$  regarding the distribution of the reference reports, when:

- $n$  out of the other  $N - 1$  agents observe the high quality signal, 1
- the other  $N - 1$  agents are reporting according to the strategy profile  $s_{-i}$ ;

Given  $n$  and  $s_{-i}$ , agent  $i$  believes with probability  $\mu[n, s_{-i}](x)$  that  $x$  reference reports are positive.

Finally, I will assume that reporting feedback does not cost anything ( $C_r = 0$ ) and that  $\Lambda$  is an upper bound on all external benefits an agent may obtain by lying. Under these new assumptions the incentive-compatibility constraints become:

$$\begin{aligned} \sum_{n=0}^{N-1} Pr[n|1] (\tau(1, n) - \tau(0, n)) &\geq \Lambda; \\ \sum_{n=0}^{N-1} Pr[n|0] (\tau(0, n) - \tau(1, n)) &\geq \Lambda; \end{aligned} \tag{3.18}$$

and the expected payment to an honest reporter is:

$$E[V(\bar{s}_i, \bar{s}_{-i})] = Pr[1] \sum_{n=0}^{N-1} Pr[n|1] \tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0] \tau(0, n); \tag{3.19}$$

The optimization problem defining the payment mechanism that minimizes the expected payment to an honest reporter is given by LP 3.4.1 and has the following simplified form:

**LP 3.6.1**

$$\begin{aligned}
\min \quad & E[V(\bar{s}_i, \bar{s}_{-i})] = Pr[1] \sum_{n=0}^{N-1} Pr[n|1] \tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0] \tau(0, n); \\
\text{s.t.} \quad & \sum_{n=0}^{N-1} Pr[n|1] (\tau(1, n) - \tau(0, n)) \geq \Lambda; \\
& \sum_{n=0}^{N-1} Pr[n|0] (\tau(0, n) - \tau(1, n)) \geq \Lambda; \\
& \tau(0, n), \tau(1, n) \geq 0; \forall n = \{0, 1, \dots, N-1\};
\end{aligned}$$

The proposition below characterizes the solution of LP 3.6.1, and reveals the vulnerability of the resulting payment mechanism to lying coalitions:

**Proposition 3.6.2** *The incentive-compatible payment scheme that minimizes the expected payment to an honest reporter (defined by LP 3.6.1) is:*

$$\begin{aligned}
\tau(0, n) &= 0, \forall n \neq 0; \quad \tau(1, n) = 0, \forall n \neq N-1 \\
\tau(0, 0) &= \Lambda \frac{Pr[N-1|0] + Pr[N-1|1]}{Pr[N-1|1]Pr[0|0] - Pr[N-1|0]Pr[0|1]}; \\
\tau(1, N-1) &= \Lambda \frac{Pr[0|0] + Pr[0|1]}{Pr[N-1|1]Pr[0|0] - Pr[N-1|0]Pr[0|1]};
\end{aligned}$$

Before starting the proof, I will use the following lemma:

**Lemma 3.6.1** *Given any set of types  $\Theta$ , probability distributions  $Pr[1|\theta]$ , prior belief over types  $Pr[\theta]$  and number of agents  $N$ , we have  $\frac{Pr[n|1]}{Pr[n|0]} < \frac{Pr[n+1|1]}{Pr[n+1|0]}$  for all  $n = 0 \dots N-1$ .*

PROOF. see Appendix 3.D ■

PROOF.(Proposition 3.6.2)

Let us write the corresponding dual problem:

$$\begin{aligned}
\max \quad & \Lambda y_0 + \Lambda y_1; \\
\text{s.t.} \quad & Pr[n|0]y_0 - Pr[n|1]y_1 \leq Pr[0]Pr[n|0] \\
& Pr[n|1]y_1 - Pr[n|0]y_0 \leq Pr[1]Pr[n|1] \\
& \forall n \in \{0, \dots, N-1\};
\end{aligned}$$

where  $y_0$  (respectively  $y_1$ ) is the dual variable corresponding to the constraint where the agent observes 0 (respectively 1). By dividing the first set of constraints by  $Pr[n|0]$  and the second set of constraints by  $Pr[n|1]$ , we have:

$$\begin{aligned}
y_0 - y_1 Pr[n|1]/Pr[n|0] &\leq Pr[0], \forall n \in \{0, \dots, N-1\}; \\
y_1 - y_0 Pr[n|0]/Pr[n|1] &\leq Pr[1], \forall n \in \{0, \dots, N-1\};
\end{aligned}$$

Clearly, among the  $2(N - 1)$  constraints of the dual problem, only two are active, corresponding to:  $n_1 = \arg \min_n \frac{Pr[n|1]}{Pr[n|0]}$ , and  $n_2 = \arg \min_n \frac{Pr[n|0]}{Pr[n|1]}$ . By Lemma 3.6.1 we know that  $n_1 = 0$  and  $n_2 = N - 1$ . Therefore, the only two variables of LP 3.6.1 that have non-zero values are  $\tau(0, 0)$  and  $\tau(1, N - 1)$ , which satisfy the linear equations:

$$\begin{aligned} Pr[0|0]\tau(0, 0) - Pr[N - 1|0]\tau(1, N - 1) &= \Lambda; \\ -Pr[0|1]\tau(0, 0) + Pr[N - 1|1]\tau(1, N - 1) &= \Lambda; \end{aligned}$$

hence the result of the proposition. ■

Intuitively, the optimal payment mechanism rewards the consensus between all reporters, and it is exactly this property that makes it vulnerable to collusion. Adding supplementary constraints to the design problem (an idea I will investigate in the following subsections) or changing the objective function<sup>10</sup> may prevent lying equilibria where every agent reports the same signal. Nevertheless, all payments that satisfy the incentive compatibility constraints have the following property: there must be at least two values of the reference reports,  $n_1 < n_2$ , such that:

$$\begin{aligned} \tau(0, n_1) &> \tau(1, n_2) \\ \tau(1, n_2) &> \tau(0, n_2); \end{aligned}$$

This property will prove essential in Section 3.6.4 where I will prove that all incentive-compatible payment mechanisms also accept lying equilibria.

### 3.6.2 Automated Design of Collusion-resistant Payments

Using several reference reports does not make incentive compatible payments, by default, collusion-resistant. The result of Proposition 3.6.2 shows that the incentive-compatible constraints alone, also generate reward schemes that are vulnerable to conformity rating (i.e, everybody reports the same thing). In most cases, nevertheless, payment schemes based on several reference reports are not *constrained* to reward agreement, so one could specify further conditions, which added to the design problem generate collusion-resistant mechanisms.

On the other hand,  $N$  agents have more colluding strategies than just 2 agents. Deterring all possible coalitions requires a careful analysis of the options available to the reporters, and the consequent choice of constraints added to the design problem. Once the constraints have been specified, my approach for designing the actual payments is mostly numerical, and based on *automated mechanism design* (Conitzer and Sandholm, 2002): I define the appropriate mechanism by an optimization problem that minimizes the expected budget of the reputation mechanism under the constraints of incentive-compatibility and collusion resistance. Nevertheless, I also consider the complexity of the design problem, and whenever necessary, I give alternatives that trade performance for computational efficiency. Since the payment mechanism must be (re)computed every time the information in the system changes, I believe it is a strong practical requirement to have fast design algorithms.

The ideal reward mechanism deters any coalition, no matter how big, even when every colluder may use a different strategy and side-payments are possible. Such a mechanism, unfortunately, is trivially impossible: given that all agents may collude and use side-payments to subsidize the agents that might otherwise quit the coalition, the payment mechanism doesn't have any leverage to encourage honest reporting. Whatever the payment scheme, the coalition will adopt the strategy that maximizes the total revenue, regardless of the truth.

<sup>10</sup> One might wish, for example, to design a mechanism that minimizes the expected budget paid to all  $N$  buyers. In this case, the objective function of the problem LP 3.6.1 is:  $B = \sum_{n=0}^N Pr[n](n \cdot \tau(1, n - 1) + (N - n) \cdot \tau(0, n))$ , where  $Pr[n]$  is the prior probability that  $n$  out of  $N$  buyers observe high quality;

Positive results may be obtained only by imposing further restrictions on possible lying coalitions. The first restriction is that not all agents can collude. Some agents are altruistic in nature and report honestly for moral or social reasons. Other agents are not aware of collusion opportunities, or cannot be contacted by a forming coalition. Social or legal norms against collusion may furthermore create prejudices that deter some agents from entering the coalition.

The second restriction addresses the complexity of the coordination among colluders. Symmetric collusion strategies prescribe that all colluders are reporting according to the same strategy. The coordination on symmetric strategies is very simple, and requires one anonymous access to a publicly available source of information that specifies the colluding strategy. Intuitively, the role of the coordination device may be played by a public blog which analyzes the mechanisms and informs potential colluders on the most profitable symmetric colluding strategy. Asymmetric collusion strategies, on the other hand, require significantly more complex coordination. Since every colluder may use a different reporting strategy, the coordination device must know the identity of the colluder before instructing on a collusion strategy. This is often unfeasible, either because colluders might not want to reveal their identity and thus create a trace of their misbehavior, or because identity of the colluders cannot be known at all before the actual reporting takes place.

The third restriction addresses the availability of side-payments between colluders (or transferable utilities). Even when the rewards offered by the reputation mechanism are monetary, the kind of micro-payments that would be required among the colluders are difficult and expensive to implement. Side-payments are even less feasible when the rewards offered by the reputation mechanism are in kind, or in some currency under the control of the reputation mechanism (e.g., Yahoo points or Slashdot karma cannot be transferred even if users wanted to). The conversion of such subjective resources to real money that can afterwards be transferred is even more difficult than the transfer itself.

One notable exception where side-payments are feasible is when the same strategic entity controls a number of online identities, or “sybils” (Cheng and Friedman, 2005). Here, the controlling agent is interested in maximizing his overall revenue (i.e., the sum of the revenues obtained by the sybils), so side-payments do not have to physically occur<sup>11</sup>.

To summarize, we address collusion scenarios where:

- all or only some of the agents can become part of a lying coalition,
- colluders can coordinate or not on using different strategies,
- colluders can make or not side-payments to other colluders.

From the remaining seven restricted collusion scenarios (see Table 3.5) I am only addressing five. I will not consider the settings where all or some of the agents may make side-payments among themselves, but can only collude on symmetric strategies. As discussed in the previous paragraph, transferable utilities should mostly be assumed for a group of online identities controlled by the same strategic agent. Restricting the agent from coordinating its sybils on asymmetric strategies seems to us unreasonable.

For all scenarios involving non-transferable utilities, collusion resistance can emerge as a consequence of having honest reporting as the only (or the most attractive) equilibrium. When all agents may collude, an honest reporting dominant equilibrium is impossible. Therefore, I will resort to designing reward schemes where honest reporting is a unique Nash equilibrium, or the Pareto-optimal Nash equilibrium.

<sup>11</sup>Whenever rewards are non-monetary, the overall utility of the controlling agent is usually less than the sum of utilities of the sybils. On Slashdot, for example, ten users with bad karma are not worth as one user with good karma. Nevertheless, we will keep for simplicity the assumption of additive utilities for the controlling agent.

	Non-Transferable Utilities		Transferable Utilities	
	symmetric strategies	asymmetric strategies	symmetric strategies	asymmetric strategies
all agents collude	Section 3.6.3	Section 3.6.4	unreasonable assumption	impossible
some agents collude	Section 3.6.5	Section 3.6.6	unreasonable assumption	Section 3.6.7

**Table 3.5:** Different collusion scenarios.

When only a fraction of the agents may collude (non-colluders are assumed to report honestly) I also consider designing rewards that make honest reporting the dominant strategy for the colluders.

I will restrict the analysis to *pure* reporting strategies and pure strategy equilibria. The reason behind this choice is grounded in practical considerations: mixed strategies and mixed strategy equilibria are more complex and difficult to understand, and therefore unlikely to be observed in practical applications. Acknowledging the limitations brought in by this assumption, I still believe these results are valuable for a number of practical scenarios. The following subsections address each one collusion scenario, and describe possible methods for designing collusion-resistant, incentive-compatible reward mechanisms.

### 3.6.3 Full Coalitions on Symmetric Strategies, Non-Transferable Utilities

I assume that agents (i) can only coordinate once (before any of them purchases the product) on the same (pure) reporting strategy, and (ii) cannot make side-payments from one to another. This simple form of coordination between colluders considerably simplifies the problem of the mechanism designer; the only supplementary constraint on the incentive-compatible payment mechanism is to ensure that none of the pure symmetric strategy profiles is a NE.

The set of pure strategies is finite (and contains 3 *lying* strategies) therefore we can exhaustively enumerate the constraints that prevent the corresponding symmetric lying strategy profiles to be NE:

- $s^{pos}$  (always reporting 1) is not NE when a rational agent would rather report 0 instead of 1 given that all other agents follow  $s^{pos}$ :

$$\tau(0, N - 1) > \tau(1, N - 1); \quad (3.20)$$

- $s^{neg}$  (always reporting 0) is not NE when a rational agent would rather report 1 instead of 0 given that all other agents follow  $s^{neg}$ ;

$$\tau(1, 0) > \tau(0, 0); \quad (3.21)$$

- $s^{lie}$  is not NE when at least one agent (either observing 1 or 0) would rather report the truth. Given that other agents always lie,  $N - 1 - n$  reference reports will be positive whenever  $n$  high quality signals were actually observed:

$$\begin{aligned}
& \text{either } \sum_{n=0}^{N-1} Pr[n|0](\tau(0, N-1-n) - \tau(1, N-1-n)) > 0; \\
& \text{or } \sum_{n=0}^{N-1} Pr[n|1](\tau(1, N-1-n) - \tau(0, N-1-n)) > 0;
\end{aligned} \tag{3.22}$$

The objective function from Eq. (3.19), and the constraints (3.18), (3.20), (3.21) and (3.22) define the optimal incentive-compatible payment mechanism that is also collusion-resistant in the sense explained in the beginning of this section (i.e., honest reporting is the unique pure-strategy symmetric NE). To compute the payments, the mechanism designer must solve two linear optimization problems, one corresponding to each branch of the constraint (3.22).

**Proposition 3.6.3** *Collusion-resistant, incentive-compatible rewards require minimum  $N = 4$  agents.*

PROOF. From (3.21), (3.20) and (3.18) we know that there must be  $n_1 \in \{1, \dots, N-1\}$  and  $n_2 \in \{0, \dots, N-2\}$  such that  $n_1 < n_2$ ,  $\tau(0, n_1) > \tau(1, n_1) = 0$  and  $\tau(1, n_2) > \tau(0, n_2) = 0$ . Clearly, this may happen only when  $N \geq 4$ . ■

Taking the plumber example described in Section 3.3.1 and  $N = 4$  agents, the conditional distribution of the reference reports and the optimal collusion-resistant, incentive-compatible payment mechanism are the following:

$Pr[\cdot \cdot]$	0	1	2	3	$\tau(\cdot, \cdot)$	0	1	2	3
0	0.4179	0.2297	0.1168	0.2356	0	0	12.37	0	$\varepsilon$
1	0.0255	0.0389	0.2356	0.7	1	$\varepsilon$	0	6.29	0

where  $\varepsilon$  is a small positive value, and the guaranteed margin for truth-telling is  $\Lambda = 1$ . For any  $N > 4$  the payment mechanism looks the same and rewards a report if all but one of the other agents agree with the submitted report. In the same time, opposing consensus is rewarded by a small amount  $\varepsilon$ .

Payments with exactly the same structure represent a general solution of the design problem in this context. Moreover, the payments always exist:

**Proposition 3.6.4** *Given any set of types  $\Theta$ , probability distributions  $Pr[1|\theta]$ , prior belief over types  $Pr[\theta]$ , and number of agents  $N \geq 4$ , the following payment system has honest reporting as the unique symmetric NE:*

$$\begin{aligned}
& \tau(0, n) = 0, \forall n \neq 1, N-1; \quad \tau(1, n) = 0, \forall n \neq 0, N-2; \quad \tau(0, N-1) = \tau(1, 0) = \varepsilon \\
& \tau(0, 1) = \begin{cases} \Lambda \frac{Pr[1|1]}{Pr[1|0]Pr[1|1] - Pr[N-2|0]Pr[N-2|1]} & \text{if condition A} \\ \Lambda \frac{Pr[1|0]}{Pr[N-2|0]Pr[N-2|1] - Pr[1|0]Pr[1|1]} & \text{if condition B} \\ \Lambda \frac{Pr[N-2|1] + Pr[N-2|0]}{Pr[1|0]Pr[N-2|1] - Pr[N-2|0]Pr[1|1]} & \text{otherwise} \end{cases} \\
& \tau(1, N-2) = \begin{cases} \Lambda \frac{Pr[N-2|1]}{Pr[1|0]Pr[1|1] - Pr[N-2|0]Pr[N-2|1]} & \text{if condition A} \\ \Lambda \frac{Pr[N-2|0]}{Pr[N-2|0]Pr[N-2|1] - Pr[1|0]Pr[1|1]} & \text{if condition B} \\ \Lambda \frac{Pr[1|1] + Pr[1|0]}{Pr[1|0]Pr[N-2|1] - Pr[N-2|0]Pr[1|1]} & \text{otherwise} \end{cases}
\end{aligned}$$

$$A = \bigwedge \left( \begin{array}{l} Pr[1|0]Pr[1|1] > Pr[N-2|0]Pr[N-2|1] \\ Pr[N-2|1] > Pr[1|1] \\ Pr[N-2|1]^2 - Pr[1|1]^2 > Pr[1|0]Pr[1|1] - Pr[N-2|0]Pr[N-2|1] \end{array} \right)$$

$$B = \bigwedge \left( \begin{array}{l} Pr[N-2|0]Pr[N-2|1] > Pr[1|0]Pr[1|1] \\ Pr[1|0] > Pr[N-2|0] \\ Pr[1|0]^2 - Pr[N-2|0]^2 > Pr[N-2|0]Pr[N-2|1] - Pr[1|0]Pr[1|1] \end{array} \right)$$

PROOF. It is straight-forward to check that for  $\varepsilon$  small enough, the payments described in the proposition verify the constraints (3.18), (3.20) and (3.21). Moreover, these payments minimize the expected payment to an honest reporter. ■

A less strict notion of collusion resistance requires honest reporting to be the Pareto-optimal NE. The intuition is that any stable (i.e., equilibrium) coalition will necessarily make some colluders worse off than in the honest equilibrium. Assuming non-transferable utilities, colluders that benefit from the coalition cannot subsidize the ones that make a loss, and hopefully, the latter will refuse to join the coalition in the first place.

The payment mechanism that has honest reporting as the Pareto-optimal equilibrium solves the following optimization problem:

#### LP 3.6.2

$$\begin{aligned} \min \quad & E[V(\bar{s}_i, \bar{s}_{-i})] = Pr[1] \sum_{n=0}^{N-1} Pr[n|1]\tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0]\tau(0, n); \\ \text{s.t.} \quad & \sum_{n=0}^{N-1} Pr[n|1](\tau(1, n) - \tau(0, n)) \geq \Lambda; \\ & \sum_{n=0}^{N-1} Pr[n|0](\tau(0, n) - \tau(1, n)) \geq \Lambda; \\ & \tau(1, N-1) < E[V(\bar{s}_i, \bar{s}_{-i})]; \\ & \tau(0, 0) < E[V(\bar{s}_i, \bar{s}_{-i})]; \\ & \bigvee \left( \begin{array}{l} \sum_{n=0}^{N-1} Pr[n|0](\tau(0, N-1-n) - \tau(1, N-1-n)) > 0 \\ \sum_{n=0}^{N-1} Pr[n|1](\tau(1, N-1-n) - \tau(0, N-1-n)) > 0 \\ E[V(s_i^{lie}, s_{-i}^{lie})] < E[V(\bar{s}_i, \bar{s}_{-i})]; \end{array} \right) \\ & \tau(0, n), \tau(1, n) \geq 0; \forall n = \{0, 1, \dots, N-1\}; \end{aligned}$$

The first two constraints make honest reporting a Nash equilibrium. The next two constraints prevent lying colluders on  $s^{pos}$  or  $s^{neg}$  to get higher rewards than in the honest equilibrium. These constraints are always easier to satisfy than the constraints (3.20), (3.21) which prevent equilibria on  $s^{pos}$  and  $s^{neg}$ . The last constraint requires that the symmetric profile where every agent lies is either not a NE, or that it generates an expected payoff lower than the honest equilibrium. The expected payoff to an agent reporting according to  $s^{lie}$  when everybody else reports according to  $s^{lie}$  is:

$$E[V(s_i^{lie}, s_{-i}^{lie})] = \sum_{n=0}^{N-1} \left( Pr[n|0]Pr[0]\tau(1, N-1-n) + Pr[n|1]Pr[1]\tau(0, N-1-n) \right);$$

Note that for coalitions on  $s^{lie}$  it is not sufficient to limit the payoff to a colluder below the expected payoff in the honest equilibrium (the third inequality of the last disjunctive constraint of LP 3.6.2). Numerical simulations performed on randomly generated problems (Appendix 3.E describes the details on how the problems were generated) show that for 40- 50% of the problems the collusion-resistant payments are cheaper by eliminating altogether the symmetric lying equilibrium on  $s^{lie}$ : i.e., either the first, or the second inequality from the last constraint of LP 3.6.2 are easier to satisfy than the third inequality.

In either case, the resulting optimal payments have the following structure:

- $\tau(0, 0) = \tau(1, N - 1) = E[V(\bar{s}_i, \bar{s}_{-i})] - \varepsilon$ . These values prevent the lying coalitions on  $s^{pos}$  or  $s^{neg}$  to Pareto-dominate the honest reporting equilibrium;
- $\tau(0, 1) > 0$  and  $\tau(1, N - 2) > 0$  are scaled to satisfy the incentive-compatibility constraints, and the “easiest” of the three inequalities that prevent a coalition on  $s^{lie}$ ;
- the other payments are 0.

For the plumber example in Section 3.3.1 the payments are the following:

$\tau(\cdot, \cdot)$	0	1	2	3
0	1.30	4.52	0	0
1	0	0	1.26	1.30

### 3.6.4 Full Coalitions on Asymmetric Strategies, Non-Transferable Utilities

In the next collusion scenario all  $N$  agents can coordinate on asymmetric collusion strategies, without being able to make side-payments from one to another. Each of the  $N$  agents can have a different reporting strategy, and the collusion strategy profile is denoted by  $s = (s_i)$ ,  $i = 1, \dots, N$ , where  $s_i \in S$  is the reporting strategy of agent  $i$ .

I distinguish between two cases, when the communication (and therefore the coordination on the collusion strategy profile) happens *before* or *after* the agents perceive the quality signals from the product they purchase. In the former case, no payment scheme can satisfy the incentive-compatibility constraints. In the latter case, honest reporting can never be a unique Nash equilibrium of the mechanism; however, honest reporting can be the Pareto-optimal Nash equilibrium.

**Proposition 3.6.5** *When agents communicate and coordinate their reports after perceiving the quality signals, strict incentive-compatible payment mechanisms do not exist.*

PROOF. Consider two settings, that are identical except for the observation of agent  $i$ . In setting  $I$ , agent  $i$  observes  $o_i = 0$ , in setting  $II$ , agent  $i$  observes  $o_i = 1$ ; in both settings the other agents observe  $n$  high quality signals. An incentive-compatible mechanism requires  $i$  to report 0 in setting  $I$ , and 1 in setting  $II$ . Assume all other agents report truthfully; during the communication phase (happening after signals have been perceived) agent  $i$  learns in both settings that the reference reports contain  $n$  positive reports. An incentive-compatible payment mechanism requires that:

- $\tau(0, n) > \tau(1, n)$  - honest reporting is strictly better for  $i$  in setting  $I$  ;



- $\tau(1, n) > \tau(0, n)$  - honest reporting is strictly better for  $i$  in setting  $II$ ;

Clearly this is impossible. ■

The previous proposition formalizes the intuition that truth-telling may only be an ex-ante Nash equilibrium. The reference reports must be unknown to the agent in order to allow the design of incentive-compatible payments.

When the communication takes place before the agents observe the signals, incentive-compatible payments do exist, but always accept other Nash equilibria where agents lie:

**Proposition 3.6.6** *When agents communicate and coordinate their reports before perceiving the quality signals, no payment mechanism has a unique honest reporting Nash equilibrium.*

PROOF. The proof shows that a full coalition can always find a profile of constant reporting strategies,  $s = (s_i)$ ,  $i = 1, \dots, N$ ,  $s_i \in \{s^{neg}, s^{pos}\}$  that is a NE.

Let  $s(n) = (s_i)$  be the family of reporting strategy profiles where  $n$  out of  $N$  agents always report 1, and the other  $N - n$  agents always report 0: i.e.,

$$\begin{aligned} s_i &= s^{pos}, \forall i \in A_1; & s_i &= s^{neg}, \forall i \in A_0; \\ |A_1| &= n, & |A_2| &= N - n; \\ A_1 \cap A_0 &= \emptyset; & A_1 \cup A_0 &= \{1, 2, \dots, N\}; \end{aligned} \tag{3.23}$$

Assume that the payment mechanism defined by  $\tau(\cdot, \cdot)$  accepts honest reporting as the unique NE. We have seen in Section 3.6.1 that the incentive-compatible constraints (3.18) imply the existence of  $n_1 < n_2 \in \{0, 1, \dots, N - 1\}$  such that  $\tau(0, n_1) > \tau(1, n_1)$ , and  $\tau(1, n_2) > \tau(0, n_2)$ .

With non-transferable utilities, the strategy profile  $s(n_2 + 1)$  is not a NE if and only if one of the  $n_2 + 1$  agents that should report 1 would rather report 0:

$$\tau(0, n_2) > \tau(1, n_2);$$

or one of the  $N - n_2 - 1$  agents that should report 0 would rather report 1:

$$\tau(1, n_2 + 1) > \tau(0, n_2 + 1);$$

The first inequality cannot be true by the choice of  $n_2$ ; therefore, it must be that  $\tau(1, n_2 + 1) > \tau(0, n_2 + 1)$ .

Similarly,  $s(n_2 + 2)$  is not a NE iff either  $\tau(0, n_2 + 1) > \tau(1, n_2 + 1)$  (impossible), or  $\tau(1, n_2 + 2) > \tau(0, n_2 + 2)$ . Continuing this argument we find that  $\tau(1, N - 1) > \tau(0, N - 1)$  which makes  $s(N)$  (i.e., all agents report 1) a Nash equilibrium. Hence the result of the proposition. ■

Proposition 3.6.6 holds regardless the number of reports,  $N$ , available to the reputation mechanism. The proof shows that all incentive-compatible reward schemes have the property that for some  $n \in \{0, \dots, N - 1\}$ , either  $\tau(1, n) > 0$  and  $\tau(1, n + 1) < \tau(0, n + 1)$ , or  $\tau(1, N - 1) > 0$ . In the first case, the coalition can adopt the lying strategy where  $n + 1$  agents always report 1, and  $N - n - 1$  agents always report 0. The structure of the payments makes such a coalition stable, as no agent finds it profitable to deviate from the coalition. In the second case the payment scheme is vulnerable to everybody always reporting 1.

While lying equilibria always exist in this scenario, they do not necessarily *Pareto-dominate* the honest reporting NE. Take for example the incentive-compatible payments that solve LP 3.6.1, with the additional constraints that  $\tau(0, 0) = 0$  and  $\tau(1, N - 1) = 0$ . A stable coalition can form on the strategy profiles  $s(n_2 + 1)$  or  $s(n_1)$ , where  $n_2 + 1$  (respectively  $n_1$ ) agents always report 1 and the others always report 0 regardless of their observation. This equilibrium, however, does not Pareto-dominate the truthful one: the agents that report 0 do not get any reward, whereas they do get rewarded in the honest equilibrium.

The payment mechanism can be further improved by setting  $\tau(0, n_1 - 1) = \tau(1, n_2 + 1) = \varepsilon$ , where  $\varepsilon$  is some small positive value. This modification eliminates the equilibria  $s(n_2 + 1)$  and  $s(n_1)$  and instead introduces the equilibria  $s(n_2 + 2)$  and  $s(n_1 - 1)$ . Both these equilibria are extremely unattractive (some agents get paid  $\varepsilon$ , while others don't get paid at all) and are dominated by the honest equilibrium.

**Proposition 3.6.7** *Given the set of types  $\Theta$ , the conditional probabilities  $Pr[1|\theta]$ , the prior belief over types  $Pr[\theta]$ , and  $N = 4$  agents, the following payment scheme has honest reporting as the Pareto-optimal Nash equilibrium:*

$\tau(\cdot, \cdot)$	0	1	2	3
0	$\varepsilon$	$x > 0$	0	0
1	0	0	$y > 0$	$\varepsilon$

The values  $x$  and  $y$  depend on the probabilities  $Pr[1|\theta]$  and  $Pr[\theta]$ , and  $\varepsilon$  has a small positive value.

PROOF. The payments here are similar to those of Proposition 3.6.4 except that consensus is rewarded with some small amount  $\varepsilon$  instead of being discouraged. In this way the mechanism has only three NE: honest reporting, always reporting 1 or always reporting 0. Both lying equilibria, however, generate much lower revenues (assuming, of course, that  $\varepsilon$  is small enough); therefore, honest reporting is the Pareto-optimal equilibrium. The proof that the mechanism has only 3 NE is based on brute force: for  $x$  and  $y$  taking the values specified in Proposition 3.6.4, we verify that no other strategy profile is a NE. The details are presented in Appendix 3.F. ■

For general reward mechanisms based on  $N > 4$  reports, honest reporting can become the Pareto-optimal NE by considering all lying strategy profiles,  $s$ , and adding to the design problem either of the following linear constraints:

$$\begin{aligned} V(s_i, s_{-i}|o_i) &< V(s_i^*, s_{-i}|o_i) \text{ for some } i, o_i \text{ and } s_i^*; \\ E[V(s_i, s_{-i})] &< E[V(\bar{s}_i, \bar{s}_{-i})] \text{ for some } i; \end{aligned} \tag{3.24}$$

The first constraint ensures that the strategy profile  $s$  is not a NE, and consists of a disjunction of at most 8 linear inequalities: for each reporting strategy  $s_i = (s_i(0), s_i(1)) \in S$ , the agent reporting according to  $s_i$  has the incentive to deviate either when observing 0, or when observing 1. There are four strategies in  $S$  and only one possible deviation for each observed signal, hence the 8 inequalities. The second constraint ensures that  $s$  does not Pareto-dominate the honest equilibrium, and consists of a similar disjunction of at most 4 inequalities. Note that any two strategy profiles that represent different permutations of the same set of  $N$  reporting strategies will generate the same constraints.

Therefore, there are  $\binom{N+3}{3}$  different constraints imposing honesty as a Pareto-optimal NE, each consisting of a disjunction of at most 12 linear inequations. The resulting optimization problem is a disjunctive linear program which can be transformed into a mixed integer linear program (Sherali and Shetty, 1980).

Unfortunately, the complexity of the resulting optimization problem is exponential in the number  $N$  of reporters considered by the payment mechanism. Since the payment mechanism depends on the current belief over the types  $\theta$ , the reputation mechanism might be required to frequently update the payments in order to reflect the changing beliefs. For large values of  $N$  this is clearly infeasible.

I therefore consider a special family of payment mechanisms that can be designed efficiently to make honest reporting the Pareto-optimal NE. The basic idea is to consider payments similar to those from Proposition 3.6.7, that reward a report only when all but one of the reference reports agree. Consensus on the positive or negative feedback is also rewarded by a small amount  $\varepsilon$ , but all other payments are zero:

$\tau(\cdot, \cdot)$	0	1	2...N-3	N-2	N-1
0	$\varepsilon$	$x > 0$	0...0	0	0
1	0	0	0...0	$y > 0$	$\varepsilon$

**Figure 3.9:** Payment mechanism for  $N > 4$  agents.

The payment mechanism now depends on only 2 parameters,  $x$  and  $y$  that must be scaled to prevent any other lying strategy profile to become a NE Pareto-dominating the honest equilibrium. Note that no strategy profile where more than one agent reports according to  $s^{pos}$  or  $s^{neg}$  can become a successful collusion strategy. If at least two agents always report 1, none of the other agents will ever want to report 0 (as  $\tau(0, n) = 0$  for any  $n \geq 2$ ). Similarly if at least two agents always report 0, none of the other agents will ever want to report 1. Nevertheless, both consensus equilibria yield very small payoffs, significantly lower than the payoff of the honest reporting equilibrium.

Following the intuition from the proof of Proposition 3.6.7, many of the remaining lying strategy profiles cannot be a NE regardless of the values of  $x$  and  $y$ . Let us consider the set of potential lying equilibrium strategy profiles:

$$\tilde{S} = \{(n_0 \times s^{neg}, n_1 \times s^{pos}, \bar{n} \times \bar{s}, n_l \times s^{lie}) \mid n_0 + n_1 + \bar{n} + n_l = N\}; \quad (3.25)$$

where  $n_0 \in \{0, 1\}$  agents always report 0,  $n_1 \in \{0, 1\}$  agents always report 1,  $\bar{n} \notin \{N-1, N\}$  agents report honestly and  $n_l$  agents always lie. The cardinality of  $\tilde{S}$  is  $4(N-1)$ . The profile  $s \in \tilde{S}$  is a NE if and only if for any strategy  $s_i \in s$ , the agent reporting according to  $s_i$  does not have the incentive to deviate to another reporting strategy given that all other agents keep reporting according to  $s_{-i}$ . Let  $o_i \in Q_2$  be the signal observed by agent  $i$ . The report prescribed by strategy  $s_i$  is  $r_i = s_i(o_i) \in Q_2$ , and given that  $\varepsilon$  is small enough to be ignored, the expected payoff to agent  $i$  is:

$$\begin{aligned} &Pr[1|o_i, s_{-i}] \cdot x \text{ if } r_i = 0 \\ &Pr[N-2|o_i, s_{-i}] \cdot y \text{ if } r_i = 1 \end{aligned}$$

where  $Pr[1|o_i, s_{-i}]$  and  $Pr[N-2|o_i, s_{-i}]$  are the probabilities that exactly 1, respectively  $N-2$  of the other  $N-1$  agents will report positively given the observation  $o_i$  and the strategy profile  $s_{-i}$ .

The deviation to reporting  $1 - r_i$  is not profitable for some observation  $o_i$  if and only if:

$$\begin{aligned} &Pr[1|o_i, s_{-i}] \cdot x - Pr[N-2|o_i, s_{-i}] \cdot y > 0 \text{ if } r_i = 0 \\ &Pr[N-2|o_i, s_{-i}] \cdot y - Pr[1|o_i, s_{-i}] \cdot x > 0 \text{ if } r_i = 1 \end{aligned}$$

The conditions that make  $s$  a NE can therefore be expressed as a set of at most 8 inequalities with the following structure:

$$\begin{aligned} &a_j x - b_j y > 0; \quad a_j, b_j > 0 \\ &-a_k x + b_k y > 0; \quad a_k, b_k > 0 \end{aligned}$$

If  $\max_j \frac{b_j}{a_j} > \min_k \frac{b_k}{a_k}$  the above system of inequations is infeasible, so that for any positive values of  $x$  and  $y$  the corresponding strategy profile  $s$  can not be a NE. However, when  $\max_j \frac{b_j}{a_j} < \min_k \frac{b_k}{a_k}$ , there are values of  $x$  and  $y$  which can make  $s$  a NE, and therefore, in the design problem we must specify a constraint that prevents  $s$  from Pareto-dominating the honest NE. The corresponding constraint will be a disjunction of inequalities: 2 for restricting  $x$  and  $y$  to values that do not make  $s$  a NE, at most 3 that limit the expected payments of colluders below the expected payment of the honest equilibrium.

Since there are  $4(N - 1)$  potential lying strategy profiles, the optimization problem defining  $x$  and  $y$  can have up to  $4N - 2$  constraints: 2 linear incentive-compatible constraints and up to  $4(N - 1)$  disjunctive linear constraints. The transformation to a mixed integer linear program involves adding up to  $4(N - 1)$  integer variables, which in the worst case, can result in exponential-time (in  $N$ ) complexity of the design problem.

Fortunately, most of the strategy profiles in  $\tilde{S}$  can be eliminated analytically. It turns out that the payment mechanism from Figure 3.9 does not accept as a Nash Equilibrium any strategy profile where at least one agent reports truthfully and another agent reports according to  $s^{lie}$ :

**Proposition 3.6.8** *Let  $s = (n_0 \times s^{neg}, n_1 \times s^{pos}, \bar{n} \times \bar{s}, n_l \times s^{lie}) \in \tilde{S}$  be a strategy profile where  $n_0$  agents always report 0,  $n_1$  agents always report 1,  $\bar{n}$  agents report honestly and  $n_l$  agents always lie. If  $(\bar{n} \neq 0 \wedge n_l \neq 0)$  or  $(n_0 = 1 \wedge n_1 = 1)$ ,  $s$  cannot be a Nash equilibrium of the payment mechanism described in Figure 3.9.*

PROOF. For the reasons explained above, the number of agents always reporting 0 or 1 can be restricted to the following cases: (i)  $n_0 = 0, n_1 = 0$ , (ii)  $n_0 = 1, n_1 = 0$ , (iii)  $n_0 = 0, n_1 = 1$  and (iv)  $n_0 = 1, n_1 = 1$ . For all cases, consider the strategy profiles where  $\bar{n} \geq 1$  agents are honest, and the remaining agents lie according to  $s^{lie}$ . For each such profile I will show that no values of  $x$  and  $y$  can simultaneously satisfy the equilibrium constraints of both an honest and a lying agent. Moreover, when  $n_0 = n_1 = 1$  the strategy profiles where all other agents are honest or all other agents lie, cannot be Nash equilibria. The technical details of the proof are given in Appendix 3.G. ■

The remaining lying strategy profiles to be considered for computing the values of  $x$  and  $y$  are the following:

- $s_1 = (N \times s^{lie})$  when all agents lie; the constraints to prevent this equilibrium are also considered in LP 3.6.2;
- $s_2 = (s^{neg}, (N - 1) \times s^{lie})$  when one agent always reports 0, and all other agents lie;
- $s_3 = (s^{pos}, (N - 1) \times s^{lie})$  when one agent always reports 1, and all other agents lie;

A solution for  $x$  and  $y$  can therefore be found in constant time.

### 3.6.5 Partial Coalitions on Symmetric Strategies, Non-Transferable Utilities

For most practical applications it is reasonable to assume that only a fraction of the agents will be able to collude. The non-colluders are assumed to report honestly, and their reports can be used by the mechanism to deter any “partial” lying coalition. Note, however, that honest reports cannot be identified and selectively used by the reputation mechanism.

I start as in Section 3.6.3, by assuming only symmetric collusion strategies and no side-payments available among colluders. The number of colluders is  $N_{col} < N$ , and the remaining  $\bar{N} = N - N_{col}$

report honestly. There are 3 symmetric pure lying strategies, and appropriate constraints can ensure that none of them becomes a Nash equilibrium, or Pareto-dominates the honest equilibrium.

Concretely, let  $\bar{P}r[\cdot|\cdot]$  be the probability distribution of the reports submitted by non-colluders, such that  $\bar{P}r[n|o_i]$  is the probability that  $n$  out of  $\bar{N}$  agents report positively given the observation  $o_i \in Q_2$ . Likewise, let  $\hat{P}r[\cdot|\cdot]$  be the probability distribution of the reports submitted by the other colluders: i.e.,  $\hat{P}r[n|o_i]$  is the probability that  $n$  out of  $N_{col} - 1$  colluders report positively.

The payment scheme that makes honest reporting the unique Nash equilibrium for the colluders and minimizes the expected payment to an honest reporter solves the following optimization problem:

$$\begin{aligned}
\min \quad & E[V(\bar{s}_i, \bar{s}_{-i})]; \\
\text{s.t.} \quad & \sum_{n=0}^N Pr[n|0](\tau(0, n) - \tau(1, n)) \geq \Lambda; \\
& \sum_{n=0}^N Pr[n|1](\tau(1, n) - \tau(0, n)) \geq \Lambda; \\
& \bigvee_{o_i \in Q_2} \left( \sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i](\tau(1, n + N_{col} - 1) - \tau(0, n + N_{col} - 1)) < 0 \right) \\
& \bigvee_{o_i \in Q_2} \left( \sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i](\tau(0, n) - \tau(1, n)) < 0 \right) \\
& \bigvee_{o_i \in Q_2} \left( \sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i] \sum_{x=0}^{N_{col}-1} \hat{P}r[N_{col} - 1 - x|o_i](\tau(1 - o_i, n + x) - \tau(o_i, n + x)) < 0 \right) \\
& \tau(0, n), \tau(1, n) \geq 0; \forall n = \{0, 1, \dots, N - 1\};
\end{aligned}$$

where besides the first two incentive-compatibility constraints, the third, fourth and fifth constraints encourage deviations from the symmetric collusion on  $s^{pos}$ ,  $s^{neg}$  and  $s^{lie}$  respectively. The resulting optimization problem is a disjunctive linear program.

Finally, honest reporting can be made the Pareto-optimal equilibrium by modifying the optimization problem such that the disjunctive constraints preventing an equilibrium on the lying symmetric strategies, also specify inequalities for limiting the payoff received by a colluder below the expected payment of an honest reporter:

- colluders on  $s^{pos}$  gain less than in the honest equilibrium:

$$\sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i] \tau(1, n + N_{col} - 1) < E[V(\bar{s}_i, \bar{s}_{-i})];$$

- colluders on  $s^{neg}$  gain less than in the honest equilibrium:

$$\sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i] \tau(0, n) < E[V(\bar{s}_i, \bar{s}_{-i})];$$

- colluders on  $s^{lie}$  expect to gain less than in the honest equilibrium:

$$\sum_{o_i=0,1} Pr[o_i] \sum_{n=0}^{\bar{N}} \bar{P}r[n|o_i] \sum_{x=0}^{N_{col}-1} \hat{P}r[N_{col} - 1 - x|o_i] \tau(1 - o_i, n + x) < E[V(\bar{s}_i, \bar{s}_{-i})];$$

### 3.6.6 Partial Coalitions on Asymmetric Strategies, Non-Transferable Utilities

A more practical scenario is to remove the restriction on symmetric collusion strategies, and consider all strategy profiles  $s = (s_i)$ ,  $i = 1, \dots, N_{col}$  that colluders can use to misreport feedback. The remaining  $\bar{N} = N - N_{col}$  agents are assumed to report honestly, and side-payments are not available among colluders.

The main difference from the previous scenarios is that here, we can use a stronger equilibrium concept than Nash equilibrium. When the fraction of colluders is small enough, honest reporting can be made a dominant strategy for the colluders, such that regardless of what the other  $N_{col} - 1$  agents report, truth-telling is rewarded more than lying by at least some margin  $\Lambda$ .

If  $\bar{Pr}[\cdot|\cdot]$  describes the probability distribution of the  $\bar{N}$  honest reports, and  $c$  is the number of positive reports submitted by the other  $N_{col} - 1$  colluders, the payments  $\tau(\cdot, \cdot)$  that make honest reporting the dominant strategy, and minimize the payment to an honest reporter, are defined by the following optimization problem:

#### LP 3.6.3

$$\begin{aligned}
 \min \quad & E[V(\bar{s}_i, \bar{s}_{-i})] = Pr[1] \sum_{n=0}^{N-1} Pr[n|1] \tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0] \tau(0, n); \\
 \text{s.t.} \quad & \sum_{n=0}^{\bar{N}} \bar{Pr}[n|0] (\tau(0, n+c) - \tau(1, n+c)) \geq \Lambda; \\
 & \sum_{n=0}^{\bar{N}} \bar{Pr}[n|1] (\tau(1, n+c) - \tau(0, n+c)) \geq \Lambda; \\
 & \forall c \in \{0, \dots, N_{col} - 1\}, \\
 & \tau(0, n), \tau(1, n) \geq 0; \forall n = \{0, 1, \dots, N - 1\};
 \end{aligned}$$

The remaining question is how large may the colluding fraction be, such that collusion-resistant, incentive-compatible mechanisms exist.

**Proposition 3.6.9** *When more than half of the agents collude, (i.e.,  $N_{col} > N/2$ ), no incentive-compatible payment mechanism can make truth-telling the dominant strategy for the colluders.*

PROOF. The intuition behind the proof is the following: When  $N_{col} > N/2$ , the  $N_{col} - 1$  colluders submit at least as many reports as the remaining  $N - N_{col}$  honest reporters. Therefore, any sequence of honest reports, can be ‘corrected’ by a carefully chosen sequence of colluding reports, such that lying is profitable.

Formally, let us extract from the system of inequalities defined in LP 3.6.3, the subset corresponding

to  $c = \{0, \dots, \bar{N}\}$ . This subset exists since  $\bar{N} < N_{col} - 1$ . Let us form the following optimization problem:

$$\begin{aligned}
\min \quad & Pr[1] \sum_{n=0}^{N-1} Pr[n|1] \tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0] \tau(0, n); \\
\text{s.t.} \quad & \sum_{n=0}^{\bar{N}} \bar{Pr}[n|0] (\tau(0, n+c) - \tau(1, n+c)) \geq \Lambda; \quad \forall c = 0 \dots \bar{N} \\
& \sum_{n=0}^{\bar{N}} \bar{Pr}[n|1] (\tau(1, n+c) - \tau(0, n+c)) \geq \Lambda; \quad \forall c = 0 \dots \bar{N} \\
& \tau(0, n), \tau(1, n) \geq 0; \quad \forall n = \{0, 1, \dots, N-1\};
\end{aligned}$$

Let  $y_c^0$  and  $y_c^1$  be the dual variables corresponding to the constraints where the colluding agents report  $c$  positive signals, and the agent observes 0, respectively 1. The dual problem becomes:

$$\begin{aligned}
\max \quad & \Lambda \sum_{c=0}^{N_{col}-1} (y_c^0 + y_c^1); \\
\text{s.t.} \quad & \sum_{c=0}^n \bar{Pr}[n-c|0] y_c^0 - \bar{Pr}[n-c|1] y_c^1 \leq Pr[0] Pr[n|0]; \quad \forall n = 0 \dots \bar{N} \\
& \sum_{c=0}^n \bar{Pr}[n-c|1] y_c^1 - \bar{Pr}[n-c|0] y_c^0 \leq Pr[1] Pr[n|1]; \quad \forall n = 0 \dots \bar{N} \\
& \sum_{c=n+1}^{\bar{N}} \bar{Pr}[\bar{N}+1-c|0] y_c^0 - \bar{Pr}[\bar{N}+1-c|1] y_c^1 \leq Pr[0] Pr[\bar{N}+n+1|0]; \quad \forall n = 0 \dots \bar{N} \\
& \sum_{c=n+1}^{\bar{N}} \bar{Pr}[\bar{N}+1-c|1] y_c^1 - \bar{Pr}[\bar{N}+1-c|0] y_c^0 \leq Pr[1] Pr[\bar{N}+n+1|1]; \quad \forall n = 0 \dots \bar{N}
\end{aligned}$$

One can easily verify that the dual problem accepts as solutions:

$$\begin{aligned}
y_c^1 &= \bar{Pr}[c|0] \cdot \text{const}, \\
y_c^0 &= \bar{Pr}[c|1] \cdot \text{const};
\end{aligned} \tag{3.26}$$

for any positive constants. The dual problem is therefore unbounded, which makes the primal infeasible. ■

The bound from Proposition 3.6.9 is also tight. Consider the plumber example presented in Section 3.3.1, and assume the reputation mechanism has  $N = 4$  reports. The following payments are resistant to the collusion of  $N_{col} = 2$  agents:


$\tau(\cdot, \cdot)$	0	1	2	3
0	1.575	3.575	0	0
1	0	0	2.203	0.943

For example, if Alice observes 1, reporting 1 is better than reporting 0 for any report of the other colluder:

$$\begin{aligned}
\bar{Pr}[0|1] \tau(1, 0) + \bar{Pr}[1|1] \tau(1, 1) + \bar{Pr}[2|1] \tau(1, 2) &= 1.715; \\
\bar{Pr}[0|1] \tau(0, 0) + \bar{Pr}[1|1] \tau(0, 1) + \bar{Pr}[2|1] \tau(0, 2) &= 0.715; \\
\bar{Pr}[0|1] \tau(1, 1) + \bar{Pr}[1|1] \tau(1, 2) + \bar{Pr}[2|1] \tau(1, 3) &= 1.138; \\
\bar{Pr}[0|1] \tau(0, 1) + \bar{Pr}[1|1] \tau(0, 2) + \bar{Pr}[2|1] \tau(0, 3) &= 0.138;
\end{aligned}$$

where  $\bar{Pr}[0|1] = 0.0385$ ,  $\bar{Pr}[1|1] = 0.1830$  and  $\bar{Pr}[2|1] = 0.7785$  are the probabilities that 0, 1, or 2 out of the  $\bar{N} = 2$  non-colluders report positively, given that Alice observed high quality.

In the general case, the design problem has  $2N_{col}$  different constraints, and therefore, we should expect the budget required by the reputation mechanism to grow with  $N_{col}$ . I resort to numerical simulations and study the average cost of an incentive-compatible, collusion-resistant reputation mechanism as the fraction of colluders increases. I randomly generated 5000 problems as described in Appendix 3.E, and considered reward mechanisms for 5, 10, 15, 20 and 25 agents. For each problem and every number of agents the number of colluders was varied from 1 to  $N/2$ . Figure 3.10 plots the average normalized<sup>12</sup> cost of the collusion-resistant mechanism as a function of the colluding fraction,  $N_{col}/N$ . One can see that collusion resistance comes almost for free as long as less than one third of the population colludes. Above this bound the cost increases exponentially, which makes most such mechanisms impractical.



**Figure 3.10:** The average cost of the mechanism as we increase the colluding fraction.

I used the same numerical simulations to investigate the tightness of the bound set by Proposition 3.6.9. Table 3.6 presents the distribution of the maximum collusion threshold for the randomly generated problems. For more than 95% of the problems we were actually able to compute payment mechanisms that resist the maximum coalition size  $N_{col}^{max} = \lfloor N/2 \rfloor$  described by Proposition 3.6.9.

**Table 3.6:** Distribution of the maximum coalition bound.  $N_{col}^{max} = \lfloor N/2 \rfloor$ .

	Distribution of max coalition size (in %) over $[N_{col}^{max}, N_{col}^{max} - 1, \dots, 1]$
$N = 5, N_{col}^{max} = 2$	[99.98, 0.02]
$N = 10, N_{col}^{max} = 5$	[99.5, 0.36, 0.1, 0.04, 0]
$N = 15, N_{col}^{max} = 7$	[98.88, 0.54, 0.38, 0.08, 0.1, 0.02, 0]
$N = 20, N_{col}^{max} = 10$	[97.1, 0.86, 0.78, 0.56, 0.34, 0.2, 0.1, 0.04, 0.02, 0]
$N = 25, N_{col}^{max} = 12$	[96.3, 0.98, 0.76, 0.58, 0.48, 0.4, 0.24, 0.1, 0.1, 0.04, 0.02, 0]

The one-half collusion threshold set by Proposition 3.6.9 can be overcome by making honest reporting a Nash equilibrium, rather than a dominant strategy, for colluders. Consider the payment mechanism described by Proposition 3.6.7 where the reputation mechanism considers  $N = 4$  reports,

<sup>12</sup>The cost is normalized to the cost of the corresponding incentive-compatible mechanism that is not collusion-resistant.



rewards consensus with a small positive payment  $\varepsilon$ , and otherwise pays a report if and only if all but one of the other agents submit the same report. As proven above, this payment scheme accepts three Nash equilibria, where: (i) all agents report honestly, (ii) all agents always reports 0, and (iii) all agents always report 1. The fact that one agent cannot become part of a coalition and hence reports the truth, restricts the set of equilibria to only one, namely the honest reporting equilibrium. Even when the remaining three agents collude, their only NE under this payment mechanism is to report the truth.

General payment mechanisms based on  $N > 4$  agents can be designed similarly to those from Section 3.6.4: for all strategy profiles the colluders can use, constraints must be added to the design problem such that (i) no lying strategy profile is a NE, or (ii) no NE lying strategy profile Pareto-dominates the honest reporting NE.

Concretely, let  $S_{N_{col}}$  be the set of strategy profiles the colluders can use:

$$S_{N_{col}} = \{(n_0 \times s^{neg}, n_1 \times s^{pos}, \bar{n} \times \bar{s}, n_l \times s^{lie}) | n_0 + n_1 + \bar{n} + n_l = N_{col}\}$$

where  $s = (n_0 \times s^{neg}, n_1 \times s^{pos}, \bar{n} \times \bar{s}, n_l \times s^{lie}) \in S_{N_{col}}$  is the strategy profile where  $n_0$  out of  $N_{col}$  colluders always report 0,  $n_1$  colluders always report 1,  $\bar{n}$  colluders report honestly, and  $n_l$  colluders always lie. When colluders report according to the strategy profile  $s$ , let  $(s, \bar{s}) = (n_0 \times s^{neg}, n_1 \times s^{pos}, (\bar{n} + \bar{N}) \times \bar{s}, n_l \times s^{lie})$  be the strategy profile used by the  $N$  agents, where the  $\bar{N} = N - N_{col}$  non-colluders report honestly.

Honest reporting is the unique Nash equilibrium for colluders if and only if:

$$V(s_i, (s_{-i}, \bar{s}) | o_i) < V(s_i^*, (s_{-i}, \bar{s}) | o_i); \quad (3.27)$$

for some colluder  $i$  and observation  $o_i$ . Similarly,  $s$  does not Pareto-dominate the honest reporting equilibrium if:

$$E[V(s_i, (s_{-i}, \bar{s}))] < E[V(\bar{s}_i, \bar{s}_{-i})] \quad (3.28)$$

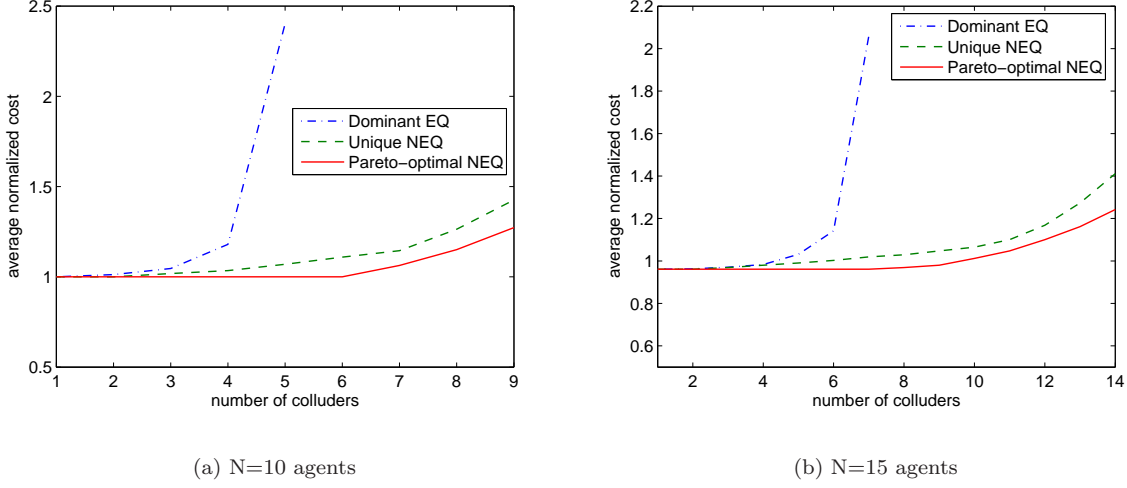
for some colluder  $i$ .

When compared to the similar constraints (3.24) of Section 3.6.4, there are two important differences. First, the inequalities apply only to the  $N_{col}$  colluders, not to the entire set of agents, so that a strategy profile that is not a NE for the  $N$  agents, might still be a Nash equilibrium for colluders. Take, for example, the case where all colluders lie: i.e.,  $s = (N_{col} \times s^{lie})$ . The strategy profile of all agents is therefore  $(s, \bar{s}) = (N_{col} \times s^{lie}, \bar{N} \times \bar{s})$ . It may well be possible that:

- a lying colluder finds it optimal to report according to  $s^{lie}$  given that  $\bar{N}$  agents report honestly and  $N_{col} - 1$  agents report lies;
- an honest reporter would rather file a negative report after a positive experience, given that  $\bar{N} - 1$  agents report honestly and  $N_{col}$  agents lie.

So  $(s, \bar{s})$  is not a NE when considering all agents, but it is an equilibrium for the subset of colluders. Similarly, it is possible that colluders gain better under  $(s, \bar{s})$  and honest reporters gain less, such that  $(s, \bar{s})$  Pareto-dominates the honest equilibrium for the colluders, but not for all agents. The constraints (3.27) and (3.28) are therefore stricter than their counterparts (3.24), as non-colluders are assumed to unconditionally report the truth without taking into account the actions of a lying coalition.

Second, we can separately consider the constraint (3.27), when honest reporting is to be enforced as the unique NEQ for colluders, or both constraints (3.27) and (3.28), when honest reporting is to be enforced as the Pareto-optimal NEQ for colluders. The presence of honest reports makes it possible to design payment mechanisms where honesty is the unique NE, an alternative that was not available under



**Figure 3.11:** Average normalized cost of collusion-resistant payment mechanism. Different equilibrium concepts.

the assumptions of Section 3.6.4. In both cases, the constraints preventing lying equilibria (or preventing lying equilibria from dominating the honest equilibrium) can be represented by a disjunction of linear inequalities, and consequently, by a conjunction of mixed integer linear constraints. The resulting design problem is a MILP, and, as discussed in Section 3.6.4, has a worst-time complexity that grows exponentially with the number  $N$  of agents.

Figure 3.11 compares the average normalized cost of a collusion-resistant payment mechanism when honest reporting is: (i) the dominant strategy, (ii) the unique NE, or (iii) the Pareto-optimal NE. The plots were generated by solving 100 randomly generated problems, for  $N = 10$  and  $N = 15$  agents. Computing the payment mechanism which satisfies the constraints (3.27) and (3.28) requires significantly more time, hence the lower number of generated problems. Moreover, the capabilities of our solver were exceeded for payments using more than 15 agents. Nevertheless, the loss in computational efficiency is clearly rewarded by both lower cost of the mechanism, and coverage of greater coalitions.

As in Section 3.6.4, the payment mechanism from Figure 3.9 can be used to reduce the complexity of the design problem when honest reporting is the unique or the Pareto-optimal NE. From Proposition 3.6.8, it is known that:

- in any NE, at most one colluder reports according to  $s^{neg}$  or  $s^{pos}$ ;
- an honest reporter and a liar cannot both regard their strategies as optimal given that at most one of the other agents reports according to  $s^{neg}$  or  $s^{pos}$

Therefore, the remaining colluding strategy profiles that must be considered when designing the payments  $x$  and  $y$  are the following:

- $(N_{col} \times s^{lie})$  when all colluders lie;
- $(s^{neg}, (N_{col} - 1) \times s^{lie})$  when one colluder always reports 0 and the others always lie;

- $(s^{pos}, (N_{col} - 1) \times s^{lie})$  when one colluder always reports 1 and the others always lie;

### 3.6.7 Partial Coalitions on Asymmetric Strategies, Transferable Utilities

As a last scenario I assume that one strategic agent controls a number of fake online identities, or *sybils*. From the agent's perspective, the individual revenues obtained by each sybil is irrelevant; the objective of the agent is to maximize the cumulated revenue obtained by all sybils.

The fact that utilities are transferable makes the problem of the mechanism designer significantly harder. In all previous scenarios, the constraints that made an incentive-compatible mechanism collusion-resistant ensured that lying coalitions are either unstable or unprofitable. However, transferable utilities allow some colluders to subsidize others, such that non-equilibrium colluding strategies can still exist. Therefore, the necessary (and sufficient) condition for collusion resistance in this context requires that the cumulated revenue of the coalition is maximized when reporting the truth.

Another difference from the settings in Sections 3.6.4 and 3.6.6 is that colluders coordinate their reporting strategy *after* observing the quality signals. This assumption is supported by the interpretation that one strategic entity controls several fake online identities.

Concretely, the payment mechanism must have the following property: whenever  $N_{col}$  colluding agents observe  $c$  high quality signals, their cumulated revenue is maximized when reporting  $c$  positive reports. An underlying assumption is that non-colluders (the other  $\bar{N} = N - N_{col}$  agents) are reporting honestly. The revenue of the coalition that reports  $r$  (out of  $N_{col}$ ) can be computed as follows. The  $r$  colluders that report positively are rewarded  $\tau(1, r - 1 + n)$ , while the  $N_{col} - r$  colluders that report negatively are rewarded  $\tau(0, r + n)$ ;  $n$  is the number of positive reports submitted by the (honest) non-colluders. The expected revenue of the coalition is therefore:

$$V(r|c) = \sum_{n=0}^{\bar{N}} \bar{Pr}[n|c] \left( r \cdot \tau(1, r - 1 + n) + (N_{col} - r) \cdot \tau(0, r + n) \right); \quad (3.29)$$

where  $\bar{Pr}[n|c]$  is the probability that  $n$  out of  $\bar{N}$  honest agents report positively, given that  $c$  out of  $N_{col}$  colluders observed high quality signals.

Honest reporting is the best strategy for the coalition, when for all  $c \in \{0, \dots, N_{col}\}$ ,  $\arg \max_r V(r|c) = c$ :

$$\sum_{n=0}^{\bar{N}} \bar{Pr}[n|c] \left( c \cdot \tau(1, c - 1 + n) + (N_{col} - c) \cdot \tau(0, c + n) - r \cdot \tau(1, r - 1 + n) - (N_{col} - r) \cdot \tau(0, r + n) \right) \geq \Lambda; \forall r \neq c \in \{0, \dots, N_{col}\} \quad (3.30)$$

The cheapest incentive-compatible, collusion-resistant payment mechanism minimizes the objective function (3.19) under the linear constraints (3.30):

#### LP 3.6.4

$$\begin{aligned} \min \quad & E[V(\bar{s}_i, \bar{s}_{-i})] = Pr[1] \sum_{n=0}^{N-1} Pr[n|1] \tau(1, n) + Pr[0] \sum_{n=0}^{N-1} Pr[n|0] \tau(0, n); \\ \text{s.t.} \quad & (3.30) \text{ is true, } \forall c, r \in \{0, \dots, N_{col}\}, c \neq r \\ & \tau(0, n), \tau(1, n) \geq 0; \forall n = \{0, 1, \dots, N - 1\}; \end{aligned}$$

For the plumber example from Section 3.3.1, assuming that Alice controls  $N_{col} = 3$  different online identities that may submit feedback about Bob, the following payments based on  $N = 6$  reports deter Alice from lying:

$\tau(\cdot, \cdot)$	0	1	2	3	4	5
0	20.85	0	0	0	4.40	9.98
1	45.54	28.78	0	0	0	4.31

Even if Alice controlled  $N_{col} = 5$  out of the  $N = 6$  reports, there still are payments that make honest reporting rational. These payments, however, are significantly higher:

$\tau(\cdot, \cdot)$	0	1	2	3	4	5
0	3455	0	1378	615	0	1125
1	1530	5569	4674	3736	0	2585

It turns out that for the general case, one honest report is enough to allow the design of incentive-compatible payments that also deter sybil attacks of size  $N - 1$ . An example of such payments are presented in the proposition below:

**Proposition 3.6.10** *Given the set of types  $\Theta$ , the conditional probabilities  $Pr[1|\theta]$ , the prior belief over types  $Pr[\theta]$  and a number  $N$  of reports, the following payments encourage honest reporting from a strategic agent who controls  $N - 1$  different reports:*

$$\begin{aligned} \tau(0, 0) &= \frac{ScR(0, 0)}{N_{col}}; \tau(0, 1) = \frac{ScR(1, 0)}{N_{col}}; \tau(1, N) = \frac{ScR(1, N - 1)}{N_{col}}; \\ \tau(0, x + 1) &= \frac{(x + 1)ScR(1, x) - xScR(0, x + 1)}{N_{col}}; \quad x = 1 \dots N - 1 \\ \tau(1, x) &= \frac{(N - 1 - x)ScR(0, x + 1) - (N - 2 - x)ScR(1, x)}{N_{col}}; \quad x = 1 \dots N - 1 \end{aligned}$$

where  $ScR(i, j)$ ,  $i \in \{0, 1\}$ ,  $j = \{0, \dots, N - 1\}$  is a proper scoring rule: e.g., the logarithmic proper scoring rule from Eq. (3.3), the quadratical proper scoring rule from Eq. (3.5) or the spherical proper scoring rule from Eq. (3.4).

PROOF. The expected payment of an agent who controls  $N - 1$  different identities, observes  $c$  out of  $N - 1$  positive signals and reports  $r$  positive reports to the reputation mechanism can be computed as in Eq. (3.29):


$$\begin{aligned} V(r|c) &= Pr[0|c] \left( r \cdot \tau(1, r - 1) + (N - 1 - r) \cdot \tau(0, r) \right) + \\ &\quad Pr[1|c] \left( r \cdot \tau(1, r) + (N - 1 - r) \cdot \tau(0, r + 1) \right) \\ &= \dots = Pr[0|c] ScR(0, r) + Pr[1|c] ScR(1, r); \end{aligned}$$

which by the definition of a proper scoring rule is strictly maximized when  $r = c$ : i.e.  $V(c|c) - V(r|c) > 0$  for all  $r \neq c$ . By scaling the scoring rule appropriately (i.e., multiplication and addition with a constant), honest reporting can be made better by lying with at least the margin  $\Lambda$ . ■

Proposition 3.6.10 proves the existence of incentive-compatible collusion resistant rewards when all but one report are controlled by the same strategic agent. However, as we have seen for the example in

Section 3.3.1, such payments can be very expensive, and hence impractical. Numerical simulations can be used to evaluate the marginal cost of increasing collusion resistance as the number of colluders (i.e., reports controlled by the same agent) increases. As in Section 3.6.6, I generated 5000 random problems and computed the optimal payments for  $N = 5, 10, 15, 20$  and 25 reports. For each case, the coalition size (i.e.,  $N_{col}$ ) is increased from 1 to  $N - 1$ .

Figure 3.12 plots the average normalized<sup>13</sup> cost of the collusion-resistant mechanism as a function of the coalition fraction. The cost grows linearly for coalitions that span up to one half of the population; for larger coalitions, the cost grows exponentially. Nevertheless, by comparing Figures 3.12 and 3.10, one can see that for the same coalition size, the collusion-resistant payments are cheaper if assuming a setting with non-transferable utilities.



**Figure 3.12:** The average cost of the mechanism as we increase the colluding fraction (setting with transferable utilities).

## 3.7 Related Work

One interesting alternative to payment schemes that encourage honest feedback is to develop mechanisms that make it in the best interest of the providers to truthfully reveal their hidden quality attributes. The truthful declaration of quality eliminates the need for reputation mechanisms and significantly reduces the cost of trust management.

Braynov and Sandholm (2002), for example, consider exchanges of goods for money and prove that a market in which agents are trusted to the degree they deserve to be trusted is equally efficient as a market with complete trustworthiness. By scaling the amount of the traded product, the authors prove that it is possible to make it rational for sellers to truthfully declare their trustworthiness. However, the assumptions made about the trading environment (i.e. the form of the cost function and the selling price which is supposed to be smaller than the marginal cost) are not common in most electronic markets.

Another interesting work that addresses the trustworthiness of reputation information is the Goodwill Hunting mechanism of Dellarocas (2002). The mechanism works for eBay-like markets and provides a way to make sellers indifferent between lying or truthfully declaring the quality of the good offered for sale. The particularity of this work is that the goods are advertised to the buyers through the

<sup>13</sup>The cost is normalized to the cost of the corresponding incentive-compatible mechanism that is not collusion resistant.

reputation mechanism, which can modify the asking price initially set by the seller. The reputation mechanism thus compensates the momentary gains or losses made by the seller from misstating the quality of the good, and creates an equilibrium where all sellers find it rational to truthfully announce the quality. A major advantage of the mechanism is that it works even when the sellers offer various goods with different values.

Mechanisms for encouraging honest reporting are also present in a number of commercial applications. The most famous is perhaps the *ESP Game* (von Ahn and Dabbish, 2004), designed to encourage human users to label web images. The game<sup>14</sup> pairs two users at random, and shows them the same image. Each player must individually write tags about the image, without being able to see the tags written by the partner. As soon as the two player write the same tag, they gain points and can pass to the next picture. The goal is to get as many points as possible in a fixed amount of time. Intuitively, this game has a very simple strategy: players must write as many correct tags as possible, since the image they see is the only synchronization device that allows them to reach agreement on a tag. The game is very successful, and the authors claim that in this way, all images on the web can be tagged in several months.

The incentive mechanism behind the ESP game has, however, several problems. First, it is vulnerable to cheating strategies where a group of players agree to reach agreement on a very simple tag like “a” or “the”. This strategy could be posted to a popular blog and exposed rapidly to the ESP players. As discussed in Section 3.6.1, these simple collusion strategies will give colluders a significant competitive advantage, at the detriment of the game designers who collect only garbage tags. The problem can be partly addressed by “taboo” lists containing all confirmed tags which were already submitted about the picture.

The second problem is that rewards are equal for all possible tags. For a Picasso, the players that match on the tag “*painting*” are equally rewarded as the players who correctly identify that the painting is a Picasso. This gives incentives to the players to concentrate on the simplest possible tags like “*person*”, “*man*”, “*woman*”, etc, without spending effort to provide more informative tags. This problem has been partly corrected by the Google Image Labeler<sup>15</sup>, a franchise of the ESP Game, which rewards players inversely proportional to the frequency of the tag they agree on. However, the exact algorithm for computing the rewards is not public. Yahoo! is also known to use a version of the ESP Game to tag their collection of images.

The ESP game can directly profit from the results of this chapter. The rewards used right now are very simple payment mechanisms that specify positive amounts on matching tags. However, the payments awarded to the players can be scaled according to the design constraints mentioned in Section 3.3.

Another example of a commercial application using payment mechanisms to encourage honest reporting is Amazon’s Mechanical Turk<sup>16</sup>. The role of the system is to provide a marketplace in which human users can solve tasks that are very difficult for machines, but easy for people (i.e., short translations, tagging, face recognition, natural language search, etc). Task owners can pay the workers for answering their tasks, and can also specify payment rules: e.g., a worker gets paid (or receives a bonus) only if the answer is confirmed by a different worker solving the same task. Here again, the amounts and the rules for paying the workers can be the object of incentive-compatible design.

A number of feedback forums reward raters independently based on the impact of their reviews on the other users. ePinion.com, for example, has professional reviewers who get paid depending on the votes expressed by normal users, and on the purchases made after reading the reviews. Another

---

<sup>14</sup><http://www.espgame.org>

<sup>15</sup><http://images.google.com/imagelabeler/>

<sup>16</sup><http://www.mturk.com/mturk/welcome>

example is the startup Friend2Friend.com<sup>17</sup> who allows users to gain commissions by recommending products to their friends.

Central to the results of this chapter is the principle of automated mechanism design (AMD). The mechanism is created automatically (using optimization algorithms) for the specific problem instance, given the specific information available to the mechanism designer. The idea has important advantages since (a) it can be used to address classes of problems for which there are no known manually designed mechanisms, (b) it can circumvent impossibility results by restricting the mechanism to one particular setting, (c) it can generate better mechanisms by capitalizing on the specific information available in the present setting, and (d) it shifts the effort of mechanism design to a machine.

Since first introduced by Conitzer and Sandholm (2002), AMD was used to generate several impressive results. Conitzer and Sandholm (2003a) (a) reinvented the Myerson auction which maximizes the seller's expected revenue in a single-object auction, (b) created expected revenue maximizing combinatorial auctions, and (c) created optimal mechanisms for a public good problem. Guo and Conitzer (2007) use AMD to optimally redistribute the payments generated by the VCG mechanism, Conitzer and Sandholm (2007) incrementally design incentive compatible mechanisms, while Hajiaghayi et al. (2007) focus on AMD for online settings. Conitzer and Sandholm (2003b) show that the AMD can potentially be exponentially faster for settings with structured preferences that allow a concise representation of the input. Conitzer and Sandholm (2004) describe an efficient algorithm for AMD when the mechanism is deterministic, does not allow payments and there is only one type-reporting agent. AMD can also be used to design multi-stage mechanisms that reduce the burden of information elicitation by querying the agents only for relevant information (Sandholm et al., 2007). The results of Sections 3.3 and 3.6 add to this already long list of results obtained through AMD.

The results of Section 3.6 are mostly related to the literature on implementation theory and incentive contracts for principle-(multi)agent settings. The main goal of implementation theory is to characterize the space of social choice rules that are implementable by some mechanism given a game-theoretic equilibrium concept. For complete information settings, well established results characterize the necessary and sufficient conditions for a social choice rule (SCR) to be implementable in dominant strategy or in Nash equilibrium. For example, SCRs can be implemented in dominant strategies only if they are strategy-proof (Gibbard, 1973), while the SCRs that can be Nash-implemented must satisfy the property of *monotonicity* and *no veto power* (Maskin, 1999). Unfortunately, SCRs of practical interest do not satisfy the monotonicity requirement. Fortunately, non-monotonic SCRs can be implemented in undominated Nash equilibria (Palfrey and Srivastava, 1991), or in subgame perfect equilibria by multi-stage mechanisms. Another relaxation that extends the set of implementable SCRs is to consider *virtual* implementation, where the socially optimal outcome is required to occur only with probability close to one (Matsushima, 1988; Abreu and Sen, 1991).

In environments with incomplete information agents have private information that is not shared by other agents. The truthful revelation of the private information can only be ensured by social choice rules that are Bayesian incentive-compatible. Moreover, a Bayesian monotonicity condition is necessary for Bayesian implementation (Jackson, 1991). Moore and Repullo (2005) characterize SCRs that can be *virtually* Bayesian implemented in pure strategies, and derive the necessary and sufficient conditions of incentive compatibility and *virtual monotonicity*.

However, applying the implementation theory to the feedback reporting setting (an environment with incomplete information) provides nothing more than the constraints on the payment function such that honest reporting is the unique Bayesian Nash equilibrium. In implementation theory terms, the set of possible world states consists of all combinations of  $N$  privately perceived quality signal (one signal for each agent). The outcome space contains all possible sets of  $N$  feedback reports and all possible combinations of  $N$  positive payments made to the  $N$  agents. The desirable SCR contains all social choice

---

<sup>17</sup><http://www.friend2friend.com/>

functions that map the possible states of the world (i.e., the set of privately perceived signals) to the outcomes where the reported feedback correspond to the privately perceived signals). Implementation theory tells that the SCR must be incentive compatible (i.e., the social choice functions prescribe outcomes where the payments to the agents make them truthfully reveal their private information) and Bayesian monotone (i.e., the social choice functions prescribe outcomes where the payments received by the agents make honest reporting the unique equilibrium). The results of Section 3.6 translate these requirements into practical constraints that allow the computation of payment functions (and therefore social choice functions) that are Bayesian Nash implementable.

A number of papers discuss incentive contracts that a principal should offer to several agents whose effort levels are private. The reward received by each agent depends on the output observed by the principal, and on the declarations of other agents. Holmström (1982), Ma (1988), and Li and Balachandran (2000) show that efficient contracts exist that are also incentive-compatible and collusion-proof. While the feedback reporting problem is similar, it differs in one major aspect: the reputation mechanism designer (i.e., the principal) does not observe a direct signal which is correlated to the reporters' (i.e., agents') private information.

### 3.8 Summary of Results

Two problems commonly associated with online feedback mechanisms are (i) the low proportion of users that leave feedback, and (ii) the temptation to provide false or biased information. Both can be addressed by reward schemes which (i) cover the cost of obtaining and reporting feedback, and (ii) maximize the expected reward of a rational agent who reports truthfully. I address in this chapter the design of such "truthful" rewards for signaling reputation mechanisms. Here, the correlation between the private signal of an agent and her belief regarding the reports of other agents can be exploited to make honest reporting a Nash equilibrium.

Section 3.2.1 reviews the work of Miller et al. (2005) and explains how incentive-compatible payment mechanisms can be computed based on proper scoring rules. Section 3.3 uses the idea of automated mechanism design to construct payment schemes that minimize the expected cost to the reputation mechanism, while offsetting both the cost of reporting and the external gains an agent could obtain from lying. The payments can be computed effectively by solving a linear optimization problem, and reduce approximately three times the expected budget required by the reputation mechanism.

Section 3.4 investigates two methods that can further decrease incentive-compatible feedback payments. The first requires the use of several reference reports. I prove that higher numbers of reference reports lead to lower costs; however, experiments show that little benefit can be obtained by using more than 4 reference reports. The second method involves probabilistic filtering mechanisms that discard some feedback reports. The key idea is to design the payment and the filtering mechanism together, by considering both the information available to the agents, and the similarity between peer reports. The resulting mechanism has much lower cost (expected payments are up to an order of magnitude lower) without affecting the quality of reputation information.

Section 3.5 extends the design of incentive-compatible payments to settings with uncertain information. If buyers have private information about the true type of a product or seller, the default payments computed by the reputation mechanism are not always incentive-compatible. One solution is to ask the buyers to reveal their private information, and have the reputation mechanism compute customized payments for each reporter. Besides the cost of the added communication, this method is also vulnerable to false declaration of private information. As an alternative, I consider payment mechanisms that are incentive-compatible for a range of private beliefs.



Incentive-compatible payments can easily be adapted to address the collusion between buyers and sellers. The collusion can happen when the benefit obtained by the seller from a false report offsets the payment returned to the client in exchange for lying. The feedback payments presented in Sections 3.3 and 3.4 make sure that no seller can afford to buy false reports from rational buyers.

The collusion between reporters, on the other hand, is addressed in Section 3.6. First, I show that general incentive-compatible payments have several equilibria besides the truthful one. Some of the lying equilibria generate higher payoffs than the honest equilibrium, which might motivate selfish agents to coordinate their reports and game the mechanism. Fortunately, supplementary constraints can be added to the design problem such that honest reporting becomes the unique, or the pareto-optimal equilibrium. For different collusion scenarios, I describe algorithms for computing such incentive-compatible, collusion-resistant payments.

One remaining problem, however, is that the seller could create fake buyer identities, (or “bribe” real buyers that never purchased the product) in order to bias the reputation information. The problem can be addressed by security mechanisms that connect feedback reports to real transaction IDs. For example, a site like Expedia.com can make sure that no client leaves feedback about a hotel without actually paying for a room. Hotels could, of course, create fake bookings, but the repeated payment of Expedia commission fees makes the manipulation of information very expensive. On the other hand, social norms and legislation (i.e., providers that try to bribe clients risk being excluded from the market) could further avoid provider side manipulation.

One hidden assumption behind the results of this chapter is that all agents are risk-neutral. Honest reporting is rational because, in expectation, it brings higher revenues. However, risk-averse buyers prefer the “sure” benefit from lying to the probabilistic feedback payment, and thus misreport. Fortunately, all mechanisms described above can be adapted to any risk-model of the buyers. Nevertheless, (i) risk-averse buyers require significantly higher payments from the mechanism, and (ii) non-linear risk models may lead to non-linear design problems.

The motivation to minimize feedback payments might not be clear when the budget of the mechanism is covered by buyer subscription fees. Higher fees will be matched (in expectation) by higher feedback payments. However, this holds only for risk-neutral buyers. Real users (probably risk-averse) will regard the fixed fees as more expensive than the revenue expected from honest reporting; and higher subscription fees are increasingly more expensive. Moreover, no real-world system implemented today charges users for reputation information. Introducing reputation information fees could seriously deter participation.

An interesting question is what happens as more and more reports are recorded by the reputation mechanism. When the type of the product does not change, the beliefs of the buyers rapidly converge towards the true type (Figure 3.6). As more information becomes available to buyers, the private quality signal they observe triggers smaller and smaller changes of the prior belief. As a consequence, the probability distributions for the reference reports conditional on the private observation (i.e.,  $Pr[q_k|\cdot]$ ) become closer, and the payments needed to guarantee a minimal expected loss from lying increase. Fortunately, external benefits from lying also decrease: the effect of a false report on reputation information tends to 0. Depending on the particular context, lying incentives decrease faster, respectively slower than the distance between the conditional probability distributions for the reference reports, and thus, truth-telling becomes easier respectively harder to guarantee. Whatever the case, it makes sense to stop collecting feedback when the beliefs of the buyers are sufficiently precise.

In real settings, however, the true type of a product or service does actually change in time: e.g., initial bugs get eliminated, the technology improves, etc. Existing incentive compatible payment scheme rely on the fact that future buyers obtain exactly the same thing as the present ones. Fortunately, changes of type are unlikely to be frequent, so that by dividing the time into “rounds” (the exact length

of a round will be an application dependent parameter) we can safely assume that changes of type may only occur between rounds.

There are several directions for extending the results of this chapter in the future. First, to make the mechanisms more practical, one might need to introduce further constraints in the design problem. For example, a designer might wish to limit the highest payment to an agent, or require all payments to take integer values so that they can be mapped to application dependent “points”. Understanding the full repercussions on the resulting mechanism requires further study. Second, the robustness guarantees of Section 3.5 could be extended to other models of private information.

Third, the extension of Section 3.6 to general settings with  $n$ -ary feedback is an interesting endeavor. While conceptually straight-forward, the resulting design problem becomes exponentially more complex as we increase the number of possible feedback values. The problem becomes to find efficient algorithms that are able to quickly compute incentive-compatible collusion-resistant payment mechanisms for non-binary feedback mechanisms.

Finally, another direction for future research is to design payment mechanisms that are resistant against a mixed collusion scenario, where, for example, several strategic agents, each controlling several fake identities, try to manipulate the reporting mechanism.

## Appendix 3.A Summary of Notation

Symbol	Meaning
$\Theta, \theta$	$\Theta$ is the set of possible types; $\theta \in \Theta$ is one type;
$\Theta_2 = \{\theta_G, \theta_B\} = \{1, 0\}$	in binary settings, the good, respectively the bad type;
$Pr[\theta]$	the prior probability of the type $\theta$ ;
$Pr[q_j \theta]$	the probability of observing the signal $q_j$ given the type $\theta$ ;
$O_i$	the random signal denoting the observation of the buyer $i$ ;
$o_i \in Q$	the signal observed by the buyer $i$ ;
$r_i \in Q$	the signal reported by the buyer $i$ ;
$C_r$	the cost of reporting;
$\Lambda(q_j, q_h)$	the benefit from misreporting $q_h$ instead of the truth $q_j$ ;
$\hat{\Lambda}(q_j, q_h)$	the benefit from misreporting, discounted by using filtering;
$B$	the maximum budget of the reputation mechanism;
$s_i = (s_i(0), \dots, s_i(M-1))$	reporting strategy of buyer $i$ ; $s_i(j) \in Q$ is the report submitted when $O_i = q_j$ ;
$\bar{s} = (q_0, q_1, \dots, q_{M-1})$	honest reporting strategy;
$s^* \neq \bar{s}$	some lying strategy;
$ref(i)$	the reference report(s) of agent $i$
$\tau(r_i, r_{ref(i)})$	the payment received by agent $i$ when she report $r_i$ and the reference reporters report $r_{ref(i)}$ ;
$V(s_i, s_{ref(i)} o_i)$	the expected payment of an agent who observes the signal $O_i = o_i$ , reports according to the strategy $s_i$ , and the reference reporters report according to the strategy $s_{ref(i)}$ : $V(s_i, s_{ref(i)} o_i) = \sum_{q_k \in Q} Pr[O_{ref(i)} = q_k O_i = o_i] \tau(s_i(o_i), s_{ref(i)}(q_k))$
$ScR(\cdot)$	a scoring rule;
$ScR_{log}(\cdot)$	the logarithmic scoring rule;
$ScR_{sph}(\cdot)$	the spherical scoring rule;
$ScR_{quad}(\cdot)$	the quadratic scoring rule;
$N_{ref}$	the number of reference reporters;
$Q(N_{ref})$	the set containing all unordered sequences of $N_{ref}$ reference reports;
$\bar{q}_k \in Q(N_{ref})$	a set of signals observed or reported by $N_{ref}$ reference reporters;
$N_{ftr}$	the number of filtering reports;
$Q(N_{ftr})$	the set containing all unordered sequences of $N_{ftr}$ filtering reports;
$\hat{q}_k \in Q(N_{ftr})$	a set of signals observed or reported by $N_{ftr}$ filtering reporters;
$\pi(r_i, r_{ftr(i)})$	the probability of accepting the report submitted by agent $i$ ;
$\gamma_{infLoss}$	threshold probability of dropping useful information;
$const$	some constant;
$\varepsilon$	some small, positive number;

Symbol	Meaning
$Pr^*[\theta] = Pr[\theta] + \varepsilon_\theta$ $ \varepsilon_\theta  < \gamma_{maxPI}$ $Pr'[\cdot]$	private belief; bound on the amount of variation; some lying probability distribution;
$S$ $\bar{s} = (0, 1)$ $s^{lie}, s^{pos}, s^{neg}$ $s = (s_i, s_{-i})$ $(\bar{s}_i, \bar{s}_{-i})$	set of binary reporting strategies: $S = \{(s(0), s(1))   s(0), s(1) \in \Delta(Q_2)\}$ ; honest reporting strategy; lying strategies; a strategy profile, one strategy for each agent; honest reporting strategy profile;
$r_i, r_{-i}$ $N$ ( $N_{ref} = N - 1$ ) $n$ $\mu[n, s_{-i}]$ $N_{col}$	the report of agent $i$ , the reports of all the other agents except $i$ ; the number of agents available to the reputation mechanism; the number of positive reports in the reference reports; belief given $n$ positive observations and reporting strategy $s_{-i}$ ; number of colluding agents;

## Appendix 3.B Generating Random Settings

We consider settings where  $M$  possible product types are each characterized by one quality signal: i.e., the sets  $Q$  and  $\Theta$  have the same number of elements, and every type  $\theta_j \in \Theta$  is characterized by one quality signal  $q_j \in Q$ . The conditional probability distribution for the signals observed by the buyers is computed as:

$$Pr[q_k | \theta_j] = \begin{cases} 1 - \varepsilon & \text{if } k = j; \\ \varepsilon / (M - 1) & \text{if } k \neq j; \end{cases}$$

where  $\varepsilon$  is the probability that a buyer misinterprets the true quality of the product (all mistakes are equally likely). We take  $\varepsilon = 10\%$ .

The prior belief is randomly generated in the following way: for every  $\theta_j \in \Theta$ ,  $p(\theta_j)$  is a random number, uniformly distributed between 0 and 1. The probability distribution over types is then computed by normalizing these random numbers:

$$Pr[\theta_j] = \frac{p(\theta_j)}{\sum_{\theta \in \Theta} p(\theta)};$$

The external benefits from lying are randomly uniformly distributed between 0 and 1.

## Appendix 3.C Cardinality of $Q(N_{ref})$

The set of all possible values for  $N_{ref}$  reference reports can be described as:

$$Q(N_{ref}) = \left\{ (n_1, \dots, n_M) \in \mathbb{N}^M \mid \sum_{j=1}^M n_j = N_{ref} \right\};$$

The elements of  $Q(N_{ref})$  can be determined by first setting  $n_1$  and then recursively finding the vectors  $(n_2, \dots, n_M)$  such that  $\sum_{j=2}^M n_j = N_{ref} - n_1$ .  $|Q(N_{ref})|$  can therefore be computed by the

recursive function:

$$\begin{aligned} |Q(N_{ref})| &= g(N, M) \\ g(N, M) &= \sum_{j=0}^N g(N-j, M-1); \\ g(0, M) &= 1; \quad g(N, 1) = 1; \quad \forall N, M \in \mathbb{N}; \end{aligned}$$

By induction over  $M$ , we show that  $g(N, M) = \mathfrak{C}_{N+M-1}^{M-1}$  for all  $N \geq 0$ . True for  $M = 1$ ; assuming  $g(N, M-1) = \mathfrak{C}_{N+M-2}^{M-2}$  for all  $N \geq 0$ , we have:

$$\begin{aligned} g(N, M) &= \sum_{j=0}^N \mathfrak{C}_{j+M-2}^{M-2} = \mathfrak{C}_{M-2}^{M-2} + \mathfrak{C}_{1+M-2}^{M-2} + \sum_{j=2}^N \mathfrak{C}_{j+M-2}^{M-2} \\ &= \mathfrak{C}_M^{M-1} + \mathfrak{C}_M^{M-2} + \sum_{j=3}^N \mathfrak{C}_{j+M-2}^{M-2} \\ &= \mathfrak{C}_{M+1}^{M-1} + \mathfrak{C}_{M+1}^{M-2} + \sum_{j=4}^N \mathfrak{C}_{j+M-2}^{M-2} = \dots = \mathfrak{C}_{M+N-1}^{M-1} \end{aligned}$$

## Appendix 3.D Proof of Lemma 3.6.1

$$\begin{aligned} Pr[n|1]Pr[n+1|0] - Pr[n|0]Pr[n+1|1] &= \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[1|\theta]}{Pr[1]} Pr[n|\theta] \right) \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[0|\theta]}{Pr[0]} Pr[n+1|\theta] \right) - \\ &\quad \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[0|\theta]}{Pr[0]} Pr[n|\theta] \right) \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[1|\theta]}{Pr[1]} Pr[n+1|\theta] \right) \\ &= \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[1|\theta]}{Pr[1]} Pr[n|\theta] \right) \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[0|\theta]}{Pr[0]} Pr[n|\theta] \frac{(N-1-n)Pr[1|\theta]}{(n+1)Pr[0|\theta]} \right) - \\ &\quad \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[0|\theta]}{Pr[0]} Pr[n|\theta] \right) \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[1|\theta]}{Pr[1]} Pr[n|\theta] \frac{(N-1-n)Pr[1|\theta]}{(n+1)Pr[0|\theta]} \right) \\ &= \frac{(N-1-n)}{(n+1)Pr[1]Pr[0]} \left( \left( \sum_{\theta \in \Theta} Pr[\theta] Pr[1|\theta] Pr[n|\theta] \right)^2 - \right. \\ &\quad \left. \left( \sum_{\theta \in \Theta} Pr[\theta] Pr[0|\theta] Pr[n|\theta] \right) \left( \sum_{\theta \in \Theta} Pr[\theta] \frac{Pr[1|\theta]^2}{Pr[0|\theta]} Pr[n|\theta] \right) \right) < 0; \end{aligned}$$

by the Cauchy-Schwartz inequality applied to the vectors  $(\sqrt{Pr[\theta]Pr[0|\theta]Pr[n|\theta]})_{\theta \in \Theta}$  and  $(\frac{Pr[1|\theta]\sqrt{Pr[\theta]Pr[n|\theta]}}{\sqrt{Pr[0|\theta]}})_{\theta \in \Theta}$ .

## Appendix 3.E Generating Random Binary Settings

Numerical simulations used to evaluate the average performance of the mechanisms described in Section 3.6 are based on random problems generated in the following way:

- the set of possible types is randomly chosen between 2 and 20;
- for each type,  $\theta$ , the probability,  $Pr[1|\theta]$ , that the buyers observe high quality is randomly chosen between 0 and 1;

- unless otherwise specified, the number of agents is chosen randomly between 2 and 30.

## Appendix 3.F Proof of Proposition 3.6.7

The idea of the proof is to show that we can find the positive values  $x$  and  $y$  such that the payment scheme defined in Proposition 3.6.7 has only three NEQ: honest reporting, everybody reporting 0 or everybody reporting 1.

*No NEQ where  $n$  agents report according to  $s^{lie}$  and  $4-n$  agents report honestly.* From Proposition 3.6.4 we know that  $x$  and  $y$  can be found to prevent an equilibrium where all agents lie. Similarly, the incentive-compatible constraints ensure that a strategy profile where one agent always lies and three agents always report the truth cannot be a NEQ. Let us show that the profile  $s = (3 \times s^{lie}, \bar{s})$  where three agents lie and one agent reports the truth. The honest reporter observing a low quality signal will report honestly if and only if:

$$Pr[2|0]x - Pr[1|0]y > 0;$$

The same honest agent reports a positive report after observing high quality if and only if:

$$-Pr[2|1]x + Pr[1|1]y > 0;$$

However, by Lemma 3.6.1 we have  $\frac{Pr[1|0]}{Pr[2|0]} > \frac{Pr[1|1]}{Pr[2|1]}$ , so the two inequalities can never be simultaneously satisfied.

Consider the profile  $s = (2 \times s^{lie}, 2 \times \bar{s})$  where two agents lie and two agents report the truth is not NEQ. One honest reporter reports the truth if and only if:

$$\begin{aligned} (3Pr[3|0] + 2Pr[1|0])x - (3Pr[0|0] + 2Pr[2|0])y &> 0; \\ -(3Pr[3|1] + 2Pr[1|1])x + (3Pr[0|1] + 2Pr[2|1])y &> 0; \end{aligned}$$

A liar, on the other hand, reports according to  $s^{lie}$  if and only if:

$$\begin{aligned} (3Pr[0|1] + 2Pr[2|1])x - (3Pr[3|1] + 2Pr[1|1])y &> 0; \\ -(3Pr[0|0] + 2Pr[2|0])x + (3Pr[3|0] + 2Pr[1|0])y &> 0; \end{aligned}$$

All 4 inequalities are satisfied if and only if

$$\begin{aligned} 3Pr[3|1] + 2Pr[1|1] &< 3Pr[0|1] + 2Pr[2|1]; \\ 3Pr[0|0] + 2Pr[2|0] &< 3Pr[3|0] + 2Pr[1|0]; \end{aligned}$$

which is impossible.

*No NEQ where one agent always reports 1,  $n$  agents report according to  $s^{lie}$  and  $3-n$  agents report honestly.* Clearly, when all 3 agents report honestly, the agent always reporting 1 has the incentive to deviate and report 0 after observing low quality. Consider the strategy profile  $s = (s^{pos}, 2 \times \bar{s}, s^{lie})$  where one agent reports according to  $s^{pos}$ , two agents report honestly and one agent reports according to  $s^{lie}$ . For the liar,  $s^{lie}$  is an equilibrium iff:

$$\begin{aligned} -\hat{Pr}[0|0]x + \hat{Pr}[1|0]y &> 0; \\ \hat{Pr}[0|1]x - \hat{Pr}[1|1]y &> 0; \end{aligned}$$

where  $\hat{Pr}[n|o_i]$  is the probability that  $n$  out of the 2 honest reporters will observe 1, given the observation  $o_i \in Q_2$ . By Lemma 3.6.1, we have  $\frac{\hat{Pr}[1|1]}{\hat{Pr}[0|1]} > \frac{\hat{Pr}[1|0]}{\hat{Pr}[0|0]}$ , so the above inequations cannot hold simultaneously.

Consider the strategy profile  $s = (s^{pos}, \bar{s}, 2 \times s^{lie})$  where one agent reports according to  $s^{pos}$ , one agent reports honestly and two agents report according to  $s^{lie}$ . The agent reporting honestly, does so iff:

$$\begin{aligned} \hat{P}r[2|0]x - \hat{P}r[0|0]y &> 0; \\ -\hat{P}r[2|1]x + \hat{P}r[0|1]y &> 0; \end{aligned}$$

where  $\hat{P}r[n|o_i]$  is the probability that  $n$  out of the 2 liars observe 1, given the observation  $o_i \in Q_2$ . This is impossible since by Lemma 3.6.1  $\frac{\hat{P}r[0|0]}{\hat{P}r[2|0]} > \frac{\hat{P}r[0|1]}{\hat{P}r[2|1]}$ .

Similar techniques can be used to prove that no strategy profile  $(s^{neg}, n \times s^{lie}, (3-n) \times \bar{s})$  or  $(s^{neg}, s^{pos}, n \times s^{lie}, (2-n) \times \bar{s})$  can be NEQ. Therefore, the only constraint (besides the incentive-compatibility constraints) acting on the payments  $x$  and  $y$  is intended to prevent the all lying equilibrium.  $x$  and  $y$  take exactly the values described by Proposition 3.6.4.

## Appendix 3.G Proof of Proposition 3.6.8

Consider the strategy profile  $s = (\bar{n} \times \bar{s}, (N - \bar{n}) \times s^{lie}) \in \tilde{S}$  where  $\bar{n} \geq 1$  agents report honestly, and the others always lie. If  $s$  were a NEQ, an honest reporter must expect a higher payment by reporting the truth, while a liar must expect a higher payment by lying. Consider an honest reporter observing 0. She will report a negative signal if and only if  $Pr[r_{-i} = 1|0]x > Pr[r_{-i} = N - 2|0]y$ , where  $Pr[r_{-i} = 1|0]$  and  $Pr[r_{-i} = N - 2|0]$  are the probabilities that exactly 1, respectively  $N - 2$  of the remaining  $N - 1$  agents report positive signals. Exactly one of the other agents reports a positive signal when:

- all but one of the other honest reporters observes low quality, and all liars observe high quality, or
- all honest reporters observe low quality, and all but one of the liars observe high quality.

$$\begin{aligned} Pr[r_{-i} = 1|0] &= \sum_{\theta \in \Theta} Pr[\theta|0] \binom{\bar{n}-1}{1} Pr[1|\theta] Pr[0|\theta]^{\bar{n}-2} \binom{N-\bar{n}}{N-\bar{n}} Pr[1|\theta]^{N-\bar{n}} + \\ &\quad \sum_{\theta \in \Theta} Pr[\theta|0] \binom{\bar{n}-1}{0} Pr[0|\theta]^{\bar{n}-1} \binom{N-\bar{n}}{N-\bar{n}-1} Pr[1|\theta]^{N-\bar{n}-1} Pr[0|\theta] \\ &= \frac{(\bar{n}-1)!(N-\bar{n}+1)!}{(N-1)!} Pr[N-\bar{n}+1|0] + \frac{(\bar{n})!(N-\bar{n})!}{(N-1)!} Pr[N-\bar{n}-1|0] \\ &= \frac{(\bar{n}-1)!(N-\bar{n})!}{(N-1)!} \left( (N-\bar{n}+1)Pr[N-\bar{n}+1|0] + \bar{n}Pr[N-\bar{n}-1|0] \right); \end{aligned}$$

Similarly,

$$Pr[r_{-i} = N - 2|0] = \frac{(\bar{n}-1)!(N-\bar{n})!}{(N-1)!} \left( (N-\bar{n}+1)Pr[\bar{n}-2|0] + \bar{n}Pr[\bar{n}|0] \right);$$

Hence the honest reporter has the incentive to truthfully submit a negative report if and only if:

$$(N - \bar{n} + 1) \left( Pr[N - \bar{n} + 1|0]x - Pr[\bar{n} - 2|0]y \right) + \bar{n} \left( Pr[N - \bar{n} - 1|0]x - Pr[\bar{n}|0]y \right) > 0;$$

On the other hand, the honest reporter will submit a positive report after observing a high quality signal if and only if:

$$(N - \bar{n} + 1) \left( Pr[\bar{n} - 2|1]y - Pr[N - \bar{n} + 1|1]x \right) + \bar{n} \left( Pr[\bar{n}|1]y - Pr[N - \bar{n} - 1|1]x \right) > 0;$$

Exactly the same reasoning leads to the following two inequations for the liar:

$$\begin{aligned} (N - \bar{n}) \left( Pr[\bar{n} - 1|0]y - Pr[N - \bar{n}|0]x \right) + (\bar{n} + 1) \left( Pr[\bar{n} + 1|0]y - Pr[N - \bar{n} - 2|0]x \right) &> 0; \\ (N - \bar{n}) \left( Pr[N - \bar{n}|1]x - Pr[\bar{n} - 1|1]y \right) + (\bar{n} + 1) \left( Pr[N - \bar{n} - 2|1]x - Pr[\bar{n} + 1|1]y \right) &> 0; \end{aligned}$$

There exist  $x$  and  $y$  such that the four inequalities are satisfied in the same time only if:

$$\begin{aligned} \frac{(N - \bar{n} + 1)Pr[\bar{n} - 2|0] + \bar{n}Pr[\bar{n}|0]}{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|0] + \bar{n}Pr[N - \bar{n} - 1|0]} &< \frac{(N - \bar{n} + 1)Pr[\bar{n} - 2|1] + \bar{n}Pr[\bar{n}|1]}{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|1] + \bar{n}Pr[N - \bar{n} - 1|1]} \\ \frac{(N - \bar{n})Pr[\bar{n} - 1|1] + (\bar{n} + 1)Pr[\bar{n} + 1|1]}{(N - \bar{n})Pr[N - \bar{n}|1] + (\bar{n} + 1)Pr[N - \bar{n} - 2|1]} &< \frac{(N - \bar{n})Pr[\bar{n} - 1|0] + (\bar{n} + 1)Pr[\bar{n} + 1|0]}{(N - \bar{n})Pr[N - \bar{n}|0] + (\bar{n} + 1)Pr[N - \bar{n} - 2|0]} \end{aligned}$$

or equivalently:

$$\begin{aligned} \frac{(N - \bar{n} + 1)Pr[\bar{n} - 2|0] + \bar{n}Pr[\bar{n}|0]}{(N - \bar{n} + 1)Pr[\bar{n} - 2|1] + \bar{n}Pr[\bar{n}|1]} &< \frac{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|0] + \bar{n}Pr[N - \bar{n} - 1|0]}{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|1] + \bar{n}Pr[N - \bar{n} - 1|1]} \\ \frac{(N - \bar{n})Pr[\bar{n} - 1|1] + (\bar{n} + 1)Pr[\bar{n} + 1|1]}{(N - \bar{n})Pr[\bar{n} - 1|0] + (\bar{n} + 1)Pr[\bar{n} + 1|0]} &< \frac{(N - \bar{n})Pr[N - \bar{n}|1] + (\bar{n} + 1)Pr[N - \bar{n} - 2|1]}{(N - \bar{n})Pr[N - \bar{n}|0] + (\bar{n} + 1)Pr[N - \bar{n} - 2|0]} \end{aligned}$$

However, one can show that:

$$\frac{(N - \bar{n})Pr[N - \bar{n}|1] + (\bar{n} + 1)Pr[N - \bar{n} - 2|1]}{(N - \bar{n})Pr[N - \bar{n}|0] + (\bar{n} + 1)Pr[N - \bar{n} - 2|0]} < \frac{(N - \bar{n} + 1)Pr[\bar{n} - 2|0] + \bar{n}Pr[\bar{n}|0]}{(N - \bar{n} + 1)Pr[\bar{n} - 2|1] + \bar{n}Pr[\bar{n}|1]}$$

and

$$\frac{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|0] + \bar{n}Pr[N - \bar{n} - 1|0]}{(N - \bar{n} + 1)Pr[N - \bar{n} + 1|1] + \bar{n}Pr[N - \bar{n} - 1|1]} < \frac{(N - \bar{n})Pr[\bar{n} - 1|1] + (\bar{n} + 1)Pr[\bar{n} + 1|1]}{(N - \bar{n})Pr[\bar{n} - 1|0] + (\bar{n} + 1)Pr[\bar{n} + 1|0]}$$

which means that the honest reporter and the liar cannot both believe that their strategies are optimal (given the strategies of the other agents).

Consider the strategy profile  $s = (s^{neg}, \bar{n} \times \bar{s}, N - \bar{n} - 1 \times s^{lie}) \in \tilde{S}$  where one agent always reports 0,  $\bar{n} \geq 1$  agents report honestly, and  $N - \bar{n} - 1 \geq 1$  agents always lie. An honest reporter and a liar both believe that  $s$  is a NEQ if and only if:

$$\begin{aligned} & \left( \bar{n}\bar{P}r[N - \bar{n} - 2|0] + (N - \bar{n})\bar{P}r[N - \bar{n}|0] \right) x - \bar{P}r[\bar{n} - 1|0]y > 0 \\ & - \left( \bar{n}\bar{P}r[N - \bar{n} - 2|1] + (N - \bar{n})\bar{P}r[N - \bar{n}|1] \right) x + \bar{P}r[\bar{n} - 1|1]y > 0 \\ & - \left( (\bar{n} + 1)\bar{P}r[N - \bar{n} - 3|0] + (N - \bar{n} - 1)\bar{P}r[N - \bar{n} - 1|0] \right) x + \bar{P}r[\bar{n}|0]y > 0 \\ & \left( (\bar{n} + 1)\bar{P}r[N - \bar{n} - 3|1] + (N - \bar{n} - 1)\bar{P}r[N - \bar{n} - 1|1] \right) x - \bar{P}r[\bar{n}|1]y > 0 \end{aligned} \tag{3.31}$$

where  $\bar{P}r[j|o_i]$  is the probability that  $j$  out of  $N - 2$  agents observe high quality signals, given the observation  $o_i$ .

Nevertheless,

$$\frac{\bar{P}r[\bar{n}|0]}{\bar{P}r[\bar{n}|1]} < \frac{\bar{P}r[\bar{n} - 1|0]}{\bar{P}r[\bar{n} - 1|1]}$$

and

$$\frac{\bar{n}\bar{P}r[N - \bar{n} - 2|0] + (N - \bar{n})\bar{P}r[N - \bar{n}|0]}{\bar{n}\bar{P}r[N - \bar{n} - 2|1] + (N - \bar{n})\bar{P}r[N - \bar{n}|1]} < \frac{(\bar{n} + 1)\bar{P}r[N - \bar{n} - 3|0] + (N - \bar{n} - 1)\bar{P}r[N - \bar{n} - 1|0]}{(\bar{n} + 1)\bar{P}r[N - \bar{n} - 3|1] + (N - \bar{n} - 1)\bar{P}r[N - \bar{n} - 1|1]}$$

which means that the inequalities in (3.31) can never be simultaneously satisfied.

Using exactly the same technique one can show that:

- $s = (s^{pos}, \bar{n} \times \bar{s}, (N - \bar{n} - 1) \times s^{lie}) \in \tilde{S}$  where one agent always reports 0,  $\bar{n} \geq 1$  agents report honestly, and  $N - \bar{n} - 1 \geq 1$  agents always lie is not a NEQ;
- $s = (s^{neg}, s^{pos}, \bar{n} \times \bar{s}, (N - \bar{n} - 2) \times s^{lie}) \in \tilde{S}$  where one agent always reports 0, one agent always reports 1,  $\bar{n} \geq 0$  agents report honestly, and  $N - \bar{n} - 1 \geq 0$  agents always lie is not a NEQ.



## Chapter 4

# Novel Applications of Signaling Reputation Mechanisms - QoS Monitoring

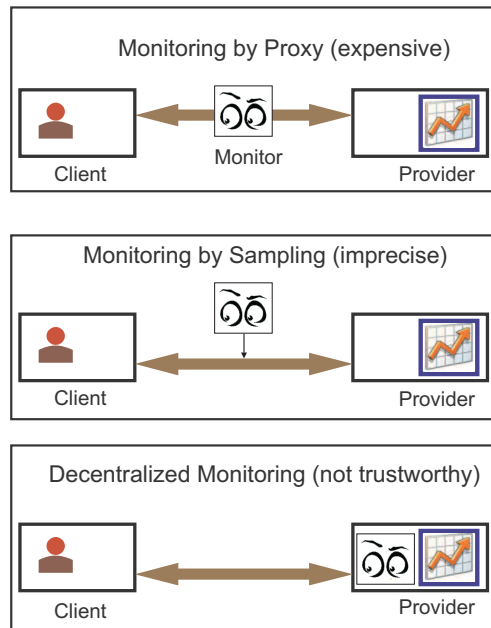
An increasing fraction of a modern economy consists of *services*. Services are generally provided under a contract that fixes the type and quality of the service to be provided as well as penalties if these are not met. We call such a contract a *Service Level Agreement* or SLA. For example, an airline provides as a service the transportation of a passenger within certain time constraints and has to pay certain penalties if this service is not delivered. Another example is providing a communication service where a certain availability and capacity is guaranteed and penalties must be paid if these are not reached.

An essential requirement for such service provisioning is to be able to *monitor* the quality of service (or QoS) that was actually delivered. As the monetary value of individual services decreases, the cost of providing accurate monitoring takes up an increasing share of the cost of providing the service itself. For example, with current technology, reliably monitoring the quality of a communication service requires constant communication with a neutral third party and would be almost as costly as providing the service itself. The cost of this monitoring remains a major obstacle to wider adoption of a service-oriented economy.

There is a large body of research addressing infrastructural facilities for SLA monitoring (Sahai et al., 2002; Ludwig et al., 2004; Dan et al., 2004; Barbon et al., 2006). Existing solutions rely on one of the following three techniques (see Figure 4.1):

- a trusted *monitor* intercepts the messages exchanged between the client and the provider and outputs an estimate of the delivered QoS.
- a trusted party periodically probes the service and outputs performance metrics.
- monitoring code runs on the provider side, as part of the service middleware. The monitoring layer intercepts the messages addressed to/originating from the provider, and estimates the delivered QoS.

The problem with the first technique is scalability. When the monitor intercepts all service invocations, it acts as a central proxy and soon becomes a performance bottleneck. Bottlenecks may of



**Figure 4.1:** Traditional QoS Monitoring Approaches.

course be avoided by only monitoring a sample of the service invocations, but the monitoring will be less precise.

The second technique is expensive and probably inaccurate. Special clients must be set up only to probe and evaluate the service. They generate supplementary service requests which unnecessarily overload service providers. Moreover, trusted clients monitor only a small sample of the total number of requests, and therefore the output results are prone to noise and errors.

The problem with the third technique is trustworthiness. The providers have obvious strategic incentives to modify the monitoring results. Unless strongly secured (which comes at a non-negligible cost), the monitoring code may be tampered with, and rendered unreliable.

This chapter has the objective of describing an alternative system and method for accurately monitoring the quality of service that has actually been delivered by a service provider. The basic principle is to have each user of the service report the quality of service she observed to a reputation mechanism. The reputation mechanism can estimate the QoS that was actually delivered by a service provider by aggregating the reports from all clients that have interacted with that provider. The QoS estimate output by the reputation mechanism can be used to check whether the provider complied with the SLA it advertised; if violations are detected, the provider may be charged penalties that are later redistributed to the community of users.

While this system can be used for a wide range of services, it is particularly suitable for monitoring services that are provided to a large group of users that are treated equally by the service provider. Different clients in the same group may receive different QoS, nevertheless, their observations can be modeled by a random signal that only depends on the overall quality delivered by the provider. From this perspective, the quality delivered by the provider within a certain time interval, to a certain client group defines the provider's fixed, unknown type. The reputation mechanism that uses the client feedback

to estimate the QoS delivered by the provider is a signaling reputation mechanism, as described in Chapter 3.

The key ingredient of the QoS monitoring system is to ensure that clients report honestly their QoS observations. As shown in the previous chapter, honest reporting incentives can be guaranteed by a payment mechanism where clients are paid for reporting their observed quality of service. To balance the incentive of clients to report bad quality of service, the payments are scaled so that truthful reports incur higher payments. These payments can be scaled to offset whatever benefit a client can obtain from misreporting. With these payments, truthful reporting becomes optimal for clients, and does not need to be further enforced. Thus, the system requires little or no direct interaction with a central monitoring authority, and can be implemented at much lower cost than previously known methods.

## Related Work

Reputation Mechanisms have emerged as efficient tools for service discovery and selection (Singh and Huhns, 2005). When electronic contracts cannot be enforced, users can protect themselves against cheating providers by looking at past behavior (i.e., the provider's reputation). Liu et al. (2004) present a QoS-based selection model that takes into account the feedback from users as well as other business related criteria. The model is extensible and dynamic. In the same spirit, Kalepu et al. (2003) proposes *verity*, a QoS measure that takes into account both reputation and the terms of the SLA. An interesting approach is proposed by Deora et al. (2003), who argue that the expectations of a client greatly influence the submitted feedback, and therefore should be used when assessing the QoS of a provider. Both Maximilien and Singh (2004) and Alunkal et al. (2003) propose concrete frameworks for service selection based on the reputation of the service provider. However, reputation-based selection gives only indirect incentives, as clients learn to avoid deceitful providers.

As opposed to the above solutions, we mainly use the feedback reported by the clients to substitute QoS monitoring. We believe that the information contained in the reports should be used directly and immediately to assess the *honesty* of the advertisement made by the provider. Moreover, this information should have direct repercussions on the gains of the provider through contractual penalties. In this way, providers get immediate incentives to exert effort.

This chapter also relates to the large body of research on monitoring and enforcing of electronic contracts. Mahbub and Spanoudakis (2004) and Robinson (2003) address the monitoring of requirements (i.e., behavioral properties), as derived from the specification of a service in a particular context. Pistore et al. (2004) propose a framework for planning the composition, and for monitoring the execution of BPEL web services. Planning techniques are exploited to synthesize a monitor of a BPEL process, which detects and signals whether the external partners behave inconsistently with the specified protocols. Barbon et al. (2006) check temporal, boolean, time-related, and statistic properties of BPEL composition, using a framework supporting RTML, the Run-time Monitoring Specification Language. Bianculli and Ghezzi (2007) address the monitoring of conversational services where the state depends on the interaction between the client and the provider. Keller and Ludwig (2002) propose WSLA and focus on monitoring QoS properties such as performance and costs. Xu and Jeusfeld (2003) and Milosevic and Dromey (2002) address the formalization and automated monitoring of contracts expressed in natural language.

Reliable information regarding the QoS of advertised services is also essential for service selection and composition. Zeng et al. (2003) and Zeng et al. (2004) present AgFlow, a middleware for quality-driven service composition. In AgFlow, the QoS of web services is evaluated using an extensible multidimensional QoS model, and the selection of individual services aims at optimizing the QoS of the composite service. Wang et al. (2006) introduces QoS-based selection of semantic web services, i.e., web services that provide well-defined, computer-interpretable semantics (McIlraith and Martin, 2003).

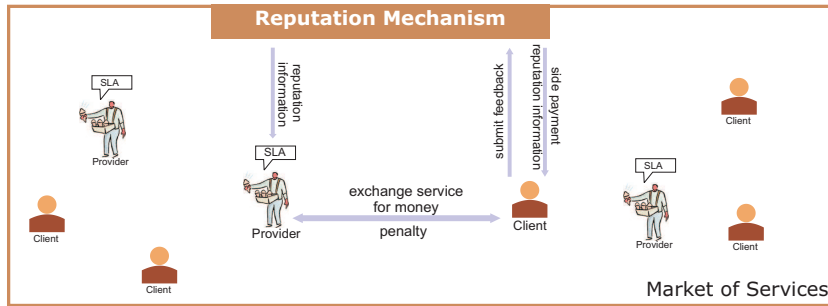


Figure 4.2: A market of web services.

## 4.1 Formal Model and Assumptions

Let us consider an online market of services pictured in Figure 4.2 where different clients interact with different service providers in a decentralized manner. There is no trusted authority or proxy intermediating the transactions between clients and providers, except that discovery is facilitated by service directories. Both clients and providers have digital identities based on public key infrastructure. The complete upper level interaction protocol is described in Section 4.2.

Services are characterized by binding SLAs specifying both functional and non-functional (quality) attributes expressed in a language such as WSDL (Dan et al., 2004) or WS-agreement (Andrieux et al., 2005). Time is divided into equal periods, and assume that the same SLA is shared by a large group of clients in any given period of time. The same service provider can have several customer groups, but all clients within the same group are treated equally (within the same period of time). The length of the time period is an application-dependent parameter, set to meet the two constraints above (i.e., large number of client requests per period, but the same SLA and service parameters).

The Quality of Service (QoS) is specified according to a common ontology. I consider *objective* quality attributes that take discrete values, and can be observed by clients for single service invocations. *ServiceIsAlive* or *InvocationFailure* are examples of such quality attributes: they are understood by all agents in the same way, can be measured for each interaction, and take Boolean values. *ResponseTime* and *Bandwidth* are both objective and observable, but usually take continuous values. For most applications, however, clients are indifferent between values that fall within some range, and therefore, they can be discretized: e.g.,  $Bandwidth \in \{DialUp, DSL, T1\}$ . On the other hand, *Availability* or *Reliability* do not meet our restriction since they are not observable for single interactions.

Second, I assume that quality properties are specified in the SLA as probability distributions over possible values for each of the quality attributes. *Availability* can therefore be indirectly expressed as a probability distribution over the boolean values of the quality attribute *ServiceIsAlive*.

Finally, the values of different quality attributes are assumed independent, with the only exception that certain values of certain quality attributes render the observation of other quality attributes impossible: e.g., if for the present invocation the *ServiceIsAlive* attribute has the value *FALSE*, the value of the *ResponseTime* attribute cannot be observed.

Formally, let  $A = \{a_1, a_2, \dots, a_{N_a}\}$  be the set of all (i.e.,  $N_a$ ) quality attributes defined by our ontology, and let  $D_j$  be the domain of values of the quality attribute  $a_j$ . Generally, the dependence between quality attributes is expressed through a (linear) correlation factor between the values of those

attributes. In our model, however, dependencies can be expressed as a relation:

$$\mathcal{R} = \{(a_j, v_h, a_k) | a_j, a_k \in A, v_h \in D_j\};$$

specifying all tuples  $(a_j, v_h, a_k)$  such that when the quality attribute  $a_j$  takes the value  $v_h \in D_j$ , the quality attribute  $a_k$  cannot be observed. For example, the response time attribute cannot have a value when the service is not alive, and therefore, the tuple  $(ServiceIsAlive, FALSE, ResponseTime)$  will be a member of this set.

A *description* of the quality attribute  $a_j \in A$  is a probability distribution  $\mu_j : D_j \rightarrow (0, 1)$  over all possible values of the attribute  $a_j$ . For example, a description of the quality attribute *ResponseTime* could be the following: “the response time is: less than 0.1s with probability 30%, less than 0.5s with probability 70%, and less than 1s with probability 100%”. A *quality advertisement* (as published by a SLA) describes a subset  $\bar{A} \subseteq A$  of quality attributes. Let  $QS$  denote the set of all possible quality advertisements, and let  $\bar{q}s$  denote a member of this set (i.e.,  $\bar{q}s$  uniquely characterizes a SLA).

After every interaction, the client observes a value for some (possibly all) of the quality attributes specified in the SLA. A *quality observation*,  $o = (v_j)$ ,  $v_j \in D_j \cup \{null\}$ ,  $a_j \in \bar{A}$ , is a vector containing a value for each of the quality attributes specified in the SLA. Since not all combinations of values can occur simultaneously (because of the constraints defined by the relation  $\mathcal{R}$ ), the quality attribute  $a_k$  will have the value  $v_k = null$  if and only if some other quality attribute  $a_j$  has the value  $v_h$ , and  $(a_j, v_h, a_k) \in \mathcal{R}$ . As quality observations give signals to the clients about the QoS delivered by the provider, we will interchangeably use the terms *signal* and *quality observation*. The set of all possible signals is denoted as in Chapter 3 by  $Q = \{q_0, q_1, \dots, q_{M-1}\}$ , where, of course, every signal is actually a vector of values, one for each attribute in the subset  $\bar{A}$ . A summary of the formal notation is given in the Appendix 3.A.

A trusted reputation mechanism (RM) is responsible for gathering and aggregating the feedback from the clients. The feedback is used to compute the delivered QoS and to update (in an application-dependent manner) the reputation information about the service provider. Feedback consists of a set of quality reports about the interactions between a client and a service provider. We assume that quality observations can be derived automatically from the messages exchanged between the client and the provider. To facilitate the reporting, the RM makes available the monitoring and reporting code that allows the clients to automatically submit feedback.

The feedback messages submitted by the clients consist of a set of quality reports about the interactions between the client and service providers. One message can thus compress information about the several transactions with several service providers. We assume that quality observations can be derived automatically from the messages exchanged between the client and the provider. To facilitate the reporting, the RM makes available the monitoring and reporting code that allows the clients to automatically submit feedback.

Although feedback is by default reported honestly, clients can tamper with the reporting code when they increase their utility by doing so. Let  $\Lambda > 0$  be an upper bound on the utility increase an agent can obtain by lying, as, for example,

- falsely reporting low quality decreases the reputation of the provider, who may be forced, in the future, to decrease the price of service;
- the decreases in reputation due to a false report may also drive away other clients, leaving the service provider more available to the requests of the lying agent;
- falsely reporting high quality could attract rewards or preferential treatment from the provider.

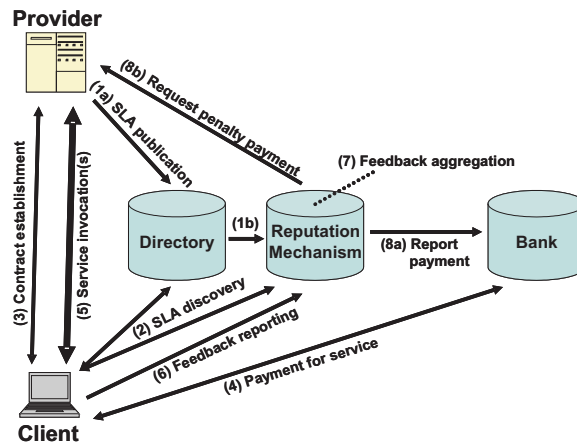


Figure 4.3: Interaction protocol involving a RM.

Tampering with the reporting code is costly, and we denote this cost by  $C$ . The same modified code can be used repeatedly or shared by several clients, and therefore, the marginal cost of one false report is often smaller than  $\Lambda$ . The potential advantage a client can obtain by lying motivates the need for measures to ensure honesty.

Unlike the traditional techniques, our approach is to make lying uninteresting, rather than impossible. We use a minimum of cryptographic tools, and propose a payment mechanism that rewards honesty. The RM will pay something for every submitted feedback, and the payments will be scaled such that, in expectation, the reward from telling the truth is better than the reward when lying by at least  $\Lambda$ . This property guarantees that no agent (or small coalition of agents) has the incentive to tamper with the reporting code.

## 4.2 Interaction Protocol

The participants in our environment are the following: *service providers* advertise SLAs and offer the corresponding services; *clients* choose SLAs and invoke the respective services; *service directories* facilitate the matching between clients and providers; *RMs* collect and aggregate feedback from the clients; a *bank* handles payments. The RMs and the bank are trusted parties. A RM can be integrated into a service directory in order to enable efficient, reputation-aware SLA selection. In this case, the service directory integrating a RM is assumed to be trusted.

Figure 4.3 illustrates the interactions between the aforementioned participants:

1. Providers advertise SLAs to a service directory (1a). Each SLA uniquely identifies the service provider (*Provider-ID*) and the service functionality, (*Serv-ID*) for example by referring to a WSDL service description, and defines the price and QoS for service invocation. The service directory assigns a suitable RM (*RM-ID*) for each SLA advertisement, which shall be used for feedback reporting. The instantiation of a RM for a new SLA (1b) requires solving a linear optimization problem, which will be discussed in Section 4.5. Advertised SLAs remain valid for a period of time specified by the provider. After expiration, they are removed from the directory. Service directories may support leases, allowing service providers to refresh SLA advertisements. Each SLA receives a unique *SLA-ID*, computed as a secure hashcode of the SLA.

2. Clients search for advertised SLAs according to functional and non-functional criteria, as well as according to reputation information. To this end, clients access a directory and a RM. If the RM is integrated within the directory, reputation-based filtering constraints can be directly included in the directory query. Clients may inspect reputation information specific to a SLA, or aggregated reputation information for a service provider.

3. The client and the chosen provider establish a contract for a given SLA, for a given period of time. The client sends a request message to the service provider, including *Client-ID*, *SLA-ID*, and the number of requested service invocations, *Nr-Invoc*. The service provider may reject the request, if it (temporarily) cannot meet the conditions of the SLA. The response message sent by the service provider is a non-forgable service invocation capability (SIC), valid for *Nr-Invoc* service invocations according to the conditions advertised in the SLA *SLA-ID*. The SIC will also be used by the client to report feedback.

4. The client pays for the agreed number of service invocations (i.e., *Nr-Invoc* times the price stated within the SLA). The payment message includes the SIC, and the bank returns the signed SIC in order to certify successful payment.

5. The client requests the service, and the provider responds. For each service invocation, the client has to provide a valid SIC signed by the bank. Hence, the service provider can easily determine that the client has payed for the SLA. The service provider keeps track of the number of service invocations for each valid SIC in order to ensure that this number does not exceed the contracted *Nr-Invoc* value. The client monitors the QoS parameters to be reported to the RM.

6. The client sends feedback to the RM. The feedback contains the SIC signed by the bank, and a timestamped series of quality reports. For each SIC, the client may send between 1 and *Nr-Invoc* reports. The quality reports need not necessarily be aggregated within a single message. I.e., for the same SIC, the client may send several messages with a varying number of quality reports. The RM does not verify whether a service was actually invoked by the client, but it ensures that the client paid for the invocation. I.e., the RM rejects reports if the SIC has not been signed by the bank.

7. The RM aggregates received feedback at the end of each time period. From all valid quality reports about a SLA, the RM estimates the actually delivered QoS by computing the distribution of values (i.e., histogram) for every quality attribute described by the SLA. Feedback can also be used to update the reputation of the service provider.

8. The RM pays valid reports as described in Section 4.5 (8a). Finally, the RM publishes the monitored QoS value for the current period and notifies the providers about the penalties they must pay (8b). Service providers who do not pay the agreed penalties may be put on a *black list* by the RM and consequently will be avoided by clients upon service selection.

## 4.3 Implementation of a Prototype

To validate the model discussed in the previous sections, we implemented a prototype of the QoS monitoring framework as a light-weight add-on on top of existing web-service middleware (Axis<sup>1</sup>). The framework exposes three types of components: directory services, reputation mechanisms, and banks. It also uses external certification authorities in order to setup a public key infrastructure (PKI). The users of the framework (i.e., the clients and the service providers) are provided with appropriate libraries in order to facilitate the deployment of applications.

---

<sup>1</sup><http://ws.apache.org/axis/>

As a general principle, all components expose two kind of interfaces:

- a web service exposing the functionality available to the users (providers and clients) of the framework. We will refer to this web service as the *public* web service (respectively the *public* interface).
- a web service exposing the functionality available to the other components within the framework. For security reasons, providers and clients do not have access to this web service. By abusing the terminology we will refer to this web service as the *private* web service (respectively the *private* interface).

The certification authority (CA) must provide the standard functionality associated to this role: creation and signing of digital X.509 certificates, validation of certificates, and revocation of expired or compromised certificates. For testing purposes we implemented a demo CA in our framework, however, any CA with a web service interface may be used. All parties (clients, providers, as well as directory services, reputation mechanisms and banks) are required to have valid identity certificates; these will be used to sign, encrypt, and authenticate exchanged messages. In our current version, we assume that CAs enforce unique identities and unique names.

The directory service is implemented as a wrapper around one or several UDDI and WSLA repositories. The public interface of the directory allows service providers to register, modify and delete service descriptions and service level agreements. Service registrations are requests by providing standard WSDL documents, signed by the provider. The directory checks the validity of the signature, and forwards the request to the UDDI repository (we used JUDDI<sup>2</sup> as the implementation of UDDI). The business key returned by the UDDI repository is returned to the provider, but is also stored by the directory next to the identity of the provider. Any subsequent modifications to existing service descriptions are first validated by the directory in order to avoid malicious corruption of WSDL documents.

Service providers may announce several SLAs for the same service. The registration of one or several SLAs is made by providing one, respectively several WSLA documents, describing the non-functional characteristics of the service. The directory first checks the validity of the business key against the identity of the provider, and then forwards the request to a proprietary WSLA repository<sup>3</sup>. The WSLA document describes the quality attributes of the service, by providing a cumulative distribution function on the values of each attribute. The quality attributes and possible values are described in an ontology.

Clients can search the directory for services that fulfill functional and non-functional requirements. Non-functional requirements are specified as a list of constraints that must be simultaneously met. Every constraint specifies a tuple  $(a_j, v_h, p_k)$ , meaning that the client expects for the quality attribute  $a_j$  a value higher than  $v_h$  with probability greater than  $p_k$ . Efficient queries of the WSLA repository can be implemented by indexing the WSLA documents according to all possible tuples  $(a_j, v_h)$ .

The private interface of the directory is used by the Bank to signal the service providers that pay (or do not pay) the required penalties. Service providers that refuse to pay the penalties are eventually placed on a black list, and are excluded from the result set returned to the clients.

Among the modules provided by our framework, the bank is the simplest one. The public interface of the bank includes the traditional operations (i.e., account creation, deposits, balance checks and withdrawals) as well as two functions required to support the interaction protocol in Figure 4.3. The first, `paySIC(SIC, signatureOfClient)` is used by clients to pay for a service invocation capability (SIC) in step 4 of the interaction protocol. The bank signs the SIC as a proof of payment, and returns

<sup>2</sup><http://ws.apache.org/juddi/>

<sup>3</sup>Our current implementation uses a MySQL database to implement the WSDL repository, with appropriate indexes to facilitate the search of services with the desired quality levels



it the client. The second, `payPenalty(bill, signatureOfProvider)` is used by providers to pay the penalties resulting from delivering lower than advertised QoS (step 8). The bills are created periodically by the reputation mechanism, and reflect the difference between the advertised and delivered quality levels. Providers can instruct the bank to automatically pay the penalty bills.

The private interface of the bank allows the reputation mechanism to announce the penalties that should be paid by a provider for not respecting the terms of the SLA. In response to such announcements the bank notifies the provider about the pending payment, or automatically pays the penalty if instructed so by the service provider.

Clients submit feedback reports by using the public interface of the reputation mechanism. One message may contain a set of reports, made up of:

- quality observations (as defined in Section 4.1) for one or several SLAs
- the corresponding SICs, signed by the bank
- the signature of the client.

The RM checks the validity of the client's signature, and verifies the signature of the bank on the SIC. All reports about a SLA beyond the number specified in the SIC are discarded.

The private interface of the reputation mechanism is used by the directory in order to query the reputation of certain service providers.


All components of the framework also include code for house-keeping operations like maintenance of databases, purging expired records, revoking expired certificates, etc. This code can run either as a separate demon process when the middleware allows it, or as part of the calls to the public or private web services.

## 4.4 Incentive-compatible Service Level Agreements

In order to deter providers from advertising higher QoS, the SLA specifies a monetary penalty that must be paid by the provider to each client at the end of a given period of time. The penalty is directly proportional to the difference between the promised QoS and the quality that has actually been delivered to the clients in a certain period of time. The SLA is incentive-compatible when the total revenue of a provider advertising a higher QoS (i.e. the price obtained by overstating the QoS minus the penalty incurred for lying) is lower than the revenue obtained by truthfully advertising the intended QoS.

**Definition 4.4.1** *A reputation-based Service Level Agreement states the following terms:*

- `per_validity`: *the period of validity. Time is indexed according to a discrete variable  $t$ ;*
- `cust_group`: *the intended customer group (e.g. silver/gold/platinum customers);*
- `QoS` (denoted as  $\bar{q}_t \in QS$ ): *the advertised quality of service;*
- `price` (denoted as  $p_t$ ): *the price of service valid for the current period;*



**Figure 4.4:** The QoS as a function of the number of requests accepted by a provider. (Experimentally determined)

- **penalty:** *the feedback-based penalty to be paid by the provider to the client for deviating from the terms of the SLA. The penalty  $\lambda_t : QS \times QS \rightarrow \mathbb{R}^+$  is a function of advertised QoS (i.e.  $\bar{q}s_t$ ) and delivered QoS (i.e. the monitored QoS,  $\hat{q}s_t$ ).  $\lambda_t(\bar{q}s_t, \hat{q}s_t) = 0$  for all  $\hat{q}s_t \geq \bar{q}s_t$  and strictly positive otherwise.*

The penalties are generally domain specific and depend on the revenues obtained and costs incurred by the provider for offering service of different qualities. For a concrete example, consider the case below

#### 4.4.1 Example of Incentive-compatible SLAs

A web service provides closing stock quotes. The web service advertises a SLA every morning and specifies the price and the quality of the service. The price is any real value, while the quality is advertised as a probability that the information (i.e., the closing quote) will be delivered before the deadline (e.g., 5 minutes after the closing time). Interested clients request the service, and then wait the answers from the service provider. They experience high quality if the answer is received before the deadline (i.e. 5 minutes after the closing time) or low quality if the answer is late or not received.

The probability of successfully answering the clients' requests depends on the infrastructure available to the provider, but also on the number of requests it must address in a given day. The infrastructure of the provider is assumed fixed; the number of requests, on the other hand, can be freely decided by the service provider every day, by declining requests once a certain threshold has been achieved. Assume that for our provider, Figure 4.4 plots the relation (experimentally determined) between the probability of successfully satisfying a request and the number of accepted requests. Let  $\phi : \mathbb{N} \rightarrow [0, 1]$  denote this function.

If  $n$  is the number of requests accepted in a given day, the provider will successfully satisfy, on the average,  $n \cdot \phi(n)$  of them. More precisely, the actually delivered QoS is assumed normally distributed around  $\phi(n)$  with a fixed variance  $\sigma^2$ .

Closing stock quotes represent mission-critical information for the clients; therefore, late or absent information attracts supplementary planning costs and lost opportunities. The price clients are willing

to pay for the service depends on the quality they expect from the provider, and is described by the function  $u : [0, 1] \rightarrow \mathbb{R}^+$ .  $u(\bar{q}s)$  is the market price for a QoS equal to  $\bar{q}s$ ; moreover, to reflect the aversion of clients to risk and failures, the price function is assumed convex.

The cost of the provider is assumed fixed, and equal to  $C$ . This value reflects the cost of maintaining the software and hardware infrastructure, and does not depend on the number of clients the provider serves. Therefore, in the day  $t$  when the provider serves  $n_t$  clients and advertises the quality  $\bar{q}s_t$ , the expected revenue of the provider is:

$$V_t(n_t, \bar{q}s_t) = E_{\hat{q}s_t} \left[ n_t \cdot (u(\bar{q}s_t) - \lambda(\bar{q}s_t, \hat{q}s_t)) - C \right];$$

where  $\hat{q}s_t \in [0, 1]$  is the quality monitored by the reputation mechanism, and  $\lambda(\bar{q}s_t, \hat{q}s_t)$  is the penalty the provider is supposed to pay to every client if the monitored quality is smaller than advertised. If clients report truthfully, the reputation mechanism will estimate the QoS delivered by the provider as:  $\hat{q}s_t = \phi(n_t) + \eta$  where  $\eta$  is the noise introduced by the environment, normally distributed around 0 with the variance  $\sigma^2$ .

**Proposition 4.4.1** *The penalty function  $\lambda : QS \times QS \rightarrow \mathbb{R}^+$  satisfying the property:*

$$\frac{\partial \lambda(qs_1, qs_2)}{\partial qs_1} \geq 2u'(qs_1); \quad \forall qs_1, qs_2 \in [0, 1];$$

*makes the SLA incentive compatible.*

PROOF. Let  $(n^*, \bar{q}s^*) = \arg \max_{(n_t, \bar{q}s_t)} V_t(n_t, \bar{q}s_t)$  be the optimal number of client requests accepted by the provider, and the optimal QoS advertised to the clients. Assuming the provider charges the market price, the first order optimality condition on  $\bar{q}s^*$  becomes:

$$\begin{aligned} \frac{1}{n_t} \frac{\partial V_t}{\partial \bar{q}s_t}(n^*, \bar{q}s^*) &= u'(\bar{q}s^*) - E \left[ \frac{\partial \lambda}{\partial \bar{q}s_t}(\bar{q}s^*, \phi(n^*) + \eta) \right] \\ &= u'(\bar{q}s^*) - \int_{qs < \bar{q}s^*} \text{normpdf}(qs | \phi(n^*), \sigma) \frac{\partial \lambda}{\partial \bar{q}s_t}(\bar{q}s^*, q) dq \\ &= 0; \end{aligned}$$

where  $\text{normpdf}(q | \phi(n^*), \sigma)$  is the normal probability distribution function with the mean  $\phi(n^*)$  and the variance  $\sigma^2$ .


By replacing the condition on  $\lambda$  from the proposition, we have:

$$\int_{qs < \bar{q}s^*} \text{normpdf}(qs | \phi(n^*), \sigma) \frac{\partial \lambda}{\partial \bar{q}s_t}(\bar{q}s^*, q) dq \leq 0.5;$$

or in other words, the normal cumulative probability distribution  $Pr[qs < \bar{q}s^* | \phi(n^*), \sigma] < 0.5$ . This is only true when  $\bar{q}s^* \leq \phi(n^*)$ , so that the provider does not accept more client requests than he can successfully satisfy with the advertised QoS,  $\bar{q}s^*$ . The penalty function  $\lambda$  therefore encourages service providers to truthfully advertise the quality they intend to deliver. ■

For a price function  $u(qs) = qs^2$ , the fixed cost  $C = 100$ , the standard deviations  $\sigma = 5\%$ , and a penalty function  $\lambda(\bar{q}s, \hat{q}s) = 2(u(\bar{q}s) - u(\hat{q}s))$ , Figure 4.5 shows the optimal revenue of the provider as a function of  $n$ . The optimal value of the payoff function is reached for  $n^* = 681$ , when  $\bar{q}s^* = 0.858 = \phi(681)$ , and the provider indeed delivers the QoS he declared.

For the general case, clients can check the constraint on the penalty function by analyzing the previous transactions concluded in the market. For every previously negotiated SLA  $(\bar{q}s_t, p_t, \lambda_t)$ , clients



**Figure 4.5:** The revenue function of the provider depending on the number of accepted requests.

infer that the market price corresponding to  $\bar{q}s_t$  must be higher than  $p_t$ : i.e.  $u(\bar{q}s_t) \geq p_t$ . Previous interactions thus establish a lower bound on the real market price that can be used to safe-check the validity of the penalty function. Note that the proof above does not make any assumptions about the market price or the cost function of the providers. Therefore, such incentive-compatible SLAs can be used for a variety of settings.

All service providers have the incentive to minimize the penalty function specified by the SLA. This happens when the constraint in Proposition 4.4.1 is satisfied up to equality. As an immediate consequence, all service providers advertise exactly the intended QoS, and will only pay minimum penalties due to noise introduced by the environment.

## 4.5 Reliable QoS Monitoring

The reliability of the monitoring mechanism relies on three assumptions:

- clients submit honest feedback;
- clients are able to submit feedback only after having interacted with the provider;
- clients submit only one feedback per interaction.

The second and third assumptions can be implemented through cryptographic mechanisms based on a public key infrastructure. As part of the interaction, providers can deliver secure certificates that can later be used by clients to provide feedback. A concrete implementation of such a security mechanism for reputation mechanisms is described by Jurca and Faltings (2003).

The first assumption, however, can be integrated into the broader context of truthful feedback elicitation. The problem can be solved by side-payments (i.e. clients get paid by the reputation mechanism for submitting feedback) which make it optimal for rational agents to report the truth.

Chapter 3 described the general theory behind building truthful reward mechanisms. In this section, however, I will also present a simplified payment mechanism that is more suited for the current setting.

Such payments can be constructed by comparing every report with some other report (called the *reference report*) submitted by a different client about the same SLA. As the two reports refer to the same service, there is a link between them; this link will be exploited such that whenever the reference report is true, it also becomes in the reporter's best interest to report the truth. Honest reporting thus becomes a Nash equilibrium in our environment (see Section 3.3).

The simplest payment rule pays a report only if it matches (i.e., has the same value as) the reference report. For example, a negative report about the SLA that describes only the attribute *ServiceIsAlive* is paid only if the reference report is negative as well. However, the payments depend on the actual value of report, and a negative report is paid differently from a positive report.

The reason why such payment rules encourage truthful reporting can be explained by the subtle changes in beliefs triggered by the private experience of a client with a given service provider. Although clients know that the service provider has all the incentives to deliver the promised QoS, they also realize that the delivered QoS will only *in expectation* equal the advertised one. Environmental noise and other unpredictable events will perturb the delivered QoS, making it higher in some rounds, and smaller in others. This is why the current experience of a client also conveys information about the immediately future interactions of other clients with the same provider.

It is therefore an acknowledged empirical fact (Prelec, 2004) (also in full accordance with Bayesian theory) that an agent's posterior belief about the observation of another client (that receives the service in the same conditions) depends on the private experience of the agent.

For example, a client that has just had a negative experience believes that the clients in the same round will probably experience similar problems. Thus, she expects that the reference report used to compute her payment from the RM will correspond to a QoS that is slightly lower than advertised. On the contrary, a satisfied client is more likely to believe that other clients will be satisfied as well; therefore, she expects a reference report corresponding to slightly higher QoS than advertised.

The payments are designed such that a negative report maximizes the expected return only when clients expect a negative reference report with probability higher than advertised. And vice-versa for a positive report. Given that the reference report is true, the client maximizes her returns by reporting honestly, which makes truth-telling a Nash equilibrium. This means that no agent can gain an advantage by deviating from the protocol. Miller et al. (2005) and Jurca and Faltings (2006) present a game theoretic analysis of such reporting scenarios and show that it is always possible to choose payments that make truth-telling optimal.

Concretely, the RM computes the payment made to every client in the following way. At the end of every period, all quality reports about the same SLA are grouped in a single set. Remember that each report corresponds to a quality observation, and therefore consists of an array of values, each corresponding to a quality attribute advertised by the provider.

For each report,  $r = (v_j^r)$ , the RM randomly chooses a reference report,  $ref = (v_j^{ref})$ , coming from a different client. Every pair of matching non-null values for the attribute  $a_j$  (i.e.,  $v_j^r = v_j^{ref} \neq null$ ) contributes with  $\tau_j(v_j^r)$  to the payment for the report  $r$ .

If  $r_i = (r^i(1), \dots, r_i(N_i))$  are all the reports submitted by client  $i$ , and  $ref^i$  denotes the set of corresponding reference reports chosen by the reputation mechanism to compute the payment due to client  $i$ , the total payment received by  $i$  is:

$$Pay(i) = \sum_{h=1}^{N_i} Pay(r^i(h), ref^i(h)); \quad Pay(r, rr) = \sum_j \tau_j(v_j^r, v_j^{ref});$$

$$\text{where: } \tau_j(v_j^r, v_j^{ref}) = \begin{cases} 0 & \text{if } v_j^r \neq v_j^{ref} \text{ or } v_j^r = null \\ \tau_j(v_j^r) & \text{if } v_j^r = v_j^{ref} \end{cases}$$

The payment mechanism is fully specified by announcing the amounts  $\tau_j(v_j)$ , paid for a report matching the reference report on the value  $v_j$  of the quality attribute  $a_j$ .

We compute the payment mechanism through *automated mechanism design* (Conitzer and Sandholm, 2002). Instead of a closed form specification, we define the mechanism through a set of constraints that act on the decision variables (i.e., the payments  $\tau(\cdot)$  in our case). By adding an objective function, we get an optimization problem that solves for the best possible payment mechanism in a given context.

The optimal payment mechanism minimizes the total cost of the RM, while guaranteeing that honesty is better than lying by at least the desired margin. The cost of the RM will depend on the SLA, so the payment mechanism must be instantiated for every SLA.

The expected cost for an honest report equals the weighted sum of all amounts  $\tau_j(v_h)$ . The probability that payment  $\tau_j(v_h)$  is made equals the probability that both the report and the reference report have the value  $v_h$  for the quality attribute  $a_j$ . Since each probability equals  $\mu_j(v_h)$  (described by the SLA) and the two events are assumed independent, we have:

$$E[Cost] = \sum_{a_j \in \bar{A}} \sum_{v_h \in D_j} \tau_j(v_h) \mu_j(v_h)^2; \quad (4.1)$$

To compute the expected revenue obtained by a client when lying or telling the truth, we must first describe the belief of a client regarding the reference report chosen by the reputation mechanism. Given the real quality observation  $o = (v_j)$ , we assume that the belief regarding the reference report changes slightly in the direction of  $o$ . If  $\mu_j(v_h)$  is the advertised probability that the attribute  $a_j$  takes the value  $v_h \in D_j$ , the belief of the client assigns:

- at least the probability  $\mu_j(v_h) + (1 - \mu_j(v_h))\gamma_{max}$  to the event that the reference report also has the value  $v_h$  for the attribute  $a_j$ , and,
- at most the probability  $\mu_j(v_k)(1 - \gamma_{min})$  to the event that the reference report has some other value  $v_k \neq v_h$  for the attribute  $a_j$ .

If  $v_h = null$  (no observation was possible for the attribute  $a_j$ ), we assume that the beliefs regarding the reference report remain unchanged. Both  $\gamma_{max}$  and  $\gamma_{min}$  take values between 0 and 1, and depend on the specific applications.

Honest reporting can be guaranteed when:

- for any quality attribute  $a_j$ , truthfully reporting maximizes the expected payment by at least  $\Lambda$ :

$$[\mu_j(v_h) + (1 - \mu_j(v_h))\gamma_{max}] \tau_j(v_h) > [\mu_j(v_k)(1 - \gamma_{min})] \tau_j(v_k) + \Lambda; \quad (4.2)$$

for all quality attributes  $a_j \in \bar{A}$  and all values  $v_h \neq v_k \in D_j$ ,

- dependencies between quality attributes do not break the honest reporting incentives (i.e., since matching *null* values do not contribute to the payment, the payments for the values that cause the *null* reports must be large enough):

$$[\mu_j(v_h) + (1 - \mu_j(v_h))\gamma_{max}] \tau_j(v_h) > [\mu_j(v_k)(1 - \gamma_{min})] \tau_j(v_k) + \mu_l(v_m) \tau_l(v_m) + \Lambda; \quad (4.3)$$

for all  $a_j, a_l \in \bar{A}$ ,  $v_h \neq v_k \in D_j$ ,  $v_m \in D_l$  such that the value  $v_h$  of the attribute  $a_j$  makes impossible the observation of the attribute  $a_l$ : i.e.,  $(a_j, v_h, a_l) \in \mathcal{R}$ .

The margin  $\Lambda$  must offset the worst case incentive for lying. This value is very conservative and can be relaxed in real applications by considering that not all lies can simultaneously attract the worst case incentives.

The objective function in (4.1) together with the constraints defined by (4.2) and (4.3) define a linear optimization problem that accepts as a solution the cheapest incentive-compatible payment mechanism for a given SLA. The number of variables is equal to the overall number of values in all domains  $D_j$ , i.e.,  $\sum_{a_j \in \bar{A}} \text{card}(D_j)$ . The number of constraints is on the order of  $\sum_{a_j \in \bar{A}} \text{card}(D_j)^2$ .

Compared to the general payments described in Chapter 3 (Sections 3.3, 3.4 and 3.6) the payments presented above implement several simplifications. First, they reward a reporter only for matching values between the report and the reference report. This decision was taken to reduce the complexity of the payment scheme. The number of possible “quality signals” observed by the clients may be huge. For example, if clients observe 5 quality attributes, each having 3 values, the set of possible observations contains  $3^5$  elements (every possible combination of values for the 5 attributes). In this case, general payments that consider all possible combinations of even 2 reports must be described by  $3^{5^2}$  values. The simplification to pay only matching reports eliminates a quadratic factor in the design process.

Second, the payments are based on comparing a report with a single reference report. In Sections 3.4 and 3.6 we have seen that the use of several reference reports both decreases the budget required by the RM, and helps to prevent lying coalitions. Nevertheless, the increased complexity of the design process forces us to limit the payments to using only one reference report.

Extending this framework to include other types of dependencies or correlations between quality attributes does not pose theoretical challenges. Optimal payments that ensure truthful reporting can still be computed by extending the optimization problem with constraints like (4.3) that limit the gains of lying on pairs of values for correlated quality attributes. Intuitively, the additional constraints isolate independent *groups* of attributes, and guarantee truth-telling incentives. However, the notation that allows the definition of such payments is complicated, and outside the scope of this thesis.

The payments naturally decrease with the margins,  $\Lambda$ , required for truth-telling. Therefore, the expected cost of the RM can be decreased either by decreasing the benefits clients may obtain by manipulating their reports, or, by increasing the cost of tampering the reporting code. While the latter direction is outside the scope of this thesis, the following two ideas can be used to address the former.

First, we can make sure that the penalties paid back by the provider do not give incentives to underrate the quality of the service. For that, we impose that the penalty paid to client  $i$  depends only on the QoS delivered to all other clients, except  $i$ . Given a large enough number of agents, the penalties paid by the provider are expected to be the same, but the feedback reported by  $i$  has no influence on the penalties paid to  $i$ .

A second interesting direction is to filter out the reports that are very far from the common distribution (see Section 3.4). Intuitively, these reports are either erroneous, or intentionally try to introduce significant perturbations towards desired values.

### 4.5.1 Example

Let us modify the example from Section 4.4 in the following way. The QoS described in the SLA now addresses two attributes: *availability*, (i.e., the probability that a request is answered before a deadline  $t_d$ ) and *correctness* the probability of returning the correct). From a practical point of view it makes perfect sense to have both criteria used together, since otherwise a service can achieve almost perfect availability by always returning the same information.

Assume the SLA advertises the value  $p_1 \in [0, 1]$  for *availability* and the value  $p_2$  for *correctness*. Formally, this SLA is expressed as the probability distribution  $\mu_1 = \{p_1, 1-p_1\}$  for the quality attribute:

$$a_1 = \text{ResponseBeforeDeadline} \in D_1 = \{0(\text{false}), 1(\text{true})\};$$

and the probability distribution  $\mu_2 = \{p_2, 1-p_2\}$  for the quality attribute:

$$a_2 = \text{InformationIsCorrect} \in D_2 = \{0(\text{false}), 1(\text{true})\};$$

Naturally, the relation  $\mathcal{R}$  defining the dependency between quality attributes contains only the tuple  $(a_1, 0, a_2)$ : if no response is received, checking for correct information is meaningless.

A quality observation (and therefore a quality report) is a vector  $o = (v_1, v_2)$  where  $v_1 \in \{0, 1\}$  and  $v_2 \in \{0, 1, \text{null}\}$ . The payment scheme used by the RM is defined by the four positive amounts  $\tau_1(1), \tau_1(0), \tau_2(1)$  and  $\tau_2(0)$ , paid when the non-null value of the attribute  $a_1$  or  $a_2$  matches the corresponding value of the reference report. The maximum benefit a client can obtain by misreporting one observation is  $\Lambda = 0.01$  (all values hereafter are normalized to the price of service, assumed 1), and the cost of tampering with the default monitoring code is  $C = 10$ . A client is assumed to generate at most  $N = 1000$  service requests within the same period of time, so the worst case truth-telling margin that must be enforced by the RM is  $\bar{\Lambda} = \Lambda - C/N/2 = 0.5\%$ .

The belief of one client regarding the value of the reference report changes by at least  $\gamma_{max} = \gamma_{min} = 20\%$  in the direction of the actual observation. The probability that the reference report contains 1 for  $a_1$  is:  $Pr_1[1|1] = p_1 + (1-p_1)\gamma_{max}$  if the client also received a response, or  $Pr_1[1|0] = p_1 - (1-p_1)\gamma_{min}$  if the client did not receive a response. Similar equations can be written for the probabilities  $Pr_2[1|1]$  and  $Pr_2[1|0]$  defining the beliefs regarding the value of the attribute  $a_2$  in the reference report. From (4.1), (4.2) and (4.3), Figure 4.6 presents the linear optimization problem that defines the minimum payments guaranteeing the truth-telling equilibrium. When  $p_1 = p_2 = 90\%$ , we obtain the payments:  $\tau_1(1) = 0.064$ ,  $\tau_1(0) = 0.680$ ,  $\tau_2(1) = 0.025$ ,  $\tau_2(0) = 0.225$ , and an expected cost of 0.081. These rather high payments can be further decreased by an order of magnitude using a filtering mechanism as in Section 3.4: (i.e., the RM probabilistically selects the reports that will contribute towards estimating delivered quality). The expected payments can thus be brought down to below 1% of the price of service, which is a practical value.

## 4.6 Deterring Malicious Coalitions

The payments defined in the previous section do not have truthful reporting as the unique equilibrium. Always reporting the same values is also an equilibrium strategy, since reports will surely match the corresponding reference reports. Moreover, it is easy to see that such *constant reporting strategies* yield higher payments than the truthful reporting. Fortunately, such coalitions on lying strategies can be rendered unprofitable when a fraction of reports are guaranteed to be honest.

We believe it is reasonable to rely on some fraction of truthful reports for several reasons. First, empirical studies show that a non-negligible fraction of users are altruists who always report the truth.



$$\begin{aligned}
\min \quad & E[Cost] = p_1^2 \tau_1(1) + (1 - p_1)^2 \tau_1(0) + p_2^2 \tau_2(1) + (1 - p_2)^2 \tau_2(0); \\
\text{s.t.} \quad & Pr_1[1|1] \tau_1(1) > Pr_1[0|1] \tau_1(0) + \bar{\Lambda}; \\
& Pr_1[0|0] \tau_1(0) > Pr_1[1|0] \tau_1(1) + \bar{\Lambda}; \\
& Pr_2[1|1] \tau_2(1) > Pr_2[0|1] \tau_2(0) + \bar{\Lambda}; \\
& Pr_2[0|0] \tau_2(0) > Pr_2[1|0] \tau_2(1) + \bar{\Lambda}; \\
& Pr_1[0|0] \tau_1(0) > Pr_1[1|0] \tau_1(1) + p_2 \tau_2(1) + \bar{\Lambda}; \\
& Pr_1[0|0] \tau_1(0) > Pr_1[1|0] \tau_1(1) + (1 - p_2) \tau_2(0) + \bar{\Lambda}; \\
& \tau_1(1), \tau_1(0), \tau_2(1), \tau_2(0) \geq 0;
\end{aligned}$$

**Figure 4.6:** Linear optimization problem defining the payment mechanism.

Second, given that the framework already provides the default (honest reporting) code, some clients won't have the knowledge to temper with the reporting code even if they want to.

The idea behind fighting lying coalition is to make them unstable. We start from the assumption that at most  $\eta_{col} \in (0, 1)$  percent of the clients can collude on a lying strategy. Then we compute a payment scheme that makes it individually better for a colluder to shift to the honest reporting strategy, knowing that  $1 - \eta_{col}$  percent of the reports are honest. Since the other coalition members cannot detect (and punish) deviators, all rational colluders will break the coalition and report honestly. The coalition is unstable in the sense that it is not profitable for coalition members to keep their commitment to the coalition.

Let us analyze the additional constraints on the optimization problem defining the payments. An honest reporter now expects that the reference report will be part of a coalition with probability at most  $\eta_{col}$ . To make sure that the client still has the incentive to truthfully report the observed value  $v_h$  instead of the collusion value  $v_k$ , the constraint in (4.2) becomes:

$$\begin{aligned}
(1 - \eta_{col}) [\mu_j(v_h) + (1 - \mu_j(v_h)) \gamma_{max}] \tau_j(v_h) > \\
(1 - \eta_{col}) [\mu_j(v_k)(1 - \gamma_{min})] \tau_j(v_k) + \eta_{col} \tau_j(v_k) + \Lambda;
\end{aligned} \tag{4.4}$$

for all quality attributes  $a_j \in \bar{A}$  and all values  $v_h \neq v_k \in D_j$ . Similarly, the constraint (4.3) becomes:

$$\begin{aligned}
(1 - \eta_{col}) [\mu_j(v_h) + (1 - \mu_j(v_h)) \gamma_{max}] \tau_j(v_h) > \\
(1 - \eta_{col}) \left[ [\mu_j(v_k)(1 - \gamma_{min})] \tau_j(v_k) + \mu_l(v_m) \tau_l(v_m) \right] + \eta_{col} (\mu_j(v_k) + \mu_l(v_m)) + \Lambda;
\end{aligned} \tag{4.5}$$

for all quality attributes  $a_j, a_l \in \bar{A}$ , values  $v_h \neq v_k \in D_j$ ,  $v_m \in D_l$ , and tuples  $(a_j, v_h, a_l) \in \mathcal{R}$ .

The linear problem that minimizes (4.1) under the set of constraints (4.4) and (4.5) defines the incentive-compatible payments that are also  $\eta_{col}$ -collusion proof.

For the example in Section 4.6, we can obtain a payment mechanism that is also robust against coalitions that cover at most a fraction  $\eta_{col}$  of the reports. The dependence of the expected cost on  $\eta_{col}$  is plotted in Figure 4.7.




Figure 4.7: Expected cost of a payment mechanism that is robust against collusion.

#### 4.6.1 Using Trusted Monitoring Infrastructure

Another way of ensuring the reliability of feedback information is to use *trusted reports* as reference reports when computing the payments due to the clients. The RM can use the trusted monitoring infrastructure that already exists in markets of services and can probatively get honest reports about a certain provider.

Without giving all the technical details, a trusted monitoring system works in two ways:

- it acts as a proxy between the client and the provider, and intercepts all messages before relaying them to the intended destination. From the conversation, the monitor can deduce the QoS the provider delivered to the client;
- the monitor itself becomes a normal client of the provider and requests the service with the sole purpose of rating it. The monitor, however, must maintain its anonymity so that service providers cannot preferentially treat the monitor.

Both methods generate truthful, trusted feedback that can be used by the RM to encourage normal clients to report honestly. Nevertheless, trusted reports come at a cost, so it is important to use as few as possible.

From a game theoretic point of view, one trusted report for every round is enough to enforce honest reporting. If all other reports are compared (and paid) against the trusted one, the incentive-compatibility constraints ensure that honest reporting is the unique optimal strategy of rational clients. However, this procedure creates strong misbehavior incentives for the reputation mechanism itself. By faking the value of the trusted report, or by strategically choosing one of several available reports, the RM can significantly reduce the payments to the buyers. The misbehavior of the mechanism cannot be detected unless some party is able to verify all trusted reports received by the mechanism.

Dealing with selfish reputation mechanisms is a complex problem in itself, beyond the scope of this chapter. However, we believe that any trustworthy implementation of a RM must ensure three properties:

- integrity of submitted reports;

- traceable rules for assigning reference reports;
- upper limits on the number of times a reference report can be used.

The first can be addressed by digital signatures, and prevents the mechanism (or anybody else) to modify the content of submitted information. Agents use their private keys to digitally encrypt submitted feedback messages (Jurca and Faltings, 2003); all agents can thus read the feedback, but no one can modify the content without being detected.

The second property prevents the RM to manipulate the choice of reference reports, without being eventually disclosed. Technical solutions can be envisaged based on secure multi-party computation protocols (Goldreich, 1998).

Finally, the third property ensures that the RM does not have the incentive to keep acquiring trusted reports until it gets one that minimizes the total payments returned to the clients. Assume, for example, that the cost of a trusted report is  $C_T$ , that the RM already has one trusted report, and that a trusted report can be used as a reference report at most  $N_R$  times. The RM can choose to use the available report and “score”  $N_R$  other reports, or can choose to spend  $C_T$  and acquire another trusted report which will hopefully lead to smaller payments to the  $N_R$  clients that will be scored against the new trusted report. By choosing to buy another trusted report, the RM can hope to obtain a decrease in payments of at most  $N_R \cdot \bar{\tau}$ , where  $\bar{\tau}$  is the maximum payment that can be made to a reporter. By setting:

$$N_R \leq \lfloor \frac{C_T}{\bar{\tau}} \rfloor;$$

the most optimistic payment saving obtained from purchasing a fresh trusted report is always smaller than the cost of the report, and the RM does not have the incentive to cheat.

The question now becomes how many trusted reports should the reputation mechanism purchase in every round to ensure that rational clients report honestly. Clearly, if there are  $N$  feedback reports in one round,  $N/N_R$  trusted reports will allow the RM to score every feedback against a trusted report. However, the number of trusted reports can be decreased even further by observing that  $\eta_{col}$ -collusion proof incentive compatible payments need a trusted report as a reference report only with probability  $\eta_{col}$  (see the constraints (4.4) and (4.5)).

If all  $N$  feedback reports obtained in a round are coming from self-interested clients, the problem of the mechanism designed becomes how to choose  $\eta_{col}$  and the corresponding  $\eta_{col}$ -collusion proof payment mechanism such that its total cost is minimized. The total expected cost to the RM is given by:

- $N$  times the expected payment to one client (under the  $\eta_{col}$ -collusion proof payment mechanism);
- $\lceil \frac{N \cdot \eta_{col}}{N_R} \rceil$  times  $C_T$ , the cost for acquiring the trusted reports;

The design problem is therefore defined by the following optimization problem:

$$\begin{aligned}
 \min \quad & E[Cost] + C_T \lceil \frac{N \cdot \eta_{col}}{N_R} \rceil; \\
 \text{s.t.} \quad & E[Cost] \text{ is defined by Eq. (4.1);} \\
 & \text{payments } \tau \text{ satisfy the constraints (4.4) and (4.5);} \\
 & N_R = \lfloor \frac{C_T}{\bar{\tau}} \rfloor \text{ where } \bar{\tau} \text{ is the maximum payment to one reporter;} \\
 & \text{the payments } \tau \text{ are positive, } \eta_{col} \in [0, 1);
 \end{aligned} \tag{4.6}$$

Another interesting observation is that a  $\eta_{col}$ -collusion proof payment mechanism requires at least a  $\eta_{col}$  colluding fraction to change the equilibrium in the market from honest reporting to some other lying strategy. This opens new opportunities for the efficient operation of reputation mechanisms. A group of clients synchronized on the honest reporting strategy tend to keep reporting the truth as long as no significant group of colluders wishes to shift the equilibrium. This property can be exploited by mechanism operators to reduce the running cost of the mechanism. Trusted reports should be used whenever new services enter the market. After several rounds, the mechanism can stop purchasing trusted reports, knowing that the clients continue to report honestly. Active enforcement should be resumed only when the mechanism has strong reasons to believe that a significant coalition tries to shift the equilibrium.

The asynchronous involvement of the reputation mechanism is also advantageous from a practical perspective. When entering the market, service providers typically have less clients; the mechanism therefore needs few trusted reports to install the honest equilibrium. As the customer base of the provider increases, new commeres will adopt the norms already existing in the market, and report the truth.

## 4.7 Summary of Results

A service market based on SLAs between service providers and clients can only function effectively if advertised SLAs are credible. However, service providers may deviate from their advertised QoS in order to reduce the costs of service provisioning. Hence, QoS monitoring is essential, but neither service-side nor client-side monitoring can be trusted.

In this chapter we presented a novel approach to achieve objective QoS monitoring by aggregating quality ratings from clients within a RM, which provides incentives for the clients to report honestly. The RM pays clients for submitting quality ratings, and the payments are designed such that lying generates expected losses that offset the potential benefits from misreporting.

The mechanism we present can also resist collusion. As long as a small enough fraction of agents collude, no lying strategy can give colluders a better payoff than honest reporting. Moreover, the reputation mechanism can also buy *trusted reports* from a traditional, centralized monitoring system. Few such trusted reports can help the RM synchronize all clients in the market on reporting the truth. Once honest reporting is the equilibrium in the market, it takes a significant fraction of colluders to shift the reporting equilibrium; therefore, the RM can stop purchasing honest reports, and resort to a passive monitoring of the system.

The QoS monitoring system described in this chapter is cheaper, more accurate, and more secure than existing solutions. First, the cumulated payments the RM has to pay to the clients for their feedback can be made low enough such that the total cost of the monitoring framework can be made several times cheaper when compared to traditional monitoring. Second, the RM will get information about most transactions, without placing important communication burdens on the network (clients can wrap several feedback reports in the same message, and do not have real-time constraints on the communication process, so that feedback messages can be transmitted whenever there is free capacity in the network). Central monitors, on the other hand, usually do not afford to monitor every interaction because of the high cost associated to proxying every transaction. Last but not least, the system is safe, since economical measures make dishonest reporting irrational.

## Appendix 4.A Summary of Notation

$A = \{a_1, \dots, a_{N_a}\}$ $N_a$ $D_j$ $v_h \in D_j$	<p>quality attributes defined by an ontology;</p> <p>total number of quality attributes;</p> <p>domain of values of attribute <math>a_j</math>;</p> <p>values belonging to the domain of attribute <math>a_j</math>;</p>
$\mathcal{R}$	$\mathcal{R} = \{(a_j, v_h, a_k)   a_j, a_k \in A, v_h \in D_j\}$ is a relation describing infeasible combinations of (attribute, value) pairs. $(a_j, v_h, a_k) \in \mathcal{R}$ iff the quality attribute $a_k$ cannot be observed when the quality attribute $a_j$ takes the value $v_h \in D_j$ ; ;
$\mu_j \in \Delta(D_j)$ $\bar{A}$ $\bar{q}_s = (\mu_1, \dots, \mu_{N_a})$ $QS$	<p>a description of the attribute <math>a_j</math>;</p> <p>a subset of <math>A</math>;</p> <p>quality of service;</p> <p>the set of possible QoSs;</p>
$o = (v_j), v_j \in D_j \cup \{null\}$ $Q = \{q_0, q_1, \dots, q_{M-1}\}$	<p>a quality observation, expressed as a vector of values, one for each attribute <math>a_j \in \bar{A}</math>;</p> <p>the set of all quality observations;</p>
$\Lambda$ $C$	<p>upper bound on the utility an agent can obtain by lying;</p> <p>the cost of modifying the default monitoring software;</p>
$\lambda(qs_1, qs_2)$ $r_i = (r_i(1), \dots, r_i(N_i))$ $ref_i = (ref_i(1), \dots, ref_i(N_i))$ $\tau_j(v_h)$ $E[Cost]$ $\gamma_{max}, \gamma_{min}$ $\eta_{col} = \frac{N_{col}}{N}$	<p>the penalty paid by the provider he advertises QoS <math>qs_1</math> but delivers <math>qs_2 &lt; qs_1</math>;</p> <p>all <math>N_i</math> reports submitted by agent <math>i</math>;</p> <p>the corresponding reference reports;</p> <p>the contribution to the payment when the reference report has the matching value <math>v_h \in D_j</math> for the quality attribute <math>a_j</math>;</p> <p>the expected payment to an honest reporter;</p> <p>thresholds for modified beliefs;</p> <p>the fraction of colluders;</p>
$C_T$ $N$ $N_R \leq \lfloor \frac{C_T}{\bar{\tau}} \rfloor$	<p>the cost of obtaining a trusted report;</p> <p>the total number of reports received by the RM in one round;</p> <p>the maximum number of times a trusted report can be used as a reference report;</p>



## Chapter 5

# Sanctioning Reputation Mechanisms

As opposed to signaling reputation mechanisms that bridge the asymmetry of information that usually exists in online trading situations, the role of *sanctioning* reputation mechanisms is to encourage cooperative behavior by inflicting indirect punishments on the users who cheat. We have seen that signaling reputation mechanisms are mainly useful in markets where the same product or service is experienced by a set of clients. The essential idea is the buyers do not know the real quality of the product, however, they do expect to be treated equally: i.e., they buy exactly the same product, or they receive a service provisioned with the same resources and effort.

This is the case, for example, in a market of web-services. Different providers possess different hardware resources and employ different algorithms; this makes certain web-services better than others. Nevertheless, all requests issued to the same web-service are treated by the same program. Some clients might experience worse service than others, but these differences are random, and not determined by the provider. This assumption of uniform quality (up to some random noise) allows future buyers to learn from the feedback of past users, and estimate with better precision the unknown quality of the service.

There are, however, online markets where different clients are not treated equally by the same seller. Think of a barber, who must skillfully shave every client that walks in his shop. The problem here is that providers must exert care (and costly effort) for satisfying every service request. Good quality can result only when enough effort was exerted, but the provider is better off by exerting less effort: e.g., clients will anyway pay for the shave, so the barber can profit by doing a sloppy job as fast as possible in order to have time for more customers.

The same situation arises in markets such as eBay or eLance.com<sup>1</sup>, where the seller (respectively the service provider) can individually decide for each buyer how much effort to invest in satisfying the request. The seller can decide to deliver the product (respectively the service) as advertised, can decide to deliver a product of significantly lower quality, but can also decide not to deliver the product at all. This is a common problem in electronic markets, as buyers pay first and thus expose themselves to cheating behavior. Dellarocas and Wood (2006) estimate that only 80% of the buyers on eBay are satisfied with the outcome of the transaction. Moreover, 46% of all consumer fraud-reports received by the Federal Trade Commission in 2005 were related to the Internet, with Internet Auctions (12%) as a top category<sup>2</sup>.

One cause behind the high fraud rates is that sellers are not motivated to fulfill their promises.

---

<sup>1</sup>eLance.com is a marketplace for professional services.

<sup>2</sup>The report is available at <http://www.consumer.gov/sentinel/pubs/Top10Fraud2005.pdf>

Providing high quality requires expensive effort, while breaking the contractual agreements usually remains unpunished: contract enforcing institutions are weak in online environments, and the expected revenue from trading again with the same buyer is too small to cover the present cost of effort. Sellers thus face a *moral hazard* situation where cheating (i.e., providing low quality) is the optimal strategy. Anticipating this behavior, rational buyer naturally chose not to transact at all.

Sanctioning reputation mechanisms can effectively restore trust in such markets by revealing information about the sellers' past behavior. The cheating of a seller is announced to the whole community of buyers, and reduces the future opportunities of that seller to engage in profitable trade. The momentary gain obtained from cheating is smaller than the future losses due to a *bad reputation*, which motivates the seller to deliver the promised quality. The market can thus function efficiently.

While similar principles have regulated impersonal trade activities ever since the Medieval period (Greif, 1989), we now have a unique opportunity to *engineer* reputation mechanisms by controlling (i) the kind of feedback requested from the participants, (ii) the definition of reputation, and (iii) the flow of information in the system (Dellarocas, 2005). On eBay, for example, feedback can be positive, negative or neutral, followed by a short textual comment. Reputation information is displayed as a summary of reports (received in the last month, received in the last 6 months, received in the last year, or ever received) followed by a list of textual comments. However, one can also imagine reputation mechanisms where feedback is only binary, or drawn from a finite set. Likewise, reputation information could be computed based on other statistics applied to the set of reports.

It therefore becomes important to understand the impact of these design decisions on the overall performance of the reputation mechanism. This chapter addresses a simplified setting with *pure* moral hazard, meaning that all sellers (or service providers) are assumed equally capable of providing high quality, given that they invest the necessary (costly) effort. This assumption does not exactly match the real markets, where, some sellers are natively more capable than others. For example, a more skilled seller *A* could deliver some quality that seller *B* cannot, or even more commonly, seller *A* can deliver some quality much cheaper than seller *B*. Unfortunately, the theoretical analysis of such environments that combine moral hazard and asymmetric information (i.e., to the buyer it is not known that seller *A* is more skilled than seller *B*) is very complex, and still open to further research. However, there are settings where the moral hazard problem is predominant; for example on eBay, all seller can be assumed equally capable of mailing the product to the winning bidder.

Dellarocas (2005) provides an extensive analysis of binary reputation mechanisms where sellers may choose between two effort levels (e.g., an eBay seller that chooses to deliver or not the item) and buyers privately perceive an imperfect binary signal about the action of the seller (e.g., the eBay buyer that receives or not the item). He describes several reputation mechanisms that maximize the overall welfare of the society and attain the theoretical efficiency bound. He also concludes that efficient reputation mechanisms may depend only on the last submitted feedback report, and that the granularity of feedback is important only up to the point it allows a more precise monitoring of the seller's action<sup>3</sup>.

This chapter generalizes the results of Dellarocas to settings where sellers may choose between an arbitrary number of effort levels, and feedback from buyers can have an arbitrary number of values. From a practical perspective this generalization is important as it allows a more precise modeling of real interactions. On eBay, for example, a binary model restricts the seller to a binary choice: ship or not the item. In reality, however, the seller may also decide regarding the accuracy of the item description, the promptness of the delivery, or the care exerted when packaging the item. Studies conducted by Pavlou and Dimoka (2006) show that eBay textual comments often refer to these other strategic decisions of the seller, and have an impact on the future trade of the seller. The design of better reputation mechanisms

---

<sup>3</sup>A larger feedback value set can be compacted into classes of positive signals and negative signals, where positive signals are more likely to be observed following high effort, and negative signals are more likely to be observed following low effort



must therefore make use of both extended strategy spaces, and finer-grained feedback values.

The analysis will be based on a simplified model of a market where the same seller transacts with a sequence of short run buyers. Prior to the interaction, buyers are assumed to query the reputation mechanism, and check the reputation of the seller. Depending on the current reputation, every buyer decides how much to pay for the good (by paying nothing the buyer refuses to trade), and the seller then decides how much effort to put into delivering the good. Following the transaction, the reputation mechanism records the feedback from the client, and updates the reputation of the seller accordingly.

Ideally, a practical reputation mechanism should have the following properties:

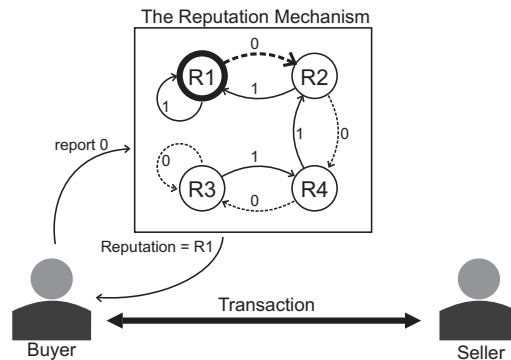
- reputation information is easy to understand, and has some intuitive representation;
- buyers are able to easily infer what is the best behavior they should adopt for a given reputation value of the seller. The reputation mechanism may eventually recommend the best course of action given the current reputation;
- clear rules specify how reputation changes when new feedback is submitted to the reputation mechanism;
- the mechanism encourages efficient trade.

Given the definition of reputation, the reputation updating rules and the behavior recommendations to the buyer, a rational seller will adopt a strategy that maximizes his life-time revenue. On the other hand, given the optimal strategy of the seller, rational buyers will respond optimally, and may disregard the recommendations of the reputation mechanism. From a game theoretic perspective, the market interactions can be modeled by a repeated game between a long-run player (the seller) and a sequence of short-run buyers that have access to a public *history* of play, as kept by the reputation mechanism. Therefore, the information displayed by the reputation mechanism, and the recommendations to the players must reflect an equilibrium strategy profile where rational buyers have the incentive to follow the recommendations, and rational sellers find it optimal to behave as the reputation mechanism predicted.

Previous results show that equilibria of repeated games can be represented by (finite) state automata (Osborne and Rubinstein, 1997; Mailath and Samuelson, 2006). We build on these results and construct reputation mechanisms that mirror the automaton of the underlying equilibrium. Thus, every state of the automaton is a possible value the reputation of the seller may take (see Figure 5.1). The outgoing transitions from a state are labeled with the possible feedback values a buyer can submit, and reflect the reputation updating rules defined by the mechanism. A behavior recommendation is associated to every state, and corresponds to the next round equilibrium strategy for the current reputation value. When the underlying automaton corresponds to an efficient equilibrium, the reputation mechanism has all of the ideal properties mentioned above.

Concretely, this chapter presents the following results. First, the two-state (i.e., *good* and *bad*) reputation mechanism of Dellarocas (2005) can be generalized to any number of effort levels and feedback values. Efficient mechanisms can be constructed using a randomization device that switches good reputation into bad with a probability that depends on the feedback submitted by the buyers. Bad reputation, however, may never become good again. In equilibrium, sellers with good reputation exert the required effort, while sellers with bad reputation do not exert any effort at all. For this reason, buyers interact only with good sellers.

Second, efficient reputation mechanisms can consider only the most recent  $N$  feedback reports, for any value of  $N$ . Every possible sequence of  $N$  reports becomes a distinct reputation value the seller may have; the price paid by the buyers depends on the current reputation of the seller, and corresponds to the expected utility obtained from the transaction with the seller. One characteristic of such mechanisms



**Figure 5.1:** The reputation mechanism defines four reputation values (R1,R2,R3 and R4) and two feedback values (0 or 1). The current reputation of the seller is R1; the reputation in the next round will be R2.

is that reputation information can be compressed to a statistic (i.e., a histogram) of the  $N$  reports. Nevertheless, the reputation mechanism must internally keep the entire sequence in order to correctly update the reputation.

Third, we show that there are reputation mechanisms where the equilibrium strategies are always pure strategies. The previous results rely on public randomization devices or mixed strategies. When practical reasons prevent the use of either, it is helpful to be able to design reputation mechanisms with pure equilibrium strategies.

Forth, we generalize the result of Dellarocas regarding the granularity of requested feedback. Given  $L$  effort levels, we show that a feedback set containing more than  $L+1$  values helps improve the efficiency of the reputation mechanism only to the extent it allows a better clustering of the  $L+1$  most significant feedback signals.

Last but not the least, we address the problem of honest reporting. As discussed in the previous chapters, empirical evidence suggests that benefits can be plundered by lying, and that some users do misreport feedback (White, 1999; Harmon, 2004). We describe an alternative to the side-payment schemes from Chapter 3 that explicitly reward truthful reports, by showing that honesty can emerge as a rational behavior when clients have a repeated presence in the market. To this end we describe a mechanism that supports an equilibrium where truthful feedback is obtained. Then we characterize the set of pareto-optimal equilibria of the mechanism, and derive an upper bound on the percentage of false reports that can be recorded by the mechanism. An important role in the existence of this bound is played by the fact that rational clients can establish a reputation for reporting honestly.

## 5.1 Related Work

The results of this chapter build on the extensive literature on repeated games with one long-run player and a sequence of short-run opponents. Fudenberg and Levine (1994) characterize the set of *perfect public equilibria* (PPE) of such games, and conclude that imperfect monitoring (i.e., buyers indirectly perceive the effort exerted by the seller through the quality of the product) impacts the efficiency of the market. Sellers are kept away from the first best payoff by an amount that is directly proportional to the chance of receiving unfair ratings. Moreover, Fudenberg and Levine show that for games describing

settings similar to ours, the set of sequential equilibrium payoffs is the same as the set of PPE payoffs.

The concept of perfect public equilibrium (Fudenberg et al., 1994) conveniently allows the analysis of infinitely repeated games using dynamic programming techniques (Abreu et al., 1990; Fudenberg et al., 1994). The game that starts in the second period is strategically equivalent to the original game, and therefore, strategies can be recursively expressed through the action of the first round, and a continuation strategy for the game starting in the second round. Equilibrium conditions require that any deviation from the strategy of the first round attracts lower continuation payoffs that make the deviation unprofitable (a direct application of the *one-shot deviation principle*).

Osborne and Rubinstein (1997) and Mailath and Samuelson (2006) express PPE equilibrium strategies by (finite) *state automata*. As agents play the game, the automaton *moves* from one state to the other depending on the history of previous rounds. Every state prescribes a strategy for the following round, and a transition table as a function of received feedback. We will be using this representation to describe efficient strategies for several reasons. First, every state of the automaton can be intuitively associated to a reputation value. Second, the transition function of the automaton nicely captures the rules for updating reputation information. Third, the one-round strategies associated to every state give an intuitive mapping between reputation values and equilibrium behavior.

Cooperative equilibria exist only when trading margins are high enough, and sellers discount future revenues sufficiently low (Dellarocas, 2005; Shapiro, 1983). The potential loss suffered by a seller who is not allowed to trade is thus high enough to offset the momentary gain from cheating. This restriction applies to our results as well.

## 5.2 The Setting

In our setting, a monopolist long-run seller (or provider) repeatedly sells the same good or service to a sequence of one-time buyers (or clients). Following the conventions, we will refer to a seller by *'he'* and to a buyer by *'she'*. In every round  $t$ , the  $t^{\text{th}}$  buyer first pays for the good, and then expects the seller to deliver the promised quality. Following payment, the seller must decide what effort to exert for producing the good. The set  $E = \{e_0, e_1, \dots, e_{L-1}\}$  of effort levels is known and does not change in time.

Exerting effort is expensive for the seller. Let  $c : E \rightarrow \mathbb{R}$  be the cost function of the seller, where  $c(e_i)$  is the cost of exerting effort  $e_i \in E$ . The cost function is assumed known. For simplicity, we order effort levels according to their cost. For any  $i < j$ ,  $e_i$  is *less effort* than  $e_j$ , and therefore,  $c(e_i) < c(e_j)$ . The continuous function obtained by the piecewise linearization of  $c$  is assumed convex, and without loss of generality, we take  $c(e_0) = 0$ ;  $e_0$  corresponds to exerting *no effort*.

The buyers *privately* perceive the quality of the delivered good as a signal taking values from the finite set  $Q = \{q_0, q_1, \dots, q_{M-1}\}$ . The perception of the buyers is influenced by the effort exerted by the seller, but also by uncontrollable environmental factors. We model the perception of buyers by an independent random signal, with an identical distribution that depends on the effort level exerted by the seller. In every round  $t$ , the signal  $q^{(t)}$  observed by the  $t^{\text{th}}$  buyer follows the distribution  $Pr[\cdot|e^{(t)}]$ , where  $Pr[q_j|e^{(t)}]$  is the probability of observing the signal  $q_j \in Q$ , given that the seller exerts effort  $e^{(t)} \in E$ . To avoid the confusion with the power operator, round indexes will always be placed as superscripts between round brackets. The distributions  $Pr[\cdot|e]$  are assumed known, and have full support: i.e.,  $Pr[q_j|e_i] \neq 0$  for all  $e_i \in E$  and all  $q_j \in Q$ .

The buyer from the  $t^{\text{th}}$  round is assumed to be the winner of a second price auction where every bidder bids her expected utility of the good. Let  $v_i : Q \rightarrow \mathbb{R}$  be the valuation function of bidder  $i$ , such

that  $v_i(q_k)$  is the value to bidder  $i$  of a good with perceived quality  $q_k \in Q$ . In round  $t$ , bidder  $i$  bids an amount equal to her expected utility of the good:

$$bid_i^{(t)} = \sum_{q_k \in Q} Pr[q_k | e^{(t)}] v_i(q_k);$$

and therefore, the price paid by the  $t^{th}$  buyer is:

$$p^{(t)} = \sum_{q_k \in Q} Pr[q_k | e^{(t)}] \tilde{v}(q_k);$$

where  $\tilde{v}(\cdot)$  is the valuation function of the second highest bidder. We assume that sufficiently many bidders bid in every round, and that their private valuation functions are drawn from a common distribution;  $\tilde{v}(q_k)$  can therefore be expected to take the same value in every round, and can be intuitively regarded as the market value of quality  $q_k$ .

Since the price in round  $t$  equals the expected utility of the second highest bidder, it can also be expressed as a function of exerted effort:  $p(e_i) = \sum_{q_k \in Q} Pr[q_k | e_i] \tilde{v}(q_k)$  is the expected price the seller would get for a good provisioned with effort  $e_i$ . We assume that  $p(e_0) = 0$ , as the market value of a good provided with no effort is 0.

The seller discounts future revenues with the factor  $\delta < 1$ , and is interested in maximizing the present-time value of the payoffs obtained over his entire life-time. If  $(p^{(t)})_{t=0, \dots, \infty}$  and  $(e^{(t)})_{t=0, \dots, \infty}$  are the sequences of prices, respectively effort levels, the lifetime revenue of the seller is:

$$\sum_{t=0}^{\infty} \delta^t (p^{(t)} - c(e^{(t)}));$$

which normalized to the one-round payoffs becomes:

$$U = (1 - \delta) \sum_{t=0}^{\infty} \delta^t (p^{(t)} - c(e^{(t)})); \quad (5.1)$$

A central reputation mechanism asks buyers to submit feedback about the quality delivered by the seller. In a first instance, we assume buyers report honestly the quality signal they observe. We do not exclude the possibility of reporting mistakes, however, these mistakes are assumed unintentional, and already factored in the distributions  $Pr[\cdot | \cdot]$ . Then, in Section 5.4 we will extend the analysis to address honest feedback reporting incentives.

The reputation mechanism is trusted by all participants; however, it cannot monitor the effort exerted by the seller, or the true quality signal observed by the buyer. After every round, the reputation mechanisms publishes the set of feedback signals reported so far, or any other statistics based on that set. A summary of the formal notation is given in the Appendix 5.A.

### 5.2.1 Example

As an example, consider an online CD retailer who auctions music CD's (of similar value) to interested buyers. All CDs are described as new, however, the seller has the choice to ship a cheaper, second-hand CD instead of a new one. The cost to the seller of a new CD is \$8, while a second hand CD costs only \$5. We assume that a second-hand CD is always available, and that shipping costs are \$1. Therefore, the three actions the seller can take at the end of ever auction are: (i) not to ship the CD, denoted as  $q_0$  and costing  $c(q_0) = 0$ , (ii) ship a second hand CD, denoted as  $q_1$  and costing  $c(q_1) = 5 + 1 = 6$ , and (iii) ship a new CD, denoted as  $q_2$ , and costing  $c(q_2) = 8 + 1 = 9$ .

Every buyer cares about getting a CD that plays well in her CD player; the market value of such a CD is \$13. CDs, however, may have scratches or other defects that may cause jitters with some devices. As these problems are clearly annoying for somebody listening to their favorite song, a CD with jitters is worth only \$2 to the owner. Naturally, the fact of not receiving any CD is worth \$0. Second hand CDs are more likely to cause jitters than new ones. We assume that a second hand CD plays well in 60% of the devices, while a new CD plays well in 95% of the devices. Moreover, the probability that a package gets lost during transportation is 1% (CDs lost during transportation are never recuperated, nor replaced by the seller).

The quality signal observed by the buyers has three possible values: (i)  $q_0$ , i.e., the CD was not received, (ii)  $q_1$ , i.e., the CD jitters, or (iii)  $q_2$ , i.e., the CD plays well. The conditional probability distribution of the quality signal observed by the buyers is presented in Figure 5.2(a). Here, for example, the probability  $Pr^*[q_1|e_1] = 0.396$  was computed as follows: given that the seller shipped a second hand CD, the buyer receives a CD that jitters with probability  $99\% \cdot 40\%$  (the probability that the CD was not lost during transportation, times the probability that a second-hand CD provokes jitters).

Given the above probabilities, we can now compute the price expected by the seller as a function of exerted effort (Figure 5.2(c)):

- $p(e_0) = 0$  when the seller does not ship the CD;
- $p(e_1) = 0 \cdot Pr^*[q_0|e_1] + 2 \cdot Pr^*[q_1|e_1] + 13 \cdot Pr^*[q_2|e_1] = 8.514$  when the seller ships a second hand CD;
- $p(e_2) = 0 \cdot Pr^*[q_0|e_2] + 2 \cdot Pr^*[q_1|e_2] + 13 \cdot Pr^*[q_2|e_2] = 12.325$  when the seller ships a new CD;

Buyers report honestly the quality signal they observed, but can make involuntary reporting mistakes. With probability 1% they report the signal  $q_j$  instead of  $q_i$ , for all  $q_j \neq q_i$ . The probability distribution of feedback signals recorded by the reputation mechanism is presented in Figure 5.2(b). Here, for example, the probability  $Pr[q_1|e_1] = 0.3941$  is computed as  $98\% \cdot 0.396 + (1 - 0.396) \cdot 1\%$ : the probability that the buyer observes  $q_1$  (given  $e_1$ ) and doesn't make a reporting mistake, plus the probability that the buyer doesn't observe  $q_1$  but makes a reporting mistake. The discount factor of the seller is assumed  $\delta = 0.9$ .

### 5.2.2 Strategies and Equilibria

The interaction between the seller and the buyers can be modeled by an infinitely repeated game between the long-run seller and a sequence of one-round buyers that win the corresponding second-price auctions. The stage game is a sequential game  $G$ , where the bidders first decide what price to pay, and the seller second decides what effort to exert. A strategy of the bidder in the stage game is a positive real value  $p \in \mathbb{R}^+$  specifying the price to pay. Second price auctions are incentive compatible, therefore rational bidders bid exactly their expected utility of the good.

A pure strategy of the seller in the game  $G$  is a mapping  $e : \mathbb{R}^+ \rightarrow E$  specifying the effort exerted as a function of the price paid by the buyer. Let  $\Delta(E)$  denote the set of all probability distributions over the elements of  $E$ . A mixed strategy of the seller,  $\epsilon(p) \in \Delta(E)$ , specifies a probability distribution over the effort exerted by the seller. If  $\epsilon(p) = \sum_{e_i \in E} \alpha_i(p) \cdot e_i$ , the seller exerts effort  $e_i$  with probability  $\alpha_i(p)$ . The sum  $\sum_{e_i \in E} \alpha_i(p)$  equals 1.

The only equilibrium of the stage game is inefficient: the seller exerts no effort (i.e.,  $\epsilon = e_0$ ) and anticipating this, bidders do not bid anything for the good ( $p = p(e_0) = 0$ ). This result is trivial, since  $e_0$  is the dominant strategy for the seller, regardless of the price paid by the buyer.

$Pr^*[\cdot \cdot]$	$e_0$	$e_1$	$e_2$
$q_0$	1	0.01	0.01
$q_1$	0	0.396	0.0495
$q_2$	0	0.594	0.9405

(a) Observation of buyers - conditional distribution of signals.

$Pr[\cdot \cdot]$	$e_0$	$e_1$	$e_2$
$q_0$	0.98	0.0197	0.0197
$q_1$	0.01	0.3941	0.058
$q_2$	0.01	0.5862	0.9223

(b) Feedback reports - conditional distribution of signals.

Cost:		
$e_0$	$e_1$	$e_2$
0	6	9

Price:		
$e_0$	$e_1$	$e_2$
0	8.5140	12.3255

(c) Cost and expected price of effort.

$e_0$	- does not ship CD
$e_1$	- ships second-hand CD
$e_2$	- ships new CD
$q_0$	- does not receive CD
$q_1$	- receives jittering CD
$q_2$	- receives good CD

(d) Legend

**Figure 5.2:** Numerical example. Conditional probability distribution of signals given effort levels, the cost of effort, and the expected price depending on effort.

The repeated game may have other equilibria when the seller cares enough about future revenues, i.e., the discount factor,  $\delta$  is large enough (Fudenberg and Levine, 1994). Intuitively, the feedback transmitted through the reputation mechanism makes it rational for the seller to exert effort. When previous feedback is positive, future buyers trust the seller and bid more for the good. The *cooperative* equilibrium is thus sustained by the fact that the future revenues made accessible by positive feedback, offset the momentary gain the seller would obtain by exerting no effort.

Formally, the feedback left by previous buyers creates a public *history* of the seller's behavior. Let  $h^{(t)} \in (\times Q)^{t-1}$  be the sequence of signals reported by the buyers up to (excluding) round  $t$ . The initial history is the empty set (i.e.,  $h^{(0)} = \emptyset$ ), and the set of all possible histories up to round  $t$  is  $H^{(t)} = (\times Q)^{t-1}$ .

In the repeated game, a strategy of the seller is a sequence  $\sigma = (\epsilon^{(t)})_{t=0, \dots, \infty}$  of mappings  $\epsilon^{(t)} : H^{(t)} \rightarrow \Delta(E)$  specifying the (possibly mixed) effort level exerted in round  $t$ , as a function of the history of the previous rounds.  $\sigma$  is a *perfect public equilibrium* (PPE) if for any round  $t$  and any history  $h^{(t)}$ , the seller cannot improve his payoff by a one-shot deviation from  $\sigma$ : i.e., there is no round  $t$  and history  $h^{(t)}$  so that the seller prefers to exert effort  $e' \in E$  instead of the equilibrium effort,  $\epsilon^{(t)}$ .

Let  $\sigma|h^{(t)}$  be the truncation of the strategy  $\sigma$  to the infinitely repeated game that starts in period  $t$ , given the history  $h^{(t)}$ ; let  $h^{(t+1)} = h^{(t)}q^{(t)}$  be the history resulting from appending the feedback  $q^{(t)} \in Q$  from round  $t$  to the history  $h^{(t)}$ .

The revenue of the seller for following  $\sigma|h^{(t)}$ , as perceived just before round  $t$  is:

$$U(\sigma|h^{(t)}) = (1 - \delta) \sum_{t'=t}^{\infty} \delta^{t'-t} (p^{(t')} - c(\epsilon^{(t')}));$$

and can further be expressed as the sum between the revenue obtained in round  $t$ , and the revenue obtained from the round  $t + 1$  onward. However, the strategy in the game that starts from round  $t + 1$  depends on the history  $h^{(t+1)} = h^{(t)}q^{(t)}$ , and therefore, the continuation payoff for the rest of the game is a function of the feedback  $q^{(t)}$  resulting from round  $t$ :

$$\begin{aligned} U(\sigma|h^{(t)}) &= (1 - \delta) \left( p^{(t)} - c(\epsilon^{(t)}) + \sum_{t'=t+1}^{\infty} \delta^{t'-t} (p^{(t')} - c(\epsilon^{(t')})) \right) \\ &= (1 - \delta) (p^{(t)} - c(\epsilon^{(t)})) + \delta \sum_{q_k \in Q} Pr[q_k|\epsilon^{(t)}] U(\sigma|h^{(t)}q^k); \end{aligned} \quad (5.2)$$

Therefore, the PPE condition for the strategy of the seller becomes:

$$U(\sigma|h^{(t)}) \geq (1 - \delta) (p^{(t)} - c(e')) + \delta \sum_{q_k \in Q} Pr[q_k|e'] U(\sigma|h^{(t)}q^k); \quad (5.3)$$

for all rounds  $t$ , histories  $h^{(t)}$  and effort levels  $e' \in E$ ,  $e' \neq \epsilon^{(t)}$ .

An equivalent, but practically more convenient representation of a strategy is given by an *automaton* (Osborne and Rubinstein, 1997). By definition, a strategy  $\sigma$  defines the action to be taken by the seller after any possible history. Since the set of all possible histories of the repeated game,  $H = \bigcup_{t=0}^{\infty} H^{(t)}$ , is infinite, the standard representation may prove cumbersome.

The automaton representation is based on two important observations. The first is that any strategy can also be represented recursively by specifying the action of the first round, and the different continuation strategies for the infinitely repeated game that starts in the second round. The continuation strategies depend on the feedback from the first round:  $\sigma = (\epsilon^{(0)}, \sigma|q_0, \dots, \sigma|q_{M-1})$  where  $\epsilon^{(0)} \in \Delta(E)$  is played in the first round, and the continuation strategy is  $\sigma|q_j$  when the feedback of the first round is  $q_j$ . Therefore,  $\sigma$  can be visually described by an infinite tree of continuation strategies. Each path in this tree is a specific play that might occur in the infinitely repeated game.

The second observation is that there may be continuation strategies, following different histories, that are the same. When recursively expanding  $\sigma$ , equal continuation strategies can be grouped together, and the tree is *condensed* to a (possibly infinite) directed graph. Every node of the graph corresponds to a (distinct) continuation strategy, and has a number of outgoing edges equal to the number of distinct feedback values that can be received after the first round of that continuation strategy. This graph visually pictures the automaton describing  $\sigma$ .

Formally, we partition the set  $H$  of histories into classes of histories that induce identical continuations. The histories  $h$  and  $h'$  from  $H$  belong to the same class if and only if the continuation strategies  $\sigma|h$  and  $\sigma|h'$  are identical. Classes are denoted by  $z$ , and the resulting partition is  $Z$ . Any strategy  $\sigma$  can be represented by the automaton  $(Z, z_s, \epsilon, trn)$  where:

- $Z$  is the set of states of the automaton;
- $z_s \in Z$  is the starting state.  $z_s$  contains the initial history,  $h^{(0)} = \emptyset$ ;

- $\epsilon : Z \rightarrow \Delta(E)$  is the output function of the automaton that associates the one-round strategy  $\epsilon(z) \in \Delta(E)$  to every state  $z \in Z$ . To simplify the notation will also use  $\epsilon_i = \epsilon(z_i)$ ;
- The transition function  $trn : Z \times Q \rightarrow Z$  identifies the next state of the automaton, given the current state and the feedback received from the buyer.  $trn(z, q_k) = z'$  if and only if  $h^{(t)} \in z$  and  $h^{(t+1)} = (h^{(t)} q_k) \in z'$  (i.e., the automaton transits to the state containing the history  $h^{(t)}$  concatenated with  $q_k$ ).

In most cases the partition  $Z$  is finite, and the strategy is represented by a finite state automaton.

Verifying the equilibrium conditions on a strategy represented by an automaton is straight-forward. Every state of the automaton can be associated with a payoff given by the continuation strategy that ‘starts’ in that state. If  $U(z_i)$  is the payoff associated to the state  $z_i \in Z$ , we have the following system of equations defining the payoffs:

$$U(z_i) = (1 - \delta) \left( p(\epsilon_i) - c(\epsilon_i) \right) + \delta \sum_{q_k \in Q} Pr[q_k | \epsilon_i] U(trn(z_i, q_k)); \quad \forall z_i \in Z;$$

As a direct application of the one-shot deviation principle, the strategy  $\sigma$  is an equilibrium if and only if for all states  $z_i \in Z$ , the seller cannot profitably deviate to some other effort  $e'$ :

$$U(z_i) \geq (1 - \delta) \left( p(\epsilon_i) - c(e') \right) + \delta \sum_{q_k \in Q} Pr[q_k | e'] U(trn(z_i, q_k));$$

In the infinitely repeated game, the seller generally has many PPE strategies. Nevertheless, online markets have the advantage that they can be designed to encourage a certain equilibrium. By controlling what information is made public, the reputation mechanism can *recommend* a certain strategy. Players are free, of course, to choose another equilibrium; however, the coordination on a new strategy profile may be difficult, and therefore, unlikely.

From the perspective of the reputation mechanism designer, it is important to know which equilibrium strategy to select, and how to construct the reputation mechanism that *recommends* that strategy.

As the most important criterion, the reputation mechanism should select an efficient equilibrium strategy that maximizes the cumulated revenue of the market participants. Since buyers pay exactly the market price, their equilibrium expected revenue is constant. Social efficiency is therefore maximized by the equilibrium strategy that gives the seller the highest possible revenue.

Among the equilibria that are socially efficient, the designer should choose the one that generates the simplest possible reputation mechanism.

It is straightforward to see that the reputation mechanism designer can implement the equilibrium strategy  $\sigma$  by revealing to the buyers only the automaton describing  $\sigma$  (i.e., set of states, the response function and transition function) and the current state of the automaton. Intuitively, the states of the automaton are the *reputation values* the seller can have. The current state of the automaton is the current reputation of the seller, the transition function defines how reputation changes as a function of submitted feedback, and the response function of the automaton prescribes the equilibrium behavior of the seller (and therefore of the buyers), given the current reputation.

A second criterion is simplicity. The reputation mechanism should be simple and intuitive to understand. Acknowledging that the perception of simplicity also depends on subjective factors like user interface and documentation, we consider two objective measures of simplicity. The first is the size



(number or states) of the automaton representing the equilibrium strategy. Fewer states lead to fewer possible values of reputation, and therefore, to a short description of the mechanism. The second is the use of randomization devices. Probabilistic decisions are assumed more complex than deterministic ones since they require the existence of trusted randomization devices.

### 5.2.3 Efficient Equilibrium Strategies

Fudenberg and Levine (1994) offer a general methodology for computing the maximum PPE payoff obtained by the seller. Let  $\mathcal{U}$  be the set of PPE payoffs the seller may obtain, and let  $\bar{U}$  be the maximum value in this set. By definition, there must be an equilibrium strategy  $\bar{\sigma}$  that generates the payoff  $\bar{U}$ . As we have seen in the previous section,  $\bar{\sigma}$  can be described by a first round strategy  $\bar{\epsilon}$  and by  $k$  different continuation strategies  $\bar{\sigma}|q_k$ , one for each possible feedback generated after the first round. Since the continuation strategies are by definition PPE strategies, the payoffs  $U_k$  generated by  $\bar{\sigma}|q_k$  must also belong to the set  $\mathcal{U}$ .

As a consequence, there must be the one round strategy  $\bar{\epsilon} \in \Delta(E)$ , and the PPE continuation payoffs  $(U_k)_{k=0, \dots, M-1} \in \mathcal{U}$  that enforce  $\bar{U}$ , so that:

- $\bar{U}$  is the payoff obtained by playing  $\bar{\epsilon}$  in the first round, and a PPE strategy that gives  $U_k$  in the subsequent rounds, depending on the feedback  $q_k$  from the first round;

$$\bar{U} = (1 - \delta)(p(\bar{\epsilon}) - c(\bar{\epsilon})) + \delta \sum_{q_k \in Q} Pr[q_k|\bar{\epsilon}]U_k; \quad (5.4)$$

- there is no profitable deviation from  $\bar{\epsilon}$  in the first round;

$$\bar{U} \geq (1 - \delta)(p(\bar{\epsilon}) - c(e')) + \delta \sum_{q_k \in Q} Pr[q_k|e']U_k; \quad \forall e' \in E \quad (5.5)$$

- all continuation payoffs are smaller or equal to  $\bar{U}$ .

$$\bar{U} \geq U_k, \quad \forall k = 1, \dots, M; \quad (5.6)$$

If  $\bar{\epsilon}$  were given, the payoff  $\bar{U}$  would be the solution of the following linear optimization problem:

**LP 5.2.1** Maximize  $\bar{U}$  as a function of  $(U_k)$ , such that (5.4), (5.5) and (5.6) hold.

However, we know from Fudenberg and Levine (1994) that  $\bar{\epsilon}$  is a pure strategy (i.e.,  $\bar{\epsilon} \in E$ ), and therefore we can find  $\bar{U}$  by solving LP 5.2.1  $L$  times, once for each element of  $E$ , and taking the maximum solution. Since linear optimization problems can be solved efficiently, the algorithm for finding  $\bar{U}$  is polynomial in the size of the sets  $Q$  and  $E$ .

Let us now compute the maximum PPE payoff the seller can obtain in the example presented in Section 5.2.1. Assuming that  $\bar{\epsilon} = e_0$  (the seller does not ship the CD), the only solution to LP 5.2.1 is  $\bar{U} = 0 = U_0 = U_1 = U_2$ . Intuitively, the seller does not gain anything in the first round (because  $p(e_0) - c(e_0) = 0$ ), and therefore, the set of constraints expressed by (5.6) force all equilibrium payoffs to 0.<sup>4</sup>

If  $\bar{\epsilon} = e_1$ , the optimization problem has the following solution:  $\bar{U} = 2.39$  and  $U_0 = 1.7, U_1 = U_2 = \bar{U}$ . If the efficient strategy required the seller to ship a second-hand CD, the fact that the buyer received a

<sup>4</sup>Assume some  $U_k$  is greater than 0. The expected continuation payoff  $\sum_{q_k \in Q} Pr[q_k|e_0]U_k$  is both smaller than  $U_k$  and greater than  $\bar{U}$ . Therefore  $\bar{U} < U_k$  which violates the constraints

CD (good or bad) is good enough to entitle the seller to the maximum PPE continuation payoff. Only when buyers report not having received the CD (signal  $q_0$ ) the seller is punished by a lower continuation payoff. These continuation payoffs make action  $e_1$  optimal for the seller: Not shipping the CD would cause an expected loss equal to  $\delta(\bar{U} - U_0)(Pr[q_0|e_0] - Pr[q_0|e_1])/(1 - \delta) = 6$ , just enough to cover the cost of shipping a second-hand CD.

The efficient strategy is however  $\bar{e} = e_2$ . The optimization problem LP 5.2.1 solves for:  $\bar{U} = 2.61$ ,  $U_0 = 1.52$ ,  $U_1 = 1.62$  and  $U_2 = \bar{U}$ . The fact that the buyer reports not having received a CD is punished by the worst continuation payoff,  $U_0$ . The report about a jittering CD is slightly better, but still reduces the continuation payoff of the seller. Only the feedback  $q_2$  (the CD plays well) entitles the seller to an efficient continuation payoff. Note that the maximum PPE payoff is quite far from the first best equilibrium payoff of the seller. If the seller were able to credibly commit to always ship a new CD, his expected payoff would be  $p(e_2) - c(e_2) = 3.32$ . The efficiency loss is a direct consequence of the imprecise monitoring process: honest sellers that always ship a new CD sometimes get punished for losses occurred during transportation or for manufacturing defects. These punishments are nonetheless necessary to deter the seller from shipping a second-hand CD, or not shipping at all.

Several properties of this example are worth mentioning. First, the feedback value  $q_2$  entitles the seller to the maximum continuation payoff. Second, the future loss suffered by the seller from exerting less effort than  $e_2$  is sometimes just enough to cover the momentary reduction of cost. For example, by shipping a second hand CD the seller would save \$3, and would expect a future loss that is currently worth:

$$\frac{\delta \left( U_2 \left( Pr[q_2|e_2] - Pr[q_2|e_1] \right) + U_1 \left( Pr[q_1|e_2] - Pr[q_1|e_1] \right) + U_0 \left( Pr[q_0|e_2] - Pr[q_0|e_1] \right) \right)}{1 - \delta} = 3;$$

Therefore, the seller is actually indifferent between shipping a new or a second-hand CD, even though the equilibrium strategy requires him to ship the new CD. It turns out that both these properties can hold for our general setting.

First, we will consider the continuation payoffs  $(U_k)_{k=0,\dots,M-1}$  that enforce the efficient payoff  $\bar{U}$ . As proven by the proposition below, at least one of these payoffs is equal to  $\bar{U}$ ; the corresponding feedback value  $q_{\bar{k}} \in Q$  entitles the seller to continue receiving  $\bar{U}$  in the rest of the game.

**Proposition 5.2.1** *Given the maximum PPE payoff  $\bar{U}$  enforced by the one round strategy  $\bar{e}$  and the continuation PPE payoffs  $(U_k)_{k=0,\dots,M-1}$ , there is  $\bar{k}$  such that  $U_{\bar{k}} = \bar{U}$ .*

PROOF. see Appendix 5.B. ■

Second, we prove that the first round strategy  $\bar{e}$  that enforces  $\bar{U}$ , is not strictly better than all other strategies the seller may play. Given the continuation payoffs  $(U_k)$ , there is at least another effort level  $e^* \neq \bar{e}$  the seller could exert in the first round, and still expect the payoff  $\bar{U}$ .

**Proposition 5.2.2** *Given the maximum PPE payoff  $\bar{U}$  enforced by the one round strategy  $\bar{e}$  and the continuation PPE payoffs  $(U_k)_{k=0,\dots,M-1}$ , there is at least another effort level  $e^* \in E$ ,  $e^* \neq \bar{e}$  such that the seller is indifferent between exerting  $\bar{e}$  or  $e^*$  in the first round.*

PROOF. see Appendix 5.C. ■

One last remark addresses the conditions when the seller's maximum PPE payoff is greater than 0. Intuitively, the seller exerts effort in the first round only when future revenues made accessible by

positive feedback offset the cost of effort. Naturally, this can happen only when future transactions generate sufficient revenue, and the discount factor  $\delta$  is high enough. The exact dependence can be computed analytically for the binary case (Dellarocas, 2005), and determined numerically for the general case addressed in this chapter.

### 5.3 Designing Efficient Reputation Mechanisms

Efficient reputation mechanisms should allow the market participants to coordinate on a PPE strategy that gives the seller the maximum revenue,  $\bar{U}$  (the surplus of the buyers is constant). Therefore, the first problem of the mechanism designer is to (i) find an efficient PPE strategy, and (ii) to choose what information to display to the agents in order to coordinate their play on the desired equilibrium strategy.

One obvious way of constructing an efficient PPE strategy is to recursively specify a strategy for the next round, and the different continuation payoffs (one for every possible feedback received after the first round) given by the equilibrium that start after the next round.  $\bar{U}$  is thus obtained when the seller plays  $\bar{\epsilon}$  in the first round, and receives one of the continuation payoffs ( $U_k$ ) for the game that starts in the second round. Every one of the payoffs  $U_k$  can be further obtained when the seller plays  $\epsilon_k^{(1)} \in \Delta(E)$  in the second round, and receives one of the continuation payoffs ( $U_j$ ) $_{j=0, \dots, M-1}$  from the third round onward<sup>5</sup>. Following the same procedure, the efficient PPE strategy can be fully determined.

However, a reputation mechanism based on a recursive description of the equilibrium has several practical disadvantages. First, as the game is repeated infinitely, the equilibrium cannot be fully expanded. Any description of the equilibrium will inform the agents that after some finite number of rounds (whose strategy is specified), the continuation strategy will be an equilibrium giving some known amount. The agents will have to trust the reputation mechanism that it will correctly choose future continuation strategies that enforce the promised payoffs.

Second, the reputation mechanism must actively compute the equilibrium in every round. Similarly, an auditor of the reputation mechanism must constantly verify that the next round strategy and the continuation payoffs follow the equilibrium conditions.

Third, there is no intuitive connection between the equilibrium strategy and the reputation of the seller. Buyers are being told that previous feedback entitles the seller to a continuation payoff that is worth a particular value. It may well happen that a sequence of apparently bad feedback entitles the seller to a high continuation payoff<sup>6</sup>, and therefore the intuition that bad feedback should correspond to low payoffs is lost.

Finally, the space of continuation payoffs is continuous between 0 and  $\bar{U}$ . Every practical computation system will make rounding errors that could affect the equilibrium conditions.

A more practical approach is to describe efficient PPE strategies through a finite state automaton, as mentioned in Section 5.2.2. Take for example the automaton  $(Z, z_s, \epsilon, trn)$  where the  $Z$  is the set of states,  $z_s$  is the starting state,  $\epsilon(z)$  describes the one-round strategy for each state  $z \in Z$ , and  $trn$  is the transition function. There is an immediate mapping to a reputation mechanism that encourages the agents to play the strategy described by  $(Z, z_s, \epsilon, trn)$ :  $Z$  can be considered the set of reputation values (every state can be given an intuitive label corresponding to the value of the state),  $z_s$  is the initial

<sup>5</sup>The continuation payoffs ( $U_j$ ) may of course be different for every  $U_k$ ; moreover, we must have  $U_k = (1 - \delta)(p(\epsilon_k^{(1)}) - c(\epsilon_k^{(1)})) + \delta \sum_{q_j \in Q} Pr[q_j | \epsilon_k^{(1)}] U_j$ , and  $U_k \geq (1 - \delta)(p(\epsilon_k^{(1)}) - c(e')) + \delta \sum_{q_j \in Q} Pr[q_j | e'] U_j$  for all  $e' \in E$ , in order to make one shot deviation unprofitable.

<sup>6</sup>This may happen, for example, when the seller just comes out of a sequence of rounds where he was punished by low prices

reputation of the seller, the function  $\epsilon$  prescribes the next round action for every reputation value, and the function  $trn$  defines how the reputation gets updated depending on the feedback of the last round.

Reputation mechanisms described in this way do not require active computation in every round, can be easily verified, and have an intuitive description. In the rest of this section we will describe different automata that correspond to efficient strategies and have, in some sense, a convenient form.

We take as the main criterion for *convenience*, the simplicity of the corresponding reputation mechanism. First, we would like the automaton to be as small as possible (in terms of number of states). Second, we would like the reputation mechanism to rely as little as possible on random (respectively) mixed strategies. It turns out there is a clear tradeoff between the two criteria. A summary of results presented in this section is the following:

- efficient PPE strategies can be described by a 2-state probabilistic automaton. The two states of the automaton correspond to the extreme values of the PPE payoff set,  $\mathcal{U}$ : 0 corresponds to the *bad* state and  $\bar{U}$  corresponds to the *good* state. The automaton is probabilistic in the sense that the transition function randomly chooses the next state. The transition probabilities, however, are a function of the feedback submitted by the buyer. Probabilistic automata can be implemented only if the reputation mechanism can take random decisions (i.e., has a random device that is trusted by the participants)
- when the reputation mechanism cannot use a randomization device, efficient PPE strategies can be described by deterministic automata of size  $M^N$  for any integer  $N$  greater or equal to 1. Such automata require the sellers to use mixed one-round strategies.
- when the reputation mechanism cannot use a randomization device, and sellers can use only pure strategies, a reputation mechanism can be designed by solving an optimization problem. Depending on the size of the resulting automaton, the reputation mechanism may not be fully efficient.


### 5.3.1 Probabilistic Reputation Mechanisms

Given the efficient payoff  $\bar{U}$  enforced by the effort level  $\bar{\epsilon}$  and the continuation payoffs  $(U_k)$ ,  $k = 0, \dots, M-1$ , consider the following automaton  $\mathcal{A}^p = (Z, z_s, \epsilon, trn)$  where:

- the set of states contains 2 states:  $Z = \{z_G, z_B\}$ , where  $z_G$  is the *good* state and  $z_B$  is the *bad* state.
- the automaton starts in the good state:  $z_s = z_G$
- $\epsilon(z_G) = \bar{\epsilon}$  and  $\epsilon(z_B) = e_0$ ; the automaton prescribes the effort level  $\bar{\epsilon}$  in the good state, and no effort (i.e.,  $e_0$ ) in the bad state.
- $trn(z_B, \cdot) = z_B$ , so that once in the bad state, the automaton stays forever in the bad state regardless of the feedback provided by the buyers. The transition from the state  $z_G$  is probabilistic:

$$trn(z_G, q_k) = \begin{cases} z_G & \text{with probability } \frac{U_k}{\bar{U}} \\ z_B & \text{with probability } \frac{\bar{U} - U_k}{\bar{U}} \end{cases}$$

The reputation mechanism that implements  $\mathcal{A}^p$  corresponds to the *randomization device mechanism* described by Dellarocas (2005) for binary reputation mechanisms. The seller has either good or bad reputation. Only sellers with good reputation are allowed to trade in the market; sellers with bad



**Figure 5.3:** Two state probabilistic reputation mechanism.

reputation are isolated and excluded for ever from the market. A seller with good reputation is always expected to exert the effort  $\bar{e}$ . Depending on the feedback submitted by the client, the reputation mechanism changes the good reputation into bad reputation with a certain probability. The probability is computed such that a deviation from the effort  $\bar{e}$  cannot bring the seller higher revenues.

Before proving that the resulting reputation mechanism is efficient, let us draw in Figure 5.3 the automaton  $\mathcal{A}^p$  corresponding to the example in Section 5.2.1. As explained, from the state  $z_B$ , the only possible transition is to  $z_B$ . From the state  $z_G$ , the transitions are as follows:

- if the buyer submits the feedback  $q_2$ , the automaton stays in  $z_G$ ;
- if the buyer submits the feedback  $q_1$ , the automaton transits to  $z_B$  with probability 38%;
- if the buyer submits the feedback  $q_0$ , the automaton transits to  $z_B$  with probability 42%;

**Proposition 5.3.1**  $\mathcal{A}^p$  defines an efficient PPE strategy.

PROOF. First, let us compute the payoff expected by the seller associated to each state  $z \in Z$  of  $\mathcal{A}^p$ . Let  $U_G$  and  $U_B$  be the payoffs associated to the states  $z_G$  and  $z_B$  respectively. When in  $z_B$  the seller obtains  $p(e_0) - c(e_0) = 0$  in the first round, and stays in  $z_B$  for the game that starts in the second round. Therefore:

$$U_B = (1 - \delta)(p(e_0) - c(e_0)) + \delta U_B = 0;$$

When in  $z_G$ , the seller obtains  $p(\bar{e}) - c(\bar{e})$  in the first round, and continues in the state  $z_G$  with probability  $\frac{U_k}{\bar{U}}$ . Therefore:

$$U_G = (1 - \delta)(p(\bar{e}) - c(\bar{e})) + \delta \sum_{q_k \in Q} Pr[q_k | \bar{e}] \left( \frac{U_k}{\bar{U}} U_G + \frac{\bar{U} - U_k}{\bar{U}} U_B \right) = \bar{U};$$

Clearly, the strategy described by  $\mathcal{A}^p$  gives the seller the maximum PPE payoff,  $\bar{U}$ .

Second, let us verify that the strategy defined by  $\mathcal{A}^p$  is an equilibrium. When in the state  $z_B$ , any deviation to more effort brings supplementary costs, and no benefits. When in state  $z_G$ , assume there is a profitable deviation to the effort level  $e' \neq \bar{e}$ . This happens only if:

$$U_G < (1 - \delta)(p(\bar{e}) - c(e')) + \delta \sum_{q_k \in Q} Pr[q_k | e'] U_k$$

This however contradicts the assumption that  $\bar{U}$  is a PPE payoff enforced by  $\bar{e}$  and  $(U_k)$ . ■

The reputation mechanism described by  $\mathcal{A}^p$  uses extreme punishments to enforce the cooperation of the seller. Any feedback value that is different from  $q_{\bar{k}}$  (the feedback that entitles the seller to a continuation payoff equal to  $\bar{U}$ , see Proposition 5.2.1) triggers the exclusion of the seller from the market with non-zero probability. Effort levels other than  $\bar{\epsilon}$  generate a higher risk of punishment, hence the equilibrium. Unfortunately, even good faith sellers that always exert the expected effort  $\bar{\epsilon}$  face a non-zero probability of being punished, and will eventually be excluded from the market. In practice, excluded buyers should be given the option to buy back their way into the market, and start fresh with good reputation. The entrance fee in this case should be equal to  $U_G$ .

A reputation mechanism based on  $\mathcal{A}^p$  can always be constructed given  $\bar{U}$ ,  $\bar{\epsilon}$  and the continuation payoffs ( $U_k$ ). Therefore, the condition mentioned in Section 5.2.3 regarding the magnitude of future revenues is also sufficient for the existence of probabilistic reputation mechanisms that induce some level of cooperation in the market.

### 5.3.2 Deterministic Reputation Mechanisms using Mixed Strategies

The random decisions taken by the reputation mechanism, coupled with the severity of the punishments that are based on those random decisions may undermine the trustworthiness of the reputation mechanism. Sellers that unfairly get punished, or buyers that do not see the consequence of their negative feedback may raise legitimate questions regarding the honesty of the reputation mechanism. The dependance on a public randomization device whose output is uncontestable can be unpractical; this justifies the need to design deterministic reputation mechanisms where every feedback triggers a precise effect in the equilibrium strategy.

One immediate observation is that a deterministic automaton must have a number of different states at least equal to the number of different continuation payoffs that enforce the desired equilibrium outcome. For simplicity, we assume in this section that all  $M$  continuation payoffs are different. If this is not the case (details in Section 5.3.4) some states will be duplicated, and can be easily grouped together.

In this section we show that we can describe efficient PPE strategies through automata with exactly  $M$  states. Consequently, there are efficient reputation mechanisms with exactly  $M$  values of reputation. Moreover, we prove that for any integer  $N > 0$ , there is at least one automaton with  $M^N$  states that corresponds to an efficient PPE strategy. All proofs are constructive and give straightforward algorithms for designing the corresponding reputation mechanism.

Intuitively, such mechanisms define the reputation of the seller based on the last  $N$  submitted feedback reports. The continuation payoffs  $U_k$  allow a ranking of the feedback signals, such that the feedback corresponding to the highest continuation payoff is considered as the best, and the feedback corresponding to the lowest continuation payoff is regarded as the worst possible feedback.

Starting from the effort  $\bar{\epsilon}$  and continuation payoffs ( $U_k$ ) enforcing  $\bar{U}$ , we will first describe the deterministic automaton with  $M$  states. The automaton  $\mathcal{A}^d = (Z, z_s, \epsilon, trn)$  is defined as follows:

- the set of states  $Z = \{z_0, z_1, \dots, z_{M-1}\}$  contains  $M$  states, one corresponding to each distinct continuation payoff  $U_k$ ;
- the starting state  $z_s = z_{\bar{k}}$  corresponds to the maximum continuation payoff  $U_{\bar{k}} = \bar{U}$  (this state exists, as shown in Proposition 5.2.1)
- for each state  $z_i$ ,  $\epsilon_i = \epsilon(z_i)$  is a mixed strategy over a subset of effort levels  $\bar{E} \subseteq E$  defined as follows:

$$\bar{E} = \{e \in E \mid \bar{U} = (1 - \delta)(p(\bar{e}) - c(e)) + \delta \sum_{q_k \in Q} Pr[q_k|e]U_k\} \quad (5.7)$$

$\bar{E}$  contains  $\bar{e}$  and all other effort levels  $e$ , such that the seller is indifferent between exerting  $\bar{e}$  or  $e$  in the first round of the equilibrium that gives him the maximum payoff  $\bar{U}$ . As shown in Proposition 5.2.2,  $\bar{E}$  contains at least another element besides  $\bar{e}$ .

Moreover, if  $\epsilon_i = \sum_{e \in \bar{E}} \alpha_i(e) \cdot e$ , where  $\alpha_i(e)$  is the probability assigned by  $\epsilon_i$  to the effort level  $e$  from  $\bar{E}$ , we have:

$$U_i = (1 - \delta)(p(\epsilon_i) - c(\epsilon_i)) + \delta \sum_{q_k \in Q} Pr[q_k|\epsilon_i]U_k; \quad (5.8)$$

- the transition function is  $trn(\cdot, q_k) = z_k$ . i.e., regardless of the current state, the feedback  $q_k$  transits the automaton in the state  $z_k$ .

For the example in Section 5.2.1, the reputation mechanism corresponding to the automaton  $\mathcal{A}^d$  is graphically presented in Figure 5.4. The seller starts in the state  $z_2$ , and expects a payoff equal to  $\bar{U}$ . In the first round the seller ships a new CD, and depending on the feedback  $q_j$  of the buyer, his reputation at the beginning of the second round will be  $z_j$ . The equilibrium strategy of the seller for the next round depends on his current reputation, and on the set  $\bar{E}$ . For the numerical values we have chosen, the set  $\bar{E} = \{e_0, e_1, e_2\}$ , as shipping a second-hand CD, or not shipping at all, causes to the seller an expected future loss that is exactly equal to the cost saved by cheating. One example of an equilibrium strategy is the following:

- the seller ships a new CD if the feedback from the last transaction was  $q_2$ ;
- if the feedback from the last transaction was  $q_1$ , the seller randomizes between shipping a new CD (with probability 20%) and not shipping anything (with probability 80%);
- if the feedback from the last transaction was  $q_0$ , the seller randomizes between shipping a new CD (with probability 11%) and not shipping anything (with probability 89%);


**Proposition 5.3.2** *if  $\mathcal{A}^d$  exists, it defines an efficient PPE strategy.*

PROOF. By definition of  $\mathcal{A}^d$ , the continuation payoff associated to any state  $z_i$  is  $U_i$ . Naturally, the starting state  $z_s$  gives the seller the efficient payoff  $\bar{U}$ .

The remaining problem is to verify that the strategy described by  $\mathcal{A}^d$  is a PPE. For that we must prove that there is no state  $z_i$  where the seller can find a profitable deviation from  $\epsilon_i$ . We do that in two steps. First, we show that for all states  $z_i$ , the seller is indifferent between exerting any effort level  $e' \in \bar{E}$ . Second we show that for all states  $z_i$ , any deviation to the effort  $e^* \in E \setminus \bar{E}$  causes a loss to the seller.

For any  $e' \in \bar{E}$  we have:

$$\delta \sum_{q_k \in Q} Pr[q_k|e']U_k = \bar{U} - (1 - \delta)(p(\bar{e}) - c(e'));$$



**Figure 5.4:** A Reputation mechanism using only the last feedback.

Given the state  $z_i$ , the revenue obtained by the seller for playing  $e' \in \bar{E}$  instead of  $\epsilon_i$  is:

$$\begin{aligned}
 U'_i &= (1 - \delta)(p(\epsilon_i) - c(e')) + \delta \sum_{q_k \in Q} Pr[q_k|e']U_k \\
 &= (1 - \delta)(p(\epsilon_i) - c(e')) + \bar{U} - (1 - \delta)(p(\bar{e}) - c(e')) \\
 &= \bar{U} - (1 - \delta)(p(\bar{e}) - p(\epsilon_i));
 \end{aligned} \tag{5.9}$$

and does not depend on  $e'$ . Since  $\epsilon_i$  is a mixed strategy with support in  $\bar{E}$ ,  $U_i = U'_i$ , and no deviation to  $e' \in \bar{E}$  is profitable.

On the other hand, for any  $e^* \in E \setminus \bar{E}$  we know from (5.5) that:

$$\delta \sum_{q_k \in Q} Pr[q_k|e^*]U_k < \bar{U} - (1 - \delta)(p(\bar{e}) - c(e^*));$$

and as a consequence, the payoff obtained by the seller for deviating in state  $z_i$  from  $\epsilon_i$  to  $e^*$  is strictly lower than  $U_i$ :

$$\begin{aligned}
 V_i^* &= (1 - \delta)(p(\epsilon_i) - c(e^*)) + \delta \sum_{q_k \in Q} Pr[q_k|e^*]U_k \\
 &< (1 - \delta)(p(\epsilon_i) - c(e^*)) + \bar{U} - (1 - \delta)(p(\bar{e}) - c(e^*)) \\
 &< \bar{U} - (1 - \delta)(p(\bar{e}) - p(\epsilon_i)) = U_i;
 \end{aligned}$$

Therefore, the strategy defined by  $\mathcal{A}^d$  is an efficient PPE. ■

$\mathcal{A}^d$  exists only when we can find the mixed one-round strategy  $\epsilon_i$  that satisfies Equation (5.8) for each state  $z_i$ . As an immediate consequence of Equation (5.9), a necessary and sufficient existence condition of  $\mathcal{A}^d$  is that all continuation payoffs that enforce  $\bar{U}$  satisfy:

$$U_k \geq \bar{U} - (1 - \delta)(p(\bar{e}) - \min_{e \in \bar{E}} p(e)); \tag{5.10}$$

This supplementary condition introduces further constraints on the optimization problem that defines  $\bar{U}$ . There will be situations where a positive payoff can theoretically be guaranteed to the seller,



but not by a reputation mechanism characterized by  $\mathcal{A}^d$ . Unlike probabilistic reputation mechanisms characterized by  $\mathcal{A}^p$ , deterministic reputation mechanisms based on  $\mathcal{A}^d$  require stricter conditions on the magnitude of future revenues than the ones mentioned in Section 5.2.3.

A reputation mechanism that remembers the last  $n > 1$  feedback reports can be constructed using the same intuition. Depending on the feedback received in the preceding round, the price obtained by the seller in the immediately next round is different: the *worse* the last feedback, the lower the next round payment. When the reputation mechanism remembers the last  $N$  reports, the punishment for bad feedback can be spread over  $N$  rounds; however, the cumulated punishment for every feedback is the same as if the reputation mechanism remembered only the last report. Let the automaton  $\mathcal{A}^{d(N)} = (Z, z_s, \epsilon, trn)$  be defined as follows:

- the set  $Z = \{z_r | r = (r_1, r_2, \dots, r_N)\}$  contains  $M^N$  states, one corresponding to every sequence  $r$  of  $N$  feedback signals.  $r_k \in Q$  is the feedback report submitted  $k$  rounds ago.
- the starting state  $z_s = z_{\bar{r}}$  where  $\bar{r} = (q_{\bar{k}}, \dots, q_{\bar{k}})$ .  $q_{\bar{k}} \in Q$  is the perfect feedback corresponding to the maximum continuation payoff  $U_{\bar{k}} = \bar{U}$  (this state exists, as shown in Proposition 5.2.1)
- for each state  $z_r$ ,  $\epsilon_r = \epsilon(z_r)$  is a mixed strategy over a subset of effort levels  $\bar{E} \subseteq E$  defined in (5.7).  $\epsilon_r = \sum_{e \in \bar{E}} \alpha_r(e) \cdot e$ , where  $\alpha_r(e)$  is the probability assigned by the strategy  $\epsilon_r$  to the effort level  $e$  from  $\bar{E}$ . Moreover,

$$\alpha_r(e) = \frac{\sum_{j=1}^N \alpha_{r_j}(e)}{\sum_{k=0}^{N-1} \delta^k}; \quad (5.11)$$

for all  $e \in \bar{E} \setminus \{\bar{e}\}$ , and  $\alpha_r(\bar{e}) = 1 - \sum_{e \in \bar{E} \setminus \{\bar{e}\}} \alpha_r(e)$ .  $\alpha_{r_j}(e)$  is the probability of effort  $e$  assigned by the strategy  $\epsilon_{r_j}$ ,  $r_j \in Q$ , and is defined by Equation (5.8).

- the transition function is  $trn(\cdot, z_r) = z_{r \oplus q_k}$ , where  $r \oplus q_k$  is the sequence of  $N - 1$  most recent reports plus the feedback  $q_k$ : i.e., if  $r = (r_1, r_2, \dots, r_N)$ , then  $r \oplus q_k = (q_k, r_1, \dots, r_{N-1})$ .

**Proposition 5.3.3** *if  $\mathcal{A}^{d(N)}$  exists, it defines an efficient PPE strategy.*

PROOF. First, we will show that the value of each state  $z_r$  of the automaton  $\mathcal{A}^{d(N)}$  is:

$$U(z_r) = \bar{U} - (1 - \delta) \sum_{j=1}^N (p(\bar{e}) - p(\epsilon_{r_j})) \frac{\sum_{k=0}^{N-j} \delta^k}{\sum_{k=0}^{N-1} \delta^k}; \quad (5.12)$$

We must verify that:

$$U(z_r) = (1 - \delta)(p(\epsilon_r) - c(\epsilon_r)) + \delta \sum_{q_k \in Q} Pr[q_k | \epsilon_r] U(z_{r \oplus q_k});$$

for all states  $z_r \in Z$ . We denote as  $\phi = \sum_{k=0}^{N-1} \delta^k$ , and derive the following preliminary results:

$$\begin{aligned} p(\epsilon_r) &= \frac{\sum_{j=1}^N p(\epsilon_{r_j}) - (N - \phi)p(\bar{e})}{\phi}; \\ c(\epsilon_r) &= \frac{\sum_{j=1}^N c(\epsilon_{r_j}) - (N - \phi)c(\bar{e})}{\phi}; \\ Pr[\cdot | \epsilon_r] &= \frac{\sum_{j=1}^N Pr[\cdot | \epsilon_{r_j}] - (N - \phi)Pr[\cdot | \bar{e}]}{\phi}; \end{aligned}$$

Therefore:

$$\begin{aligned}
U' &= (1 - \delta)(p(\epsilon_r) - c(\epsilon_r)) + \delta \sum_{q_k \in Q} Pr[q_k | \epsilon_r] U(z_{r \oplus q_k}) \\
&= \frac{1 - \delta}{\phi} \left( \sum_{j=1}^N (p(\epsilon_{r_j}) - c(\epsilon_{r_j})) - (N - \phi)(p(\bar{\epsilon}) - c(\bar{\epsilon})) \right) \\
&\quad + \frac{\delta}{\phi} \sum_{q_k \in Q} \left( \sum_{j=1}^N Pr[q_k | \epsilon_{r_j}] - (N - \phi) Pr[q_k | \bar{\epsilon}] \right) \left( \bar{U} - (1 - \delta)(p(\bar{\epsilon}) - p(\epsilon_{q_k})) \right. \\
&\quad \left. - (1 - \delta) \sum_{j=1}^{N-1} (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \frac{\sum_{k=0}^{N-1-j} \delta^i}{\phi} \right)
\end{aligned}$$

However,

$$U_{r_j} = (1 - \delta)(p(\epsilon_{r_j}) - c(\epsilon_{r_j})) + \delta \sum_{q_k \in Q} Pr[q_k | \epsilon_{r_j}] U_k = \bar{U} - (1 - \delta)(p(\bar{\epsilon}) - p(\epsilon_{r_j}));$$

therefore:

$$\begin{aligned}
U' &= \frac{1 - \delta}{\phi} \left( \sum_{j=1}^N (p(\epsilon_{r_j}) - c(\epsilon_{r_j})) - (N - \phi)(p(\bar{\epsilon}) - c(\bar{\epsilon})) \right) \\
&\quad + \frac{\delta}{\phi} \sum_{q_k \in Q} \left( \sum_{j=1}^N Pr[q_k | \epsilon_{r_j}] - (N - \phi) Pr[q_k | \bar{\epsilon}] \right) U_k \\
&\quad - \frac{\delta(1 - \delta)}{\phi} \sum_{j=1}^{N-1} (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=0}^{N-1-j} \delta^i \\
&= \frac{1}{\phi} \sum_{j=1}^N (U_{r_j} - \bar{U}) + \bar{U} - \frac{\delta(1 - \delta)}{\phi} \sum_{j=1}^{N-1} (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=0}^{N-1-j} \delta^i \\
&= \bar{U} - \frac{(1 - \delta)}{\phi} \sum_{j=1}^N (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \left( 1 + \delta \sum_{i=0}^{N-1-j} \delta^i \right) \\
&= \bar{U} - \frac{(1 - \delta)}{\phi} \sum_{j=1}^N (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=0}^{N-j} \delta^i = U(z_r)
\end{aligned}$$

Next, let us show that the strategy defined by  $\mathcal{A}^{d(N)}$  is an equilibrium. Assume that in state  $z_r$ , the seller deviates to the strategy  $e' \in \bar{E}$ . The payoff he obtains is thus:

$$\begin{aligned}
U'(z_r) &= (1 - \delta)(p(\epsilon_r) - c(e')) + \delta \sum_{q_k \in Q} Pr[q_k | e'] U(z_{r \oplus q_k}) \\
&= \frac{1 - \delta}{\phi} \left( \sum_{j=1}^N p(\epsilon_{r_j}) - (N - \phi)p(\bar{\epsilon}) - \phi \cdot c(e') \right) + \delta \sum_{q_k \in Q} Pr[q_k | e'] \\
&\quad \left( U_k - (1 - \delta) \frac{\sum_{j=1}^{N-1} (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=1}^{N-j} \delta^i}{\phi} \right) \\
&= \frac{1 - \delta}{\phi} \left( \sum_{j=1}^N p(\epsilon_{r_j}) - (N - \phi)p(\bar{\epsilon}) - \phi \cdot c(e') \right) + \bar{U} - (1 - \delta)(p(\bar{\epsilon}) - c(e')) \\
&\quad - (1 - \delta) \frac{\sum_{j=1}^{N-1} (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=1}^{N-j} \delta^i}{\phi} \\
&= \bar{U} - \frac{(1 - \delta)}{\phi} \sum_{j=1}^N (p(\bar{\epsilon}) - p(\epsilon_{r_j})) \sum_{i=0}^{N-j} \delta^i = U(z_r)
\end{aligned}$$

On the other hand, all deviations to the effort  $e^* \in E \setminus \bar{E}$  will give the seller a payoff strictly smaller than  $U(z_r)$  since by definition,  $\delta \sum_{q_k \in Q} Pr[q_k|e']U_k < \bar{U} - (1 - \delta)(p(\bar{e}) - c(e'))$ . Hence  $\mathcal{A}^{d(N)}$  describes an efficient PPE strategy. ■

One property of  $\mathcal{A}^{d(N)}$  is that the single-round strategies prescribed by it do not depend on the sequence of feedback reports. The mixed strategies  $\epsilon_r$  are computed only from the distribution of feedback reports in the sequence  $r$ .

Automaton  $\mathcal{A}^{d(N)}$  exists only if we can find the strategies  $\epsilon_r$  that satisfy the conditions mentioned in the definition. One immediate consequence of Equation (5.11) is that  $\sum_{e \in \bar{E} \setminus \{\bar{e}\}} \alpha_k(e) \leq \frac{\sum_{k=0}^{N-1} \delta^k}{N} < 1$ . Therefore, the continuation payoffs  $U_k$  must satisfy stricter constraints:

$$\begin{aligned} U_k &= \bar{U} - (1 - \delta)(p(\bar{e}) - p(\epsilon_k)) \\ &\geq \bar{U} - (1 - \delta) \min_{e \in \bar{E}} (p(\bar{e}) - p(e)) \frac{\sum_{k=0}^{N-1} \delta^k}{N}; \end{aligned}$$

The greater the value of  $N$ , the harder are the constraints on the continuation payoffs that enforce  $\bar{U}$ . Therefore, reputation mechanisms that remember more reports from the past do not increase the efficiency of the mechanism, but limit the settings where efficient mechanisms can be constructed. This conclusion agrees with the results of Dellarocas (2005) for the binary case.

### 5.3.3 Deterministic Reputation Mechanisms with Pure Strategies

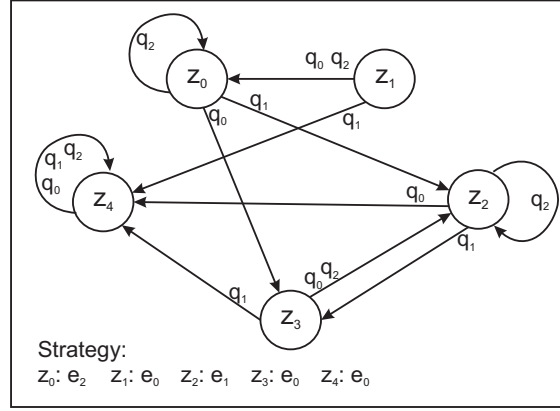
One assumption we made in Section 5.2 is that buyers bid their expected utility of the good offered by the seller. The equilibrium described by  $\mathcal{A}^{d(N)}$  prescribes a mixed strategy for the seller for most of the rounds: unless the feedback from the last  $N$  rounds is equal to  $\bar{r} = (q_{\bar{k}}, \dots, q_{\bar{k}})$ , the continuation payoff expected by the seller is smaller than  $\bar{U}$ , and requires the seller to randomize his action in the next round. Although the seller is indifferent between any effort level in the support of the mixed equilibrium strategy, the buyers must believe that the seller will choose randomly the exerted effort, following exactly the probabilities prescribed by the equilibrium.

Real buyers might have problems adopting such beliefs. Take for example the buyer before a round where the equilibrium prescribes the mixed strategy  $\alpha \bar{e} + (1 - \alpha)e_0$ , with  $\alpha$  close to 1. The seller is indifferent between exerting high effort,  $\bar{e}$ , or exerting no effort at all; in equilibrium, however, he is expected to exert high effort almost all the time. The buyer should therefore pay a price which is very close to  $p(\bar{e})$ , the expected utility of a good provisioned with high effort. Nevertheless, any aversion to risk will determine the buyer to pay significantly less than  $p(\bar{e})$ . As most people are risk-averse, a reputation mechanism that uses mixed strategies may have a serious impact on efficiency.

It therefore becomes important to find reputation mechanisms that support pure perfect public equilibrium strategies. Given a finite integer  $N$ , the problem we are addressing in this section is to find the finite state automaton  $\mathcal{A}^{pr(N)}$  with  $N$  states, describing a pure strategy PPE. Assuming that several automata exist, we would like to find the one that maximizes the revenue of the seller. The design parameter  $N$  can be intuitively seen as the maximum amount of memory available to the reputation mechanism.

Formally, given the set of  $N$  states  $Z = \{z_0, z_1, \dots, z_{N-1}\}$ , we must find a starting state  $z_s \in Z$ , a pure strategy response function  $\epsilon : Z \rightarrow E$ , and a deterministic transition function  $trn : Z \times Q \rightarrow Z$  such that the automaton  $\mathcal{A}^{pr(N)} = (Z, z_s, \epsilon, trn)$  corresponds to a PPE strategy and maximizes the revenue of the seller.

This design problem can be captured in an optimization problem, where the objective is to maximize



**Figure 5.5:** The reputation mechanism supporting a pure strategy equilibrium.

the continuation payoff of the starting state  $z_s$ , given that  $\mathcal{A}^{pr(N)}$  satisfies the equilibrium constraints: i.e., the seller cannot benefit by deviating from the strategy  $\epsilon(z_i) = \epsilon_i$  in any state  $z_i \in Z$ . Let  $U_i$  be the continuation payoff given by  $\mathcal{A}^{pr(N)}$  in state  $z_i$ . By symmetry we can assume that  $z_s = z_{N-1}$  and therefore the optimization problem becomes:

$$\begin{aligned}
 \min \quad & U_{N-1} \\
 \text{s.t.} \quad & U_i = (1 - \delta) \left( p(\epsilon_i) - c(\epsilon_i) \right) + \delta \sum_{q_k \in Q} Pr[q_k | \epsilon_i] U_{trn(z_i, q_k)}; \quad \forall z_i \in Z; \\
 & U_i \geq (1 - \delta) \left( p(\epsilon_i) - c(e) \right) + \delta \sum_{q_k \in Q} Pr[q_k | e] U_{trn(z_i, q_k)}; \quad \forall z_i, \forall e \neq \epsilon_i \in E; \\
 & U_i \leq U_{i+1} \quad \forall i = \{0, 1, \dots, N-2\};
 \end{aligned} \tag{5.13}$$

where  $U_{trn(z_i, q_k)}$  is the continuation payoff of the state  $trn(z_i, q_k)$ . The first constraint makes sure that the continuation payoffs are generated by the automaton  $\mathcal{A}^{pr(N)}$ , the second constraint expresses the equilibrium conditions, and the third constraint breaks the symmetry. The variables of the optimization problem are the strategies  $\epsilon_i$  and the transitions  $trn(z_i, q_k)$ .

For the example described in Section 5.2.1, Figure 5.5 graphically presents the best pure strategy reputation mechanism with  $N = 5$  states. The maximum payoff obtained by the seller is  $U_0 = 2.60$ , which is close to the efficient payoff  $\bar{U} = 2.61$ . The other continuation payoffs are:  $U_1 = 2.317$ ,  $U_2 = 1.608$ ,  $U_3 = 1.432$  and  $U_4 = 0$ .

This reputation mechanism has, however, several counterintuitive properties. Take for example the transitions from the state  $z_1$ . If the feedback of the buyer is  $q_0$  (the buyer did not receive the CD) or  $q_2$  (the buyer received a working CD) the automaton transits to the state  $z_0$ . If on the other hand the feedback is  $q_1$  (the CD is not working) the automaton transits to the state  $z_4$ . The state  $z_0$  corresponds to the maximum reputation value, while the state  $z_4$  corresponds to the lowest possible reputation value. The seller is therefore *rewarded* when the buyer does not receive the CD, and strongly *punished* when the buyer receives a defective CD. The same observations apply to the transitions from the state  $z_3$ .

These anomalies can be eliminated by imposing further constraints on the optimization problem (5.13). For example, one can request that the continuation payoffs following the signals  $q_0$ ,  $q_1$ , respec-

tively  $q_2$  should be increasing. As another example, one might want a reputation mechanism where in equilibrium, the sellers with better reputation exert higher effort (this again, is not the case for the mechanism presented in Figure 5.5: in state  $z_2$  the seller ships a second-hand CD, while in the state  $z_1$  the seller does not ship a CD). A number of other “practical” constraints can be added to make the reputation mechanism more intuitive, and optimal in the same time. The augmented optimization problem defines an automated algorithm for designing the best possible reputation mechanism given the desired constraints.

### 5.3.4 Feedback Granularity

We have assumed so far that the set of possible feedback values is fixed, and contains  $M$  elements. An interesting question is if the granularity of the requested feed influences the efficiency of the reputation mechanism. The answer is negative; we show that a mechanism where the buyer can submit  $L + 1$  different feedback values can be equally efficient as another mechanism where feedback can be infinitely diverse. However, having less than  $L + 1$  distinct feedback values can decrease the efficiency of the mechanism.

The main reason behind this fact is that among the  $M$  continuation payoffs that enforce the efficient payoff  $\bar{U}$ , at most  $L + 1$  take different values. If  $M > L + 1$  several continuations will give the seller either the maximum payoff,  $\bar{U}$ , or the minimum payoff 0.

**Proposition 5.3.4** *There are at most  $L + 1$  different continuation payoffs that enforce the socially efficient payoff,  $\bar{U}$ .*

PROOF. As seen in Proposition 5.2.2, the continuation payoffs  $(U_k)_{k=0,\dots,M-1}$  solve the linear problem (5.24). As a first step in this proof, note that the problem has a unique solution: for  $\delta \neq 0$ , the vector  $C$  defining the coefficients of the objective function is not parallel to any of the vectors defining the constraints.

The optimal solution therefore corresponds to one vertex of the polytope bounding the feasible region. This polytope is defined by  $2M + L - 1$  hyperplanes, where  $M + L - 1$  hyperplanes are defined by the equations  $A \cdot x = B$ , and the remaining  $M$  hyperplanes are defined by the equations  $U_k = 0$ .

Consequently, the optimal solution will satisfy  $M$  out of the  $2M + L - 1$  inequalities up to equality. When  $M > L$ , it follows that at least  $M - L + 1$  out of the  $2M$  bounding constraints  $0 \leq U_k \leq \bar{U}$  will be satisfied to equality. This means that  $M - L + 1$  continuation payoffs will be either equal to 0, or to  $\bar{U}$ . Hence, there are at most  $L + 1$  different continuation payoffs that enforce  $\bar{U}$ . ■

When different feedback signals generate the same continuation payoffs, the equilibrium strategy fails to make a distinction between them. Any of the first round feedback signals that trigger the maximum continuation payoff,  $\bar{U}$ , are equally *good*, while any of the signals that trigger the minimum continuation payoff equal to 0 are equally *bad*. Therefore, signals that trigger the same continuations can be grouped together, without loss of efficiency. A mechanism with more than  $L + 1$  different feedback signals is more efficient only up to the extent it allows a better grouping of signals that are very good or very bad. This result agrees with the conclusion of Dellarocas (2005) for the binary case.

The bound defined by Proposition 5.3.4 is also tight. There are settings where further compressing the feedback set reduces the maximum payoff the seller can get, and hence the efficiency of the market. To see this, let us extend the example from Section 5.2.1, and assume that buyers who experience jitters now examine the surface of the CD, and guess whether jitters are caused by manufacturing defects, or by scratches made by previous owners. A CD that plays well is not inspected. The buyer can now

$Pr^*[\cdot \cdot]$	$e_0$	$e_1$	$e_2$
$q_0$	1	0.01	0.01
$q_{1s}$	0	0.277	0.005
$q_{1n}$	0	0.119	0.0445
$q_2$	0	0.594	0.9405

(a) Observation of buyers - conditional distribution of signals.

$Pr[\cdot \cdot]$	$e_0$	$e_1$	$e_2$
$q_0$	0.97	0.0196	0.0196
$q_{1s}$	0.01	0.2761	0.0148
$q_{1n}$	0.01	0.1240	0.0528
$q_2$	0.01	0.5802	0.9129

(b) Feedback reports - conditional distribution of signals.

**Figure 5.6:** The buyer observes whether CD's have manufacturing defects or are scratched by ware.

report 4 quality signals, as a CD with jitters can be reported new ( $q_{1n}$ ), when the buyer believes the problems are coming from manufacturing defects, or second hand ( $q_{1s}$ ), when the buyer believes jitters are caused by scratches. On a second hand (jittering) CD, the buyer is likely to detect scratches in 70% of the cases. A new jittering CD, on the other hand, has observable scratches with much lower probability, i.e., 10%. The conditional probability distribution of the signals observed by the buyer depending on the action of the seller are displayed in Figure 5.6(a). The probability distribution of the signals recorded by the reputation mechanism (considering reporting mistakes) is given in Figure 5.6(b).

The maximum PPE payoff of the seller in this new context is  $\bar{U} = 2.97$  and is enforced by the first round strategy  $\bar{e} = e_2$  and the continuation payoffs:  $U_0 = 1.91$ ,  $U_{1s} = 1.69$ ,  $U_{1n} = U_2 = \bar{U}$ . Note that the maximum PPE payoff the seller can now obtain is higher than in the previous case (2.97 instead of 2.61). This, however, is not the consequences of a fourth quality signal – a reputation mechanism where the feedback signals  $q_{1n}$  and  $q_2$  are indistinguishable would be equally efficient. The increased efficiency is rather due to a change in the monitoring technology. By requiring additional feedback from buyers regarding the appearance of a defective CD, the reputation mechanism can more effectively punish the sellers for deviating from the equilibrium strategy, and hence increases the efficiency of the system.

Nevertheless, 4 quality signals are strictly necessary for smaller discount factors. When  $\delta = 0.79$ , the maximum PPE payoff is  $\bar{U} = 2.79$ , enforced by the effort  $\bar{e} = e_2$  and the continuation payoffs  $U_0 = 0.22$ ,  $U_{1s} = 0$ ,  $U_{1n} = 1.84$  and  $U_2 = \bar{U}$ . Any further reduction of the feedback set (i.e., by *compressing* two signals into one) strictly decreases the maximum PPE payoff the seller can obtain.

In general, when  $M > L + 1$ , the automata describing efficient equilibrium strategies will have duplicate states. All signals that trigger identical continuation payoffs can be grouped together into one signal, and therefore, the corresponding states in the automata can be collapsed into one state.

## 5.4 A Mechanism for Obtaining Reliable Feedback Reports

Signaling reputation mechanisms also face the important challenge of providing incentives to rational users to report honest feedback. Rational users can regard the private information they have observed as a valuable asset, not to be freely shared. Worse even, agents can have external incentives to misreport and thus manipulate the reputation information available to other agents. Without proper measures, the reputation mechanism will obtain unreliable information, biased by the strategic interests of the reporters.

There are well known solutions for providing honest reporting incentives for signaling reputation mechanisms (see Chapter 3). Since all clients interacting with a service receive the same quality (in a statistical sense), a client's private observation influences her belief regarding the experience of other clients. For example, if one client has a bad experience with a certain service, she is more likely to believe that other clients will also encounter problems when interacting with the same service. This correlation between the client's private belief and the feedback reported by other clients can be used to design feedback payments that make honesty a Nash equilibrium (see Section 3.3). When submitting feedback, clients get paid an amount that depends both on the value they reported and on the reports submitted by other clients. As long as others report truthfully, the expected payment of every client is maximized by the honest report – thus the equilibrium. Miller et al. (2005) and Jurca and Faltings (2006) show that incentive-compatible payments can be designed to offset both reporting costs and lying incentives.

For *sanctioning* reputation mechanisms the same payment schemes are not guaranteed to be incentive-compatible. Different clients may experience different service quality because the provider decided to exert different effort levels. The private beliefs of the reporter may no longer be correlated to the feedback of other clients, and therefore, the statistical properties exploited to obtain a truthful Nash equilibrium are no longer present.

As an alternative, we propose different incentives to motivate honest reporting based on the repeated presence of the client in the market. Game theoretic results (i.e., the *folk theorems*) show that repeated interactions support new equilibria where present deviations are made unattractive by future penalties. Even without a reputation mechanism, a client can guide her future play depending on the experience of previous interactions. As a first result of this chapter, we describe a mechanism that indeed supports a cooperative equilibrium where providers exert effort all the time. The reputation mechanism correctly records when the client received low quality.

There are certainly some applications where clients repeatedly interact with the same seller with a potential moral hazard problem. The barber shop mentioned in the beginning of this chapter is one example, as most people prefer going to the same barber (or hairdresser). Another example is a market of delivery services. Every package must be scheduled for timely delivery, and this involves a cost for the provider. Some of this cost may be saved by occasionally dropping a package, hence the moral hazard. Moreover, business clients typically rely on the same carrier to dispatch their documents or merchandise. As their own business depends on the quality and timeliness of the delivery, they do have the incentive to form a lasting relationship and get good service. Yet another example is that of a business person who repeatedly travels to an offshore client. The business person has a direct interest to repeatedly obtain good service from the hotel which is closest to the client's offices.

We assume that the quality observed by the clients is also influenced by environmental factors outside the control of, however observable by, the provider. Despite the barber's best effort, a sudden movement of the client can always generate an accidental cut that will make the client unhappy. Likewise, the delivery company may occasionally lose or damage some packages due to transportation accidents. Nevertheless, the delivery company (like the barber) eventually learns with certainty about any delays, damages or losses that entitle clients to complain about unsatisfactory service.

The mechanism we will develop in this section is quite simple. Before asking feedback from the client, the mechanism gives the provider the opportunity to acknowledge failure, and reimburse the client. Only when the provider claims good service does the reputation mechanism record the feedback of the client. Contradictory reports (the provider claims good service, but the client submits negative feedback) may only appear when one of the parties is lying, and therefore, both the client and the provider are sanctioned: the provider suffers a loss as a consequence of the negative report, while the client is given a small fine.

One equilibrium of the mechanism is when providers always do their best to deliver the promised quality, and truthfully acknowledge the failures caused by the environmental factors. Their “honest” behavior is motivated by the threat that any mistake will drive the unsatisfied client away from the market. When future transactions generate sufficient revenue, the provider does not afford to risk losing a client, hence the equilibrium.

Unfortunately, this socially desired equilibrium is not unique. Clients can occasionally accept bad service and keep returning to the same provider because they don’t have better alternatives. Moreover, since complaining for bad service is sanctioned by the reputation mechanism, clients might be reluctant to report negative feedback. Penalties for negative reports and the clients’ lack of choice drives the provider to occasionally cheat in order to increase his revenue.

As a second result, we characterize the set of pareto-optimal equilibria of our mechanism and prove that the amount of unreported cheating that can occur is limited by two factors. The first factor limits the amount of cheating in general, and is given by the quality of the alternatives available to the clients. Better alternatives increase the expectations of the clients, therefore the provider must cheat less in order to keep his customers.

The second factor limits the amount of unreported cheating, and represents the cost incurred by clients to establish a reputation for reporting the truth. By stubbornly exposing bad service when it happens, despite the fine imposed by the reputation mechanism, the client signals to the provider that she is committed to always report the truth. Such signals will eventually change the strategy of the provider to full cooperation, who will avoid the punishment for negative feedback. Having a reputation for reporting truthfully is of course, valuable to the client; therefore, a rational client accepts to lie (and give up the reputation) only when the cost of building a reputation for reporting honestly is greater than the occasional loss created by tolerated cheating. This cost is given by the ease with which the provider switches to cooperative play, and by the magnitude of the fine imposed for negative feedback.

Our mechanism is similar to the one used by Papaioannou and Stamoulis (2005a,b, 2006) for peer-to-peer applications. The authors require both transacting peers to submit feedback on their mutual performance. If reports disagree, both parties are punished, however the severity of their punishment depends on each agent’s *non-credibility* factor. This factor is maintained and evolved by the reputation mechanism. Papaioannou and Stamoulis experimentally investigate a variety of peer-to-peer situations, and show that a class of common lying strategies are successfully deterred by their scheme. The main difference from their work, however, is that we consider *all* possible pareto-optimal equilibrium strategies, and set upper bounds on the amount of untruthful information recorded by the reputation mechanism.

Concretely, we will proceed by describing in more detail our mechanism (Section 5.4.1), and by analyzing it from a game theoretic point of view (Section 5.4.2). We also establish the existence of the cooperative equilibrium, and derive an upper bound on the amount of cheating that can occur in any pareto-optimal equilibrium.

Then, we establish the cost of building a reputation for reporting honestly (Section 5.4.5), and hence compute an upper bound on the percentage of false reports recorded by the reputation mechanism in any equilibrium. Finally, Section 5.4.6 analyzes the impact of malicious buyers that explicitly try to destroy the reputation of the provider, and gives initial approximations on the worst case damage such buyers can cause to providers.



### 5.4.1 The CONFESS Mechanism

We limit the model described in Section 5.2 to a binary setting where the effort exerted by the provider can be *high* or *low*, and the quality observed by the clients is likewise, *high* or *low*. Formally, the set of effort levels is  $E_2 = \{e_0, e_1\}$  and the set of quality signals is  $Q_2 = \{q_0, q_1\}$ . Moreover, we assume that every client repeatedly interacts with the service provider, however, successive requests from the same client are always interleaved with enough requests generated by other clients. Transactions are assumed sequential, the provider does not have capacity constraints, and accepts all requests.

The price of service is  $p$  monetary units, but only high quality is valuable to the clients; the utility of high quality is  $u$ . Low quality has utility 0, and can be precisely distinguished from high quality. Before each round, the client can decide to request the service from the provider, or quit the market and resort to an outside provider that is completely trustworthy. The outside provider always delivers high quality service, but for a higher price  $p(1 + \rho)$ .

If the client decides to interact with the online provider, she issues a request to the provider, and pays for the service. The provider can now decide to exert low ( $e_0$ ) or high ( $e_1$ ) effort when treating the request. Low effort has a normalized cost of 0, but generates only low quality. High effort is expensive (normalized cost equals  $c(e_1) = c$ ) and generates high quality with probability  $\alpha < 1$ .  $\alpha$  is fixed, and depends on the environmental factors outside the control of the provider.  $\alpha p > c$ , so that it is individually rational for providers to exert effort.

After exerting effort, the provider can observe the quality of the resulting service. He can then decide to deliver the service as it is, or to acknowledge failure and roll back the transaction by fully reimbursing<sup>7</sup> the client. We assume perfect delivery channels, such that the client perceives exactly the same quality as the provider. After delivery, the client inspects the quality of service, and can accuse low quality by submitting a negative report to the reputation mechanism.

The reputation mechanism (RM) can oversee monetary transactions (i.e., payments made between clients and the provider) and can impose fines on all parties; however, the RM does not observe the effort level exerted by the provider, nor does it know the quality of the delivered service. The RM asks feedback from the client only if she chose to transact with the provider in the current round (i.e., paid the price of service to the provider) and the provider delivered the service (i.e., provider did not reimburse the client). When the client submits negative feedback, the RM punishes both the client and the provider: the client must pay a fine  $\varepsilon$ , and the provider accumulates a negative reputation report.

#### Examples

A delivery service for perishable goods (goods that lose value past a certain deadline) is one example where our mechanism can apply. Pizza, for example, must be delivered within 30 minutes, otherwise it gets cold and loses its taste. Hungry clients can order at home, or drive to a more expensive local restaurant, where they're sure to get a hot pizza. The price of a home delivered pizza is  $p = 1$ , while at the restaurant, the same pizza would cost  $p(1 + \rho) = 1.2$ . In both cases, the utility of a warm meal is  $u = 2$ .

The pizza delivery provider must exert costly effort to deliver orders within the deadline. A courier must be dispatched immediately (high effort), for an estimated cost of  $c = 0.8$ . While such action usually results in good service (the probability of a timely delivery is  $\alpha = 99\%$ ), traffic conditions and unexpected accidents (e.g., the address is not easily found) may still delay some deliveries past the

<sup>7</sup>In reality, the provider might also pay a penalty for rolling back the transaction. As long as this penalty is small, the qualitative results of this section remain valid.

deadline.

Once at the destination, the delivery person, as well as the client, know if the delivery was late or not. As it is common practice, the provider can acknowledge being late, and reimburse the client. Clients may provide feedback to a reputation mechanism, but their feedback counts only if they were not reimbursed. The client's fine for submitting a negative report can be set for example at  $\varepsilon = 0.01$ . The future loss to the provider caused by the negative report (and quantified through  $\bar{\varepsilon}$ ) depends on the reputation mechanism.

A simplified market of car garagists or plumbers could fit the same model. The provider is commissioned to repair a car (respectively the plumbing) and the quality of the work depends on the exerted effort. High effort is more costly but ensures a lasting result with high probability. Low effort is cheap, but the resulting fix is only temporary. In both cases, however, the warranty convention may specify the right of the client to ask for a reimbursement if problems reoccur within the warranty period. Reputation feedback may be submitted at the end of the warranty period, and is accepted only if reimbursements didn't occur.

An interesting emerging application comes with a new generation of web services that can optimally decide how to treat every request. For some service types, a high quality response requires the exclusive use of costly resources. For example, computation jobs require CPU time, storage requests need disk space, information requests need queries to databases. Sufficient resources, is a prerequisite, but not a guarantee for good service. Software and hardware failures may occur, however, these failures are properly signaled to the provider. Once monetary incentives become sufficiently important in such markets, intelligent providers will identify the moral hazard problem, and may act strategically as identified in our model.

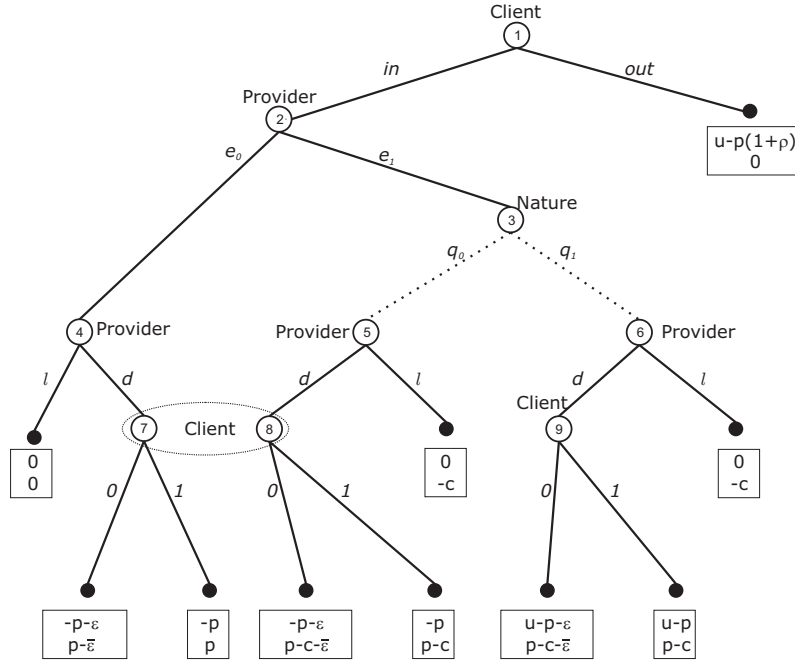
### 5.4.2 Behavior and Reporting Incentives

From game theoretic point of view, one interaction between the client and the provider can be modeled by the extensive-form game ( $G$ ) with imperfect public information, shown in Figure 5.7. The client moves first and decides (at node 1) whether to play *in* and interact with the provider, or to play *out* and resort to the trusted outside option.

Once the client plays *in*, the provider can chose at node 2 whether to exert high or low effort (i.e., plays  $e_1$  or  $e_0$  respectively). When the provider plays  $e_0$  the generated quality is low. When the provider plays  $e_1$ , nature chooses between high quality ( $q_1$ ) with probability  $\alpha$ , and low quality ( $q_0$ ) with probability  $1 - \alpha$ . The constant  $\alpha$  is assumed common knowledge in the market. Having seen the resulting quality, the provider delivers (i.e., plays  $d$ ) the service, or acknowledges low quality and rolls back the transaction (i.e., plays  $l$ ) by fully reimbursing the client. If the service is delivered, the client can report positive (1) or negative (0) feedback.

A pure strategy is a deterministic mapping describing an action for each of the player's information sets. The client has three information sets in the game  $G$ . The first information set is singleton and contains the node 1 at the beginning of game when the client must decide between playing *in* or *out*. The second information set contains the nodes 7 and 8 (the dotted oval in Figure 5.7) where the client must decide between reporting 0 or 1, given that she has received low quality,  $q_0$ . The third information set is singleton and contains the node 9 where the client must decide between reporting 0 or 1, given that she received high quality,  $q_1$ . The strategy  $in0^{q_0}1^{q_1}$ , for example, is the honest reporting strategy, specifying that the client enters the game, reports 0 when she receives low quality, and reports 1 when she receives high quality. The set of pure strategies of the client is:

$$A_C = \{out1^{q_0}1^{q_1}, out1^{q_0}0^{q_1}, out0^{q_0}1^{q_1}, out0^{q_0}0^{q_1}, in1^{q_0}1^{q_1}, in1^{q_0}0^{q_1}, in0^{q_0}1^{q_1}, in1^{q_0}1^{q_1}\};$$



**Figure 5.7:** The game representing one interaction. Empty circles represent decision nodes, edge labels represent actions, full circles represent terminal nodes and the dotted oval represents an information set. Payoffs are represented in rectangles, the top row describes the payoff of the client, the second row describes the payoff of the provider.

Similarly, the set of pure strategies of the provider is:

$$A_P = \{e_0l, e_0d, e_1l^{q_0}l^{q_1}, e_1l^{q_0}d^{q_1}, e_1d^{q_0}l^{q_1}, e_1d^{q_0}d^{q_1}\};$$

where  $e_1l^{q_0}d^{q_1}$ , for example, is the socially desired strategy: the provider exerts effort at node 2, acknowledges low quality at node 5, and delivers high quality at node 6. A pure strategy profile  $s$  is a pair  $(s_C, s_P)$  where  $s_C \in A_C$  and  $s_P \in A_P$ . If  $\Delta(A)$  denotes the set of probability distributions over the elements of  $A$ ,  $\sigma_C \in \Delta(A_C)$  and  $\sigma_P \in \Delta(A_P)$  are mixed strategies for the client, respectively the provider, and  $\sigma = (\sigma_C, \sigma_P)$  is a mixed strategy profile.

The payoffs to the players depend on the chosen strategy profile, and on the move of nature. Let  $g(\sigma) = (g_C(\sigma), g_P(\sigma))$  denote the pair of expected payoffs received by the client, respectively by the provider when playing strategy profile  $\sigma$ . The function  $g : \Delta(A_C) \times \Delta(A_P) \rightarrow \mathbb{R}^2$  is characterized in Table 5.1 and also describes the normal form transformation of  $G$ . Besides the corresponding payments made between the client and the provider, Table 5.1 also reflects the influence of the reputation mechanism, as further explained in Section 5.4.3. The four strategies of the client that involve playing *out* at node 1 generate the same outcomes, and therefore, have been collapsed for simplicity into a single row of Table 5.1.

		Client				<i>out</i>	
		<i>in1<sup>q0</sup>1<sup>q1</sup></i>	<i>in1<sup>q0</sup>0<sup>q1</sup></i>	<i>in0<sup>q0</sup>1<sup>q1</sup></i>	<i>in0<sup>q0</sup>0<sup>q1</sup></i>		
Provider	<i>e<sub>0</sub>l</i>	0	0	0	0	$u - p(1 + \rho)$	
	<i>e<sub>0</sub>d</i>	$-p$	$-p$	$-p - \varepsilon$	$-p - \varepsilon$	$u - p(1 + \rho)$	
	<i>e<sub>1</sub>l<sup>q0</sup>l<sup>q1</sup></i>	0	0	0	0	$u - p(1 + \rho)$	
	<i>e<sub>1</sub>l<sup>q0</sup>d<sup>q1</sup></i>	$\alpha(u - p)$ $\alpha p - c$	$\alpha(u - p - \varepsilon)$ $\alpha(p - \bar{\varepsilon}) - c$	$\alpha(u - p)$ $\alpha p - c$	$\alpha(u - p - \varepsilon)$ $\alpha(p - \bar{\varepsilon}) - c$	0	$u - p(1 + \rho)$
	<i>e<sub>1</sub>d<sup>q0</sup>l<sup>q1</sup></i>	$-(1 - \alpha)p$ $(1 - \alpha)p - c$	$-(1 - \alpha)p$ $(1 - \alpha)p - c$	$-(1 - \alpha)(p + \varepsilon)$ $(1 - \alpha)(p - \bar{\varepsilon}) - c$	$-(1 - \alpha)(p + \varepsilon)$ $(1 - \alpha)(p - \bar{\varepsilon}) - c$	0	$u - p(1 + \rho)$
	<i>e<sub>1</sub>d<sup>q0</sup>d<sup>q1</sup></i>	$\alpha u - p$ $p - c$	$\alpha(u - \varepsilon) - p$ $p - \alpha\bar{\varepsilon} - c$	$\alpha u - (1 - \alpha)\varepsilon - p$ $p - (1 - \alpha)\bar{\varepsilon} - c$	$\alpha u - \varepsilon - p$ $p - \bar{\varepsilon} - c$	0	$u - p(1 + \rho)$

Table 5.1: Normal transformation of the extensive form game,  $G$ 

### 5.4.3 Implementation in the Reputation Mechanism

For every interaction, the reputation mechanism records one of the three different signals it may receive: *positive* feedback when the client reports 1, *negative* feedback when the client reports 0, and *neutral* feedback when the provider rolls back the transaction and reimburses the client. In Figure 5.7 (and Table 5.1) positive and neutral feedback do not influence the payoff of the provider, while negative feedback imposes a punishment equivalent to  $\bar{\varepsilon}$ .

Two considerations made us choose this representation. First, we associate neutral and positive feedback with the same reward (0 in this case) because intuitively, the acknowledgement of failure may also be regarded as “honest” behavior on behalf of the provider. Failures occur despite best effort, and by acknowledging them, the provider shouldn’t suffer.

However, neutral feedback may also result because the provider did not exert effort. The lack of punishment for these instances contradicts the goal of the reputation mechanism to encourage exertion of effort. Fortunately, the action  $e_0l$  can be the result of rational behavior only in two circumstances, both excusable: one, when the provider defends himself against a malicious client that is expected to falsely report negative feedback (details in Section 5.4.6), and two, when the environmental noise is too big ( $\alpha$  is too small) to justify exertion of effort. Neutral feedback can be used to estimate the parameter  $\alpha$ , or to detect coalitions of malicious clients, and indirectly, may influence the revenue of the provider. However, for the simplified model presented above, positive and neutral feedback are considered the same in terms of generated payoffs.

The second argument relates to the role of the RM to constrain the revenue of the provider depending on the feedback of the client. There are several ways of doing that. In Section 5.3.1 we have seen that one way to punish the provider when the clients submit negative reports is by exclusion. After each negative report the reputation mechanism randomly decides to bans the provider from the market. The probability distribution for the random decision can be tuned such that the provider has the incentive to cooperate almost all the time, and the market stays efficient. The second method for punishing providers was presented in Section 5.3.2, where every negative report triggers the decrease of the price the next  $N$  clients will pay for the service. For lower values of  $N$  the price decrease is higher, nonetheless,  $N$  can take any value in an efficient market.

Both mechanisms work because the future losses offset the momentary gain the provider would have had by intentionally cheating on the client. Note that these penalties are given endogenously by

lost future opportunities, and require some minimum premiums for trusted providers. When margins are not high enough, providers do not care enough about future transactions, and will use the present opportunity of cheating.

Another option is to use exogenous penalties for cheating. For example, the provider may be required to buy a licence for operating in the market<sup>8</sup>. The licence is *partially destroyed* by every negative feedback. Totally destroyed licences must be restored through a new payment, and remaining parts can be sold if the provider quits the market. The price of the licence and the amount that is destroyed by a negative feedback can be scaled such that rational providers have the incentive to cooperate. Unlike the previous solutions, this mechanism does not require minimum transaction margins as punishments for negative feedback are directly subtracted from the upfront deposit.

One way or another, all reputation mechanisms foster cooperation because the provider associates value to client feedback. Let  $V(R^+)$  and  $V(R^-)$  be the value of a positive, respectively a negative report. In the game in Figure 5.7,  $V(R^+)$  is normalized to 0, and  $V(R^-)$  is  $\bar{\varepsilon}$ .  $V(R^-)$  can also be seen as the difference between the equilibrium continuation payoff expected by the provider following a negative, instead of a positive report (see Section 5.2.3). By using this notation, we abstract away the details of the reputation mechanism, and retain only the essential punishment associated with negative feedback. Any reputation mechanism can be plugged in our scheme, as long as the particular constraints (e.g., minimum margins for transactions) are satisfied.

One last aspect to be considered is the influence of the reputation mechanism on the future transactions of the client. If negative reports attract lower prices, rational long-run clients might be tempted to falsely report in order to purchase cheaper services in the future. Fortunately, some of the mechanisms designed for single-run clients, do not influence the reporting strategy of long-run clients. The reputation mechanism that only keeps the last  $N$  reports (Section 5.3.2) is one of them. A false negative report only influences the next  $N$  transactions of the provider; given that more than  $N$  other requests are interleaved between any two successive requests of the same client, a dishonest reporter cannot decrease the price for her future transactions.

The licence-based mechanism we have described above is another example. The price of service remains unchanged, therefore reporting incentives are unaffected. On the other hand, when negative feedback is punished by exclusion, clients may be more reluctant to report negatively, since they also lose a trading partner.

#### 5.4.4 Analysis of Equilibria

The one-time game presented in Figure 5.7 has only one subgame equilibrium where the client opts *out*. When asked to report feedback, the client always prefers to report 1 (reporting 0 attracts the penalty  $\varepsilon$ ). Knowing this, the best strategy for the provider is to exert low effort and deliver the service. Knowing the provider will play  $e_0d$ , it is strictly better for the client to play *out*.

The repeated game between the same client and provider may, however, have other equilibria. Before analyzing the repeated game, let us note that every interaction between a provider and a particular client can be strategically isolated and considered independently. As the provider accepts all clients and views them identically, he will maximize his expected revenue in each of the isolated repeated games.

From now on, we will only consider the repeated interaction between the provider and one client. This can be modeled by a  $T$ -fold repetition of the stage game  $G$ , denoted  $G^T$ , where  $T$  is finite or infinite. We will deal here with the infinite horizon case, however, the results obtained can also be applied with minor modifications to finitely repeated games where  $T$  is large enough.

<sup>8</sup>The reputation mechanism can buy and sell market licences

If  $\hat{\delta}$  is the per period discount factor reflecting the probability that the market ceases to exist after each round, (or the present value of future revenues), let us denote by  $\delta$  the expected discount factor in the game  $G^T$ . If our client interacts with the provider on the average every  $N$  rounds,  $\delta = \hat{\delta}^N$ .

The life-time expected payoff of the players is computed as:

$$\sum_{\tau=0}^T \delta^\tau g_i^\tau;$$

where  $i \in \{C, P\}$  is the client, respectively the provider,  $g_i^\tau$  is the expected payoff obtained by player  $i$  in the  $\tau^{th}$  interaction, and  $\delta^\tau$  is the discount applied to compute the present day value of  $g_i^\tau$ .

We will consider *normalized* life-time expected payoffs, so that payoffs in  $G$  and  $G^T$  can be expressed using the same measure:

$$V_i = (1 - \delta) \sum_{\tau=0}^T \delta^\tau g_i^\tau; \quad (5.14)$$

We define the *average continuation payoff* for player  $i$  from period  $t$  onward (and including period  $t$ ) as:

$$V_i^t = (1 - \delta) \sum_{\tau=t}^T \delta^{\tau-t} g_i^\tau; \quad (5.15)$$

The set of outcomes publicly perceived by both players after each round is:

$$Y = \{out, l, q_01, q_00, q_11, q_10\}$$


where:

- *out* is observed when the client opts *out*,
- *l* is observed when the provider acknowledges low quality and rolls back the transaction,
- $q_i j$  is observed when the provider delivers quality  $q_i \in \{q_0, q_1\}$  and the client reports  $j \in \{0, 1\}$ .

We denote by  $h^t$  a specific public history of the repeated game out of the set  $H^t = (\times Y)^t$  of all possible histories up to and including period  $t$ . In the repeated game, a public strategy  $\sigma_i$  of player  $i$  is a sequence of maps  $(\sigma_i^t)$ , where  $\sigma_i^t : H^{t-1} \rightarrow \Delta(A_i)$  prescribes the (mixed) strategy to be played in round  $t$ , after the public history  $h^{t-1} \in H^{t-1}$ . A *perfect public equilibrium* (PPE) is a profile of public strategies  $\sigma = (\sigma_C, \sigma_P)$  that, beginning at any time  $t$  and given any public history  $h^{t-1}$ , form a Nash equilibrium from that point on (Fudenberg et al., 1994).  $V_i^t(\sigma)$  is the continuation payoff to player  $i$  given by the strategy profile  $\sigma$ .

$G$  is a game with *product structure* since any public outcome can be expressed as a vector of two components  $(y_C, y_P)$  such that the distribution of  $y_i$  depends only on the actions of player  $i \in \{C, P\}$ , the client, respectively the provider. For such games, Fudenberg et al. (1994) establish a Folk Theorem proving that any feasible, individually rational payoff profile is achievable as a PPE of  $G^\infty$  when the discount factor is close enough to 1. The set of feasible, individually rational payoff profiles is characterized by:

- the minimax payoff to the client, obtained by the option *out*:  $\underline{V}_C = u - p(1 + \rho)$ ;



**Figure 5.8:** The pareto-optimal frontier of the set of feasible, individually rational payoff profiles of  $G$ .

- the minimax payoff to the provider, obtained when the provider plays  $e_0l$ :  $\underline{V}_P = 0$ ;
- the pareto optimal frontier (graphically presented in Figure 5.8) delimited by the payoffs given by (linear combination of) the strategy profiles  $(in1^{q_0}1^{q_1}, e_1 l^{q_0} d^{q_1})$ ,  $(in1^{q_0}1^{q_1}, e_1 d^{q_0} d^{q_1})$  and  $(in1^{q_0}1^{q_1}, e_0 d)$ .

and contains more than one point (i.e., the payoff when the client plays *out*) when  $\alpha(u-p) > u-p(1+\rho)$  and  $\alpha p - c > 0$ . Both conditions impose restrictions on the minimum margin generated by a transaction such that the interaction is profitable. The PPE payoff profile that gives the provider the maximum payoff is  $(\underline{V}_C, \overline{V}_P)$  where:

$$\overline{V}_P = \begin{cases} \alpha * u - c - u + p(1 + \rho) & \text{if } \rho \leq \frac{u(1-\alpha)}{p} \\ p + \frac{c(p\rho - u)}{\alpha u} & \text{if } \rho > \frac{u(1-\alpha)}{p} \end{cases}$$

and  $\underline{V}_C$  is defined above.

While completely characterizing the set of PPE payoffs for discount factors strictly smaller than 1 is outside the scope of this section, let us note the following results:

First, if the discount factor is high enough (but strictly less than 1) with respect to the profit margin obtained by the provider from one interaction, there is at least one PPE such that the reputation mechanism records only honest reports. Moreover, this equilibrium is pareto-optimal.

**Proposition 5.4.1** *When  $\delta > \frac{p}{p(1+\alpha)-c}$ , the strategy profile:*

- the provider always exerts high effort, and delivers only high quality; if the client deviates from the equilibrium, the provider switches to  $e_0d$  for the rest of the rounds;
- the client always reports 1 when asked to submit feedback; if the provider deviates, (i.e., she receives low quality), the client switches to out for the rest of the rounds.

is a pareto-optimal PPE.

PROOF. It is not profitable for the client to deviate from the equilibrium path. Reporting 0 attracts the penalty  $\varepsilon$  in the present round, and the termination of the interaction with the provider (the provider stops exerting effort from that round onwards).

The provider, on the other hand, can momentarily gain by deviating to  $e_1 d^{q_0} d^{q_1}$  or  $e_0 d$ . A deviation to  $e_1 d^{q_0} d^{q_1}$  gives an expected momentary gain of  $p(1 - \alpha)$  and an expected continuation loss of  $(1 - \alpha)(\alpha p - c)$ . A deviation to  $e_0 d$  brings an expected momentary gain equal to  $(1 - \alpha)p + c$  and an expected continuation loss of  $\alpha p - c$ . For the discount factor satisfying our hypothesis, both deviations are not profitable. The discount factor is low enough with respect to profit margins, such that the future revenues given by the equilibrium strategy offset the momentary gains obtained by deviating.

The equilibrium payoff profile is  $(V_C, V_P) = (\alpha(u - p), \alpha p - c)$ , which is pareto-optimal and socially efficient. ■

Second, we can prove that the client never reports negative feedback in any pareto-optimal PPE, regardless the value of the discount factor. The restriction to pareto-optimal is justifiable by practical reasons: assuming that the client and the provider can somehow negotiate the equilibrium they are going to play, it makes most sense to choose one of the pareto-optimal equilibria.

**Proposition 5.4.2** *The probability that the client reports negative feedback on the equilibrium path of any pareto-optimal PPE strategy is zero.*

PROOF. The full proof presented in Appendix 5.D follows the following steps. Step 1, all equilibrium payoffs can be expressed by adding the present round payoff to the discounted continuation payoff from the next round onward. Step 2, take the PPE payoff profile  $V = (V_C, V_P)$ , such that there is no other PPE payoff profile  $V' = (V'_C, V'_P)$  with  $V_C < V'_C$ . The client never reports negative feedback in the first round of the equilibrium that gives  $V$ . Step 3, the equilibrium continuation payoff after the first round also satisfies the conditions set for  $V$ . Hence, the probability that the client reports negative feedback on the equilibrium path that gives  $V$  is 0. Pareto-optimal PPE payoff profiles clearly satisfy the definition of  $V$ , hence the result of the proposition. ■

The third result we want to mention here, is that there is an upper bound on the percentage of false reports recorded by the reputation mechanism in any of the pareto-optimal equilibria.

**Proposition 5.4.3** *The upper bound on the percentage of false reports recorded by the reputation mechanism in any PPE equilibrium is:*

$$\gamma \leq \begin{cases} \frac{(1-\alpha)(p-u)+p\rho}{p} & \text{if } p\rho \leq u(1-\alpha); \\ \frac{p\rho}{u} & \text{if } p\rho > u(1-\alpha) \end{cases} \quad (5.16)$$

PROOF. The full proof presented in Appendix 5.E builds directly on the result of Proposition 5.4.2. Since clients never report negative feedback along pareto-optimal equilibria, the only false reports recorded by the reputation mechanism appear when the provider delivers low quality, and the client reports positive feedback. However, any PPE profile must give the client at least  $\underline{V}_C = u - p(1 + \rho)$ , otherwise the client is better off by resorting to the outside option. Every round in which the provider deliberately delivers low quality gives the client a payoff strictly smaller than  $u - p(1 + \rho)$ . An equilibrium payoff greater than  $\underline{V}_C$  is therefore possible only when the percentage of rounds where the provider delivers low



quality is bounded. The same bound limits the percentage of false reports recorded by the reputation mechanism. ■

For a more intuitive understanding of the results presented in this section, let us refer to the pizza delivery example detailed in Section 5.4.1. The price of a home delivered pizza is  $p = 1$ , while at the local restaurant the same pizza would cost  $p(1 + \rho) = 1.2$ . The utility of a warm pizza to the client is  $u = 2$ , the cost of delivery is  $c = 0.8$  and the probability that unexpected traffic conditions delay the delivery beyond the 30 minutes deadline (despite the best effort of the provider) is  $1 - \alpha = 0.01$ .

The client can secure a minimax payoff of  $V_C = u - p(1 + \rho) = 0.8$  by always going out to the restaurant. However, the socially desired equilibrium happens when the client orders pizza at home, and the pizza service exerts effort to deliver pizza in time: in this case the payoff of the client is  $V_C = \alpha(u - p) = 0.99$ , while the payoff of the provider is  $V_P = \alpha p - c = 0.19$ .

Proposition 5.4.1 gives a lower bound on the discount factor of the pizza delivery service such that repeated clients can expect the socially desired equilibrium. This bound is  $\delta = \frac{p}{p(1+\alpha)-c} = 0.84$ ; assuming that the daily discount factor of the pizza service is  $\hat{\delta} = 0.996$ , the same client must order pizza at home at least once every 6 weeks. The values of the discount factors can also be interpreted in terms of the minimum number of rounds the client (and the provider) will likely play the game. For example, the discount factor can be viewed as the probability that the client (respectively the provider) will “live” for another interaction in the market. It follows that the average lifetime of the provider is at least  $1/(1 - \hat{\delta}) = 250$  interactions (with all clients), while the average lifetime of the client is at least  $1/(1 - \delta) = 7$  interactions (with the same pizza delivery service). These are clearly realistic numbers.

Proposition 5.4.3 gives an upper bound on the percentage of false reports that our mechanism may record in equilibrium from the clients. As  $u(1 - \alpha) = 0.02 < 0.2 = p\rho$ , this limit is:

$$\gamma = \frac{p\rho}{u} = 0.1;$$

It follows that at least 90% of the reports recorded by our mechanism (in any equilibrium) are correct. The false reports (false positive reports) result from rare cases where the pizza delivery is intentionally delayed to save some cost but clients do not complain. The false report can be justified, for example, by the provider’s threat to refuse future orders from clients that complain. Given that late deliveries are still rare enough, clients are better off with the home delivery than with the restaurant, hence they accept the threat. As other options become available to the clients (e.g., competing delivery services) the bound  $\gamma$  will decrease.

Please note that the upper bound defined by Proposition 5.4.3 only depends on the outside alternative available to the provider, and is not influenced by the punishment  $\bar{\varepsilon}$  introduced by the reputation mechanism. This happens because the revenue of a client is independent of the interactions of other clients, and therefore, on the reputation information as reported by other clients. Equilibrium strategies are exclusively based on the direct experience of the client. In the following section, however, we will refine this bound by considering that clients can build a reputation for reporting honestly. There, the punishment  $\bar{\varepsilon}$  plays an important role.

### 5.4.5 Building a Reputation for Truthful Reporting

An immediate consequence of Propositions 5.4.2 and 5.4.3 is that the provider can extract all of the surplus created by the transactions by occasionally delivering low quality, and convincing the clients not to report negative feedback (providers can do so by promising sufficiently high continuation payoffs that prevent the client to resort to the outside provider). Assuming that the provider has more “power” in the market, he could influence the choice of the equilibrium strategy to one that gives him the most

revenue, and holds the clients close to the minimax payoff  $\underline{V}_C = u - p(1 + \rho)$  given by the outside option.<sup>9</sup>

However, a client who could commit to report honestly, (i.e., commit to play the strategy  $s_C^* = e_1 l^{q_0} d^{q_1}$ ) would benefit from cooperative trade. The provider's best response against  $s_C^*$  is to play  $e_1 l^{q_0} d^{q_1}$  repeatedly, which leads the game to the socially efficient outcome. Unfortunately the commitment to  $s_C^*$  is not credible in the complete information game, for the reasons explained in Section 5.4.4.

Following the results of Kreps et al. (1982), Fudenberg and Levine (1989) and Schmidt (1993) we know that such honest reporting commitments may become credible in a game with incomplete information. Suppose that the provider has incomplete information in  $G^\infty$ , and believes with non-negative probability that he is facing a committed client that always reports the truth. A rational client can then “fake” the committed client, and “build a reputation” for reporting honestly. When the reputation becomes credible, the provider will play  $e_1 l^{q_0} d^{q_1}$  (the best response against  $s_C^*$ ), which is better for the client than the payoff she would obtain if the provider knew she was the “rational” type.

As an effect of reputation building, the set of equilibrium points is reduced to a set where the payoff to the client is higher than the payoff obtained by a client committed to report honestly. As anticipated from Proposition 5.4.3, a smaller set of equilibrium points also reduces the bound of false reports recorded by the reputation mechanism. In certain cases, this bound can be reduced to almost zero.

Formally, incomplete information can be modeled by a perturbation of the complete information repeated game  $G^\infty$  such that in period 0 (before the first round of the game is played) the “type” of the client is drawn by nature out of a countable set  $\Theta$  according to the probability measure  $\mu$ . The client's payoff now additionally depends on her type. We say that in the perturbed game  $G^\infty(\mu)$  the provider has incomplete information because he is not sure about the true type of the client.

Two types from  $\Theta$  have particular importance:

- The “normal” type of the client, denoted by  $\theta_0$ , is the rational client who has the payoffs presented in Figure 5.7.
- The “commitment” type of the client, denoted by  $\theta^*$ , always prefers to play the commitment strategy  $s_C^*$ . From a rational perspective, the commitment type client obtains an arbitrarily high supplementary reward for reporting the truth. This external reward makes the strategy  $s_C^*$  the dominant strategy, and therefore, no commitment type client will play anything else than  $s_C^*$ .

In Theorem 5.4.1 we give an upper bound  $k_P$  on the number of times the provider delivers low quality in  $G^\infty(\mu)$ , given that he always observes the client reporting honestly.

The intuition behind this result is the following. The provider's best response to a honest reporter is  $e_1 l^{q_0} d^{q_1}$ : always exert high effort, and deliver only when the quality is high. This gives the commitment type client her maximum attainable payoff in  $G^\infty(\mu)$ , corresponding to the socially efficient outcome. The provider, however, would be better off by playing against the normal type client, against whom he can obtain an expected payoff greater than  $\alpha p - c$ .

The normal type client may be distinguished from a commitment type client only in the rounds when the provider delivers low quality: the commitment type always reports negative feedback, while

<sup>9</sup>All pareto-optimal PPE payoff profiles are also renegotiation-proof (Bernheim and Ray, 1989; Farrell and Maskin, 1989). This follows from the proof of Proposition 5.16: the continuation payoffs enforcing a pareto-optimal PPE payoff profile are also pareto-optimal. Therefore, clients falsely report positive feedback even under the more restrictive notion of negotiation-proof equilibrium.

the normal type might decide to report positive feedback in order to avoid the penalty  $\varepsilon$ . The provider can therefore decide to deliver low quality to the client in order to test her real type. The question is, how many times should the provider test the true type of the client.

Every failed test (i.e., the provider delivers low quality and the client reports negative feedback) generates a loss of  $-\bar{\varepsilon}$  to the provider, and slightly enforces the belief that the client reports honestly. Since the provider cannot wait infinitely for future payoffs, there must be a time when the provider will stop testing the type of the provider, and accepts to play the socially efficient strategy,  $e_1 l^{q_0} d^{q_1}$ .

The switch to the socially efficient strategy is not triggered by a revelation of the client's type. The provider believes that the client *behaves* as if she were a commitment type, not that the client *is* a commitment type. The client may very well be a normal type who chooses to mimic the commitment type, in the hope that she will obtain better service from the provider. However, further trying to determine the true type of the client is too costly for the provider. Therefore, the provider chooses to play  $e_1 l^{q_0} d^{q_1}$ , which is the best response to the commitment strategy  $s_C^*$ .

**Theorem 5.4.1** *If the provider has incomplete information in  $G^\infty$ , and assigns positive probability to the normal and commitment type of the client ( $\mu(\theta_0) > 0$ ,  $\mu_0^* = \mu(\theta^*) > 0$ ), there is a finite upper bound,  $k_P$ , on the number of times the provider delivers low quality in any equilibrium of  $G^\infty(\mu)$ . This upper bound is:*

$$k_P = \left\lceil \frac{\ln(\mu_0^*)}{\ln\left(\frac{\delta(\bar{V}_P - \alpha p + c) + (1-\delta)p}{\delta(\bar{V}_P - \alpha p + c) + (1-\delta)\bar{\varepsilon}}\right)} \right\rceil \quad (5.17)$$

PROOF. First, we use an important result obtained by Fudenberg and Levine (1989) about statistical inference (Lemma 1): If every previously delivered low quality service was sanctioned by a negative report, the provider must expect with increasing probability that his next low quality delivery will also be sanctioned by negative feedback. Technically, for any  $\pi < 1$ , the provider can deliver at most  $n(\pi)$  low quality services (sanctioned by negative feedback) before expecting that the  $n(\pi) + 1$  low quality delivery will also be sanctioned by negative feedback with probability greater than  $\pi$ . This number equals to:

$$n(\pi) = \left\lceil \frac{\ln \mu^*}{\ln \pi} \right\rceil;$$

As stated earlier, this lemma does not prove that the provider will become convinced that he is facing a commitment type client. It simply proves that after a finite number of rounds the provider becomes convinced that the client is playing as if she were a commitment type.

Second, if  $\pi > \frac{\delta \bar{V}_P}{\delta \bar{V}_P + (1-\delta)\bar{\varepsilon}}$  but is strictly smaller than 1, the rational provider does not deliver low quality (it is easy to verify that the maximum discounted future gain does not compensate for the risk of getting a negative feedback in the present round). By the previously mentioned lemma, it must be that in any equilibrium, the provider delivers low quality a finite number of times.

Third, let us analyze the round,  $\bar{t}$ , when the provider is about to deliver a low quality service (play  $d^{q_0}$ ) for the last time. If  $\pi$  is the belief of the provider that the client reports honestly in round  $\bar{t}$ , his expected payoff (just before deciding to deliver the low quality service) can be computed as follows:

- with probability  $\pi$  the client reports 0. Her reputation for reporting honestly becomes credible, so the provider plays  $e_1 l^{q_0} d^{q_1}$  in all subsequent rounds. The provider gains  $p - \bar{\varepsilon}$  in the current round, and expects  $\alpha p - c$  for the subsequent rounds;

- with probability  $1 - \pi$ , the client reports 1 and deviates from the commitment strategy, the provider knows he is facing a rational client, and can choose a continuation PPE strategy from the complete information game. He gains  $p$  in the current round, and expects at most  $\bar{V}_P$  in the subsequent rounds;

$$V_P \leq (1 - \delta)(p - \pi\bar{\varepsilon}) + \delta(\pi(\alpha p - c) + (1 - \pi)\bar{V}_P)$$

On the other hand, had the provider acknowledged the low quality and rolled back the transaction (i.e., play  $l^{q_0}$ ), his expected payoff would have been at least:

$$V'_P \geq (1 - \delta)0 + \delta(\alpha p - c)$$

Since the provider chooses nonetheless to play  $d^{q_0}$  it must be that  $V_P \geq V'_P$  which is equivalent to:

$$\pi \leq \bar{\pi} = \frac{\delta(\bar{V}_P - \alpha p + c) + (1 - \delta)p}{\delta(\bar{V}_P - \alpha p + c) + (1 - \delta)\bar{\varepsilon}} \quad (5.18)$$

Finally, by replacing Eq. (5.18) in the definition of  $n(\pi)$  we obtain the upper bound on the number of times the provider delivers low quality service to a client committed to report honestly. ■

The existence of  $k_P$  further reduces the possible equilibrium payoffs a client can get in  $G^\infty(\mu)$ . Consider a rational client who receives for the first time low quality. She has the following options:

- report negative feedback and attempt to build a reputation for reporting honestly. Her payoff for the current round is  $-p - \varepsilon$ . Moreover, her worst case expectation for the future is that the next  $k_P - 1$  rounds will also give her  $-p - \varepsilon$ , followed by the commitment payoff equal to  $\alpha(u - p)$ :

$$V_C|0 = (1 - \delta)(-p - \varepsilon) + \delta(1 - \delta^{k_P - 1})(-p - \varepsilon) + \delta^{k_P} \alpha(u - p); \quad (5.19)$$

- on the other hand, by reporting positive feedback she reveals to be a normal type, loses only  $p$  in the current round, and expects a continuation payoff equal to  $\hat{V}_C$  given by a PPE strategy profile of the complete information game  $G^\infty$ :

$$V_C|1 = (1 - \delta)(-p) + \delta\hat{V}_C; \quad (5.20)$$

The reputation mechanism records false reports only when clients do not have the incentive to build a reputation for reporting honestly, and  $V_C|1 > V_C|0$ ; this is true for:


$$\hat{V}_C > \delta^{k_P - 1} \alpha(u - p) - (1 - \delta^{k_P - 1})(p + \varepsilon) - \frac{1 - \delta}{\delta} \varepsilon;$$

Following the argument of Proposition 5.4.3 we can obtain a bound on the percentage of false reports recorded by the reputation mechanism in a pareto-optimal PPE that gives the client at least  $\hat{V}_C$ :

$$\hat{\gamma} = \begin{cases} \frac{\alpha(u-p) - \hat{V}_C}{p} & \text{if } \hat{V}_C \geq \alpha u - p; \\ \frac{u-p - \hat{V}_C}{u} & \text{if } \hat{V}_C < \alpha u - p \end{cases} \quad (5.21)$$

Of particular importance is the case when  $k_P = 1$ .  $\hat{V}_C$  and  $\hat{\gamma}$  become:

$$\hat{V}_C = \alpha(u - p) - \frac{1 - \delta}{\delta} \varepsilon; \quad \hat{\gamma} = \frac{(1 - \delta)\varepsilon}{\delta p}; \quad (5.22)$$



**Figure 5.9:** The upper bound  $k_P$  as a function of the prior belief  $\mu_0^*$ .


so the probability of recording a false report (after the first one) can be arbitrarily close to 0 as  $\varepsilon \rightarrow 0$ .

For the pizza delivery example introduced in Section 5.4.1, Figure 5.9 plots the bound,  $k_P$ , defined in Theorem 5.4.1, as a function of the prior belief ( $\mu_0^*$ ) of the provider that the client is an honest reporter. We have used a value of the discount factor equal to  $\delta = 0.95$ , such that on average, every client interacts  $1/(1 - \delta) = 20$  times with the same provider. The penalty for negative feedback was taken  $\bar{\varepsilon} = 2.5$ . When the provider believes that 20% of the clients always report honestly, he will deliver at most 3 times low quality. When the belief goes up to  $\mu_0^* = 40\%$  no rational provider will deliver low quality more than once.

In Figure 5.10 we plot the values of the bounds  $\gamma$  from Eq. (5.16) and  $\hat{\gamma}$  from Eq. (5.21) as a function of the prior belief  $\mu_0^*$ . The bounds simultaneously hold, therefore the maximum percentage of false reports recorded by the reputation mechanism is the minimum of the two. When  $\mu_0^*$  is less 0.25,  $k_P \geq 2$ ,  $\gamma \leq \hat{\gamma}$ , and the reputation effect does not significantly reduce the worst case percentage of false reports recorded by the mechanism. However, when  $\mu_0^* \in (0.25, 0.4)$  the reputation mechanism records (in the worst case) only half as many false reports, and as  $\mu_0^* > 0.4$ , the percentage of false reports drops to 0.005. This probability can be further decreased by decreasing the penalty  $\varepsilon$ . In the limit, as  $\varepsilon$  approaches 0, the reputation mechanism will register a false report with vanishing probability.

The result of Theorem 5.4.1 has to be interpreted as a worst case scenario. In real markets, providers that already have a small predisposition to cooperate will defect fewer times. Moreover, the mechanism is self enforcing, in the sense that the more clients act as commitment types, the higher will be the prior beliefs of the providers that new, unknown clients will report truthfully, and therefore the easier it will be for the new clients to act as truthful reporters.

As mentioned at the end of Section 5.4.4, the bound  $\hat{\gamma}$  strongly depends on the punishment  $\bar{\varepsilon}$  imposed by the reputation mechanism for a negative feedback. The higher  $\bar{\varepsilon}$ , the easier it is for clients to build a reputation, and therefore, the lower the amount of false information recorded by the reputation mechanism.



**Figure 5.10:** The maximum probability of recording a false report as a function of the prior belief  $\mu_0^*$ .

#### 5.4.6 The Threat of Malicious Clients

The mechanism described so far encourages service providers to do their best and deliver good service. The clients were assumed rational, or committed to report honestly, and in either case, they never report negative feedback unfairly. In this section, we investigate what happens when clients explicitly try to “hurt” the providers by submitting fake negative ratings to the reputation mechanism.

An immediate consequence of fake negative reports is that clients lose money. However, the costs  $\varepsilon$  of a negative report would probably be too small to deter clients with separate agendas from hurting the provider. Fortunately, the mechanism we propose naturally protects service providers from consistent attacks initiated by malicious clients.

Formally, a *malicious type* client,  $\theta_\beta \in \Theta$ , obtains a supplementary (external) payoff  $\beta$  for reporting negative feedback. Obviously,  $\beta$  has to be greater than the penalty  $\varepsilon$ , otherwise the results of Proposition 5.4.2 would apply. In the incomplete information game  $G^\infty(\mu)$ , the provider now assigns non-zero initial probability to the belief that the client is malicious.

When only the normal type,  $\theta_0$ , the honest reporter type  $\theta^*$  and the malicious type  $\theta_\beta$  have non-zero initial probability, the mechanism we describe is robust against unfair negative reports. The first false negative report exposes the client as being malicious, since neither the normal, nor the commitment type report 0 after receiving high quality. By Bayes’ Law, the provider’s updated belief following a false negative report must assign probability 1 to the malicious type. Although providers are not allowed to refuse service requests, they can protect themselves against malicious clients by playing  $e_0l$ : i.e., exert low effort and reimburse the client afterwards. The RM records neutral feedback in this case, and does not sanction the provider. Against  $e_0l$ , malicious clients are better off by quitting the market (opt out), thus stopping the attack. The RM records at most one false negative report for every malicious client, and assuming that identity changes are difficult, providers are not vulnerable to unfair punishments.

When other types (besides  $\theta_0, \theta^*$  and  $\theta_\beta$ ) have non-zero initial probability, malicious clients are harder to detect. They could masquerade client types that are normal, but accidentally misreport. It is not rational for the provider to immediately exclude (by playing  $e_0l$ ) normal clients that rarely misreport: the majority of the cooperative transactions rewarded by positive feedback still generate

positive payoffs. Let us now consider the client type  $\theta_0(\nu) \in \Theta$  that behaves exactly like the normal type, but misreports 0 instead of 1 independently with probability  $\nu$ . When interacting with the client type  $\theta_0(\nu)$ , the provider receives the maximum number of unfair negative reports when playing the efficient equilibrium: i.e.,  $e_1 l^{q_0} d^{q_1}$ . In this case, the provider's expected payoff is:

$$V_P = \alpha p - c - \nu \bar{\epsilon};$$

Since  $V_P$  has to be positive (the minimax payoff of the provider is 0, given by  $e_0 l$ ), it must be that  $\nu \leq \frac{\alpha p - c}{\bar{\epsilon}}$ .

The maximum value of  $\nu$  is also a good approximation for the maximum percentage of false negative reports the malicious type can submit to the reputation mechanism. Any significantly higher number of harmful reports exposes the malicious type and allows the provider to defend himself.

Note, however, that the malicious type can submit a fraction  $\nu$  of false reports only when the type  $\theta_0(\nu)$  has positive prior probability. When the provider does not believe that a normal client can make so many mistakes (even if the percentage of false reports is still low enough to generate positive revenues) he attributes the false reports to a malicious type, and disengages from cooperative behavior. Therefore, one method to reduce the impact of malicious clients is to make sure that normal clients make few or no mistakes. Technical means (for example by providing automated tools for formatting and submitting feedback), or improved user interfaces (that make it easier for human users to spot reporting mistakes) will greatly limit the percentage of mistakes made by normal clients, and therefore, also reduce the amount of harm done by malicious clients.

One concrete method for reducing mistakes is to solicit only negative feedback from the clients (the principle that no news is good news, also applied by Dellarocas (2005)). As reporting involves some conscious decision, mistakes will be less frequent. On the other hand, the reporting effort will add to the penalty for a negative report, and makes it harder for normal clients to establish a reputation for honest reporters. Alternative methods for reducing the harm done by malicious clients (like filtering mechanisms, etc.,) as well as tighter bounds on the percentage of false reports introduced by such clients will be further addressed in future work.

### 5.4.7 Remarks

Further benefits can be obtained if the clients' reputation for reporting honestly is shared within the market. The reports submitted by a client while interacting with other providers will change the initial beliefs of a new provider. As we have seen in Section 5.4.5, providers cheat less if they a priori expect with higher probability to encounter honest reporting clients. A client that has once built a reputation for truthfully reporting the provider's behavior will benefit from cooperative trade during her entire lifetime, without having to convince each provider separately. Therefore the upper bound on the loss a client has to withstand in order to convince a provider that she is a commitment type, becomes an upper bound on the total loss a client has to withstand during her entire lifetime in the market. How to effectively share the reputation of clients within the market remains an open issue.

Correlated with this idea is the observation that clients that use our mechanism are motivated to keep their identity. In generalized markets where agents are encouraged to play both roles (e.g. a peer-2-peer file sharing market where the fact that an agent acts only as "provider" can be interpreted as a strong indication of "double identity" with the intention of cheating) our mechanism also solves the problem signaled by Friedman and Resnick (2001) related to cheap online pseudonyms. The price to pay for the new identity is the loss due to building a reputation as truthful reporter when acting as a client.

Unlike incentive-compatible mechanism that pay reporters depending on the feedback provided by

peers, the mechanism described here is less vulnerable to collusion. The only reason individual clients would collude is to badmouth (i.e., artificially decrease the reputation of) a provider. However, as long as the punishment for negative feedback is not super-linear in the number of reports (this is usually the case), coordinating within a coalition brings no benefits for the colluders: individual actions are just as effective as the actions when part of a coalition. The collusion between the provider and client can only accelerate the synchronization of strategies on one of the PPE profiles (collusion on a non-PPE strategy profile is not stable), which is rather desirable. The only profitable collusion can happen when competitor providers incentivize normal clients to unfairly downrate their current provider. Colluding clients become *malicious* in this case, and the limits on the harm they can do are presented in Section 5.4.6.

The CONFESS mechanism is not a general solution for all online markets. In general retail e-commerce, clients don't usually interact with the same service provider more than once, which violates a crucial assumption behind our results. Nevertheless, we believe there are several scenarios of practical importance that do meet our requirements (e.g., interactions that are part of a supply chain). For these, our mechanism can be used in conjunction with other reputation mechanisms to guarantee reliable feedback and improve the overall efficiency of the market.

Although we present a setting where the service always costs the same amount, the results can be extended to scenarios where the provider may deliver different kinds of services, having different prices. As long as the provider believes that requests are randomly drawn from some distribution, the bounds presented above can be computed using the average values of  $u$ ,  $p$  and  $c$ . The constraint on the provider's belief is necessary in order to exclude some unlikely situations where the provider cheats on a one time high value transaction, knowing that the following interactions carry little revenue, and therefore, cannot impose effective punishments.

We also systematically overestimate the bounds on the worst case percentage of false reports recorded by the mechanism. The computation of tight bounds requires a precise quantitative description of the actual set of PPE payoffs the client and the provider can have in  $G^\infty$ . Fudenberg et al. (1994) and Abreu et al. (1990) pose the theoretical grounds for computing the set of PPE payoffs in an infinitely repeated game with discount factors strictly smaller than 1. However, efficient algorithms that allow us to find this set are still an open question. As research in this domain progresses, we expect to be able to significantly lower the upper bounds described in Sections 5.4.2 and 5.4.5.

One direction of future research is to study the behavior of the above mechanism when there is two-sided incomplete information: i.e. the client is also uncertain about the type of the provider. A provider type of particular importance is the "greedy" type who always likes to keep the client to a continuation payoff arbitrarily close to the minimal one. In this situation we expect to be able to find an upper bound  $k_C$  on the number of rounds in which a rational client would be willing to test the true type of the provider. The condition  $k_P < k_C$  describes the constraints on the parameters of the system for which the reputation effect will work in the favor of the client: i.e. the provider will give up first the "psychological" war and revert to a cooperative equilibrium.

The problem of involuntary reporting mistakes briefly mentioned in Section 5.4.6 needs further addressing. Besides false negative mistakes (reporting 0 instead of 1), normal clients can also make false positive mistakes (report 1 instead of the intended 0). In our present framework, one such mistake is enough to ruin the reputation of a normal type client to report honestly. This is one of the reasons why we chose a sequential model where the feedback of the client is not required if the provider acknowledges low quality. Once the reputation of the client becomes credible, the provider always rolls back the transactions that generate (accidentally or not) low quality, so the client is not required to continuously defend her reputation. Nevertheless, the consequences of reporting mistakes in the reputation building phase must be considered in more detail. Similarly, mistakes made by the provider, monitoring and communication errors will also influence the results presented here.



Last, but not the least, practical implementations of the mechanism we propose must address the problem of persistent online identities. One possible attack created by easy identity changes has been mentioned in Section 5.4.6: malicious buyers can continuously change identity in order to discredit the provider. In another attack, the provider can use fake identities to increase his revenue. When punishments for negative feedback are generated endogenously by decreased prices in a fixed number of future transactions (Dellarocas, 2005), the provider can adopt the following strategy: he cheats on all real customers, but generates a sufficient number of fake transactions in between two real transactions, such that the effect created by the real negative report disappears. An easy fix to this latter attack is to charge transaction or entrance fees. However, these measures also affect the overall efficiency of the market, and therefore, different applications will most likely need individual solutions.

## 5.5 Summary of Results

This chapter addresses the design of sanctioning reputation mechanisms for settings with moral hazard. They allow sellers to credibly commit that they will deliver the promised quality after having been paid by the buyers, by enforcing punishments on cheating sellers: negative feedback decreases the reputation, and hence the future opportunities and revenues available to the seller. The future loss is greater than the momentary gain obtained by cheating, and therefore cooperation becomes an optimal strategy for the sellers.

As opposed to signaling reputation mechanisms, sanctioning mechanisms can be designed in different ways. For example, the designer can choose the granularity of the solicited feedback, the reputation aggregation algorithm, and the amount of detail published to the other users. A deep understanding of these design parameters is essential for encouraging efficient trading in the market.

Dellarocas (2005) was the first to address the optimal design of binary reputation mechanisms where sellers can exert two effort levels, and buyers can observe two quality signals. In this chapter we extend the results of Dellarocas to general sanctioning mechanisms where sellers can exert any number of effort levels, and buyers can perceive any number of quality signals. First, we find that a reputation mechanism with two values of reputation (i.e., good or bad) can be efficient. Depending on the feedback provided by the last buyer, the reputation mechanism randomly decides to exclude the seller from the market. The probabilities governing the decision of the mechanism are scaled thus that the seller always finds it optimal to exert the equilibrium effort level.

Second, we show that a reputation mechanism that considers only the last  $N$  reports can be efficient, for any value of  $N$ . Such a mechanism has an equilibrium where the price paid by future buyers depends on the feedback submitted by the last  $N$  clients. Bad feedback decreases the prices the seller can obtain in the future  $N$  transactions, and encourages the seller to exert the expected (i.e., equilibrium) effort level.

Third, we present an algorithm for computing deterministic reputation mechanisms that only have pure equilibria. Such mechanisms can be desired, when, for example, buyers are risk averse, and mixed strategies cause significant efficiency losses to the sellers. Unfortunately, pure strategy sanctioning reputation mechanisms are hard to compute, and may generate payoffs that are strictly smaller than the optimal one (however, potentially greater than the payoffs that consider the risk aversion of buyers).

Forth, we investigate the efficiency of the reputation mechanism as a function of feedback granularity. Intuitively, as buyers are allowed to express finer grained feedback, the reputation mechanism should have more precise information which should lead to increased efficiency of the market. Surprisingly, however, a mechanism which accepts only  $L + 1$  different feedback (where  $L$  is the number of effort levels the seller can exert) can be just as efficient as a mechanism who allows infinitely fine-grained

feedback. More detailed feedback can increase the efficiency of the mechanism only to the extent it allows a better grouping of signals into a most  $L + 1$  classes.

Finally, we describe the CONFESS mechanism that can be coupled to any sanctioning reputation mechanism to encourage the submission of honest feedback. CONFESS works by comparing the feedback submitted by the buyer to the feedback implicitly submitted by the seller. After every transaction, the seller is allowed to acknowledge failures and reimburse the buyer, if the latter received bad quality. If, however, the seller does not reimburse the buyer, and the buyer submits negatived feedback, the reputation mechanism can conclude that (i) either the buyer is lying, or (ii) the seller cheated on the buyer. Therefore, it punishes both the seller and the buyer: the seller's reputation is decrease, and the buyer must pay a small fine. When future transactions generate sufficient profit, we prove that there is an equilibrium where the provider behaves as socially desired: he always exerts effort, and reimburses clients that occasionally receive bad service due to uncontrollable factors. Moreover, we analyze the set of pareto-optimal equilibria of the mechanism, and establish a limit on the maximum amount of false information recorded by the mechanism. The bound depends both on the external alternatives available to clients and on the ease with which they can commit to reporting the truth.

## Appendix 5.A Summary of Notation

Symbol	Meaning
$E = \{e_0, e_1, \dots, e_{L-1}\}$ $E_2 = \{e_0, e_1\}$	the set of $L$ effort levels; for binary settings, $e_0$ is no effort, $e_1$ is high effort;
$Q = \{q_0, q_1, \dots, q_{M-1}\}$ $Q_2 = \{q_0, q_1\} = \{0, 1\}$	the set of $M$ observable quality signals; for binary settings, the negative and the positive signals
$c(e_i)$ $v_i(q_k)$ $\tilde{v}(\cdot)$	the cost of effort $e_i$ ; the value of quality signal $q_k$ to buyer $i$ ; the valuation function of the second highest bidder;
$Pr[q_j e_i]$	the probability that a buyer perceives the quality signal $q_j$ when the seller exerts effort $e_i$ ;
$t$ $p^{(t)}, p(e_i)$	time index; the price paid by the $t^{\text{th}}$ buyer, equals the expected utility of the second highest bidder; the expected price obtained for exerting effort $e_i$ ;
$\delta$	the discount factor;
$U, U', \bar{U}, \tilde{U}, U_k$ $\mathcal{U}, \hat{\mathcal{U}}$	the expected (normalized) lifetime revenue of the seller; the set of PPE payoffs of the seller; $\hat{\mathcal{U}}$ is a discrete set;
$G$ $\epsilon \in \Delta(E)$ $h^{(t)} \in Q^{t-1}$ $H^{(t)} = Q^{t-1}, H = \bigcup_{t=0}^{\infty} H^{(t)}$  $\sigma_S = (\epsilon^{(t)})_{t=0, \dots, \infty}$ $\sigma_B = (p^{(t)})_{t=0, \dots, \infty}$ $\sigma = (\sigma_B, \sigma_S), \sigma h^{(t)}$	a game; a (mixed) strategy of the seller; the sequence of feedback reported by the buyers up to (excluding) round $t$ ; the set of all possible histories up to round $t$ ; $H$ is the set of all possible histories;  a strategy of the seller in the infinitely repeated game of $G$ ; a strategy of the buyers in the infinitely repeated game of $G$ ; a strategy profile; the truncation of $\sigma$ to the game that starts in period $t$ ;
$\mathcal{A} = (Z, z, \epsilon, trn)$	a finite state automaton;
$N$ $r = (r_1, \dots, r_N)$ $n_0(r)$ $q_j \oplus r$ $r \oplus q_j$ $Pr[r]$	number of reports considered by the reputation mechanism; a sequence of $N$ reports, $r_j \in Q_2$ ; the number of negative reports in the sequence $r$ ; the signal $q_j$ precedes the sequence $r$ ; the signal $q_j$ is appended to the sequence $r$ ; the probability of sequence $r$ ;
$\epsilon$ $\phi, \gamma$	some small positive amount; used to substitute complex expressions;

## Appendix 5.B Appendix: Proof of Proposition 5.2.1

Given the maximum PPE payoff  $\bar{U}$  enforced by the one round strategy  $\bar{e}$  and the continuation PPE payoffs  $(U_k)_{k=0, \dots, M-1}$ , there is  $\bar{k}$  such that  $U_{\bar{k}} = \bar{U}$ .

PROOF. Consider the constraints (5.5) of the linear optimization problem LP 5.2.1. These constraints make it rational for the seller to exert effort  $\bar{e}$  in the first round, rather than any other effort  $e' \neq \bar{e}$ .

Expanding  $\bar{U}$  according to (5.4) we get:

$$\delta \sum_{q_k \in Q} \left( Pr[q_k|\bar{\epsilon}] - Pr[q_k|e'] \right) U_k \geq (1 - \delta) \left( c(e') - c(\bar{\epsilon}) \right); \quad \forall e' \neq \bar{\epsilon} \quad (5.23)$$

Assume that all  $U_k$  are strictly smaller than  $\bar{U}$ , such that  $\bar{U} - U_k > \gamma > 0$ , for all  $k = 0, \dots, M - 1$ . The continuation payoffs  $U_k + \gamma$  are all smaller than  $\bar{U}$ , and therefore, they can be PPE payoffs obtained by the seller.

Since  $\sum_{q_k \in Q} Pr[q_k|\bar{\epsilon}] = \sum_{q_k \in Q} Pr[q_k|e'] = 1$ , the continuation payoffs  $U_k + \gamma$  also satisfy the inequalities (5.23). Moreover, the payoff enforced by  $\bar{\epsilon}$  and  $(U_k + \gamma)$  is  $U' = \bar{U} + \gamma$  and meets all necessary constraints to be a PPE payoff. This is however impossible, since  $\bar{U}$  is the maximum PPE payoff the seller can get. It follows that  $\gamma = 0$  and therefore, there must be  $\bar{k}$  such that  $U_{\bar{k}} = \bar{U}$ . ■

## Appendix 5.C Appendix: Proof of Proposition 5.2.2

*Given the maximum PPE payoff  $\bar{U}$  enforced by the one round strategy  $\bar{\epsilon}$  and the continuation PPE payoffs  $(U_k)_{k=0, \dots, M-1}$ , there is at least another effort level  $e^* \in E, e^* \neq \bar{\epsilon}$  such that the seller is indifferent between exerting  $\bar{\epsilon}$  or  $e^*$  in the first round.*

PROOF. Given  $\bar{\epsilon}$ , the continuation payoffs that enforce the maximum PPE payoff  $\bar{U}$  solve the linear problem (in standard matrix notation):

$$\begin{aligned} \max. \quad & C^T x \\ \text{s.t.} \quad & Ax \leq B, \quad x \geq 0 \end{aligned} \quad (5.24)$$

where  $x^T = (U_k)_{k=0, \dots, M-1}$  are the variables, the coefficients of the objective function are  $C^T = (\delta Pr[q_k|\bar{\epsilon}])_{k=0, \dots, M-1}$ , and the constraints are given by the  $(M+L-1) \times M$ , (respectively  $(M+L-1) \times 1$ ) matrices:

$$A = \begin{pmatrix} 1 - \delta Pr[q_0|\bar{\epsilon}] & -\delta Pr[q_1|\bar{\epsilon}] & \dots & -\delta Pr[q_{M-1}|\bar{\epsilon}] \\ -\delta Pr[q_0|\bar{\epsilon}] & 1 - \delta Pr[q_1|\bar{\epsilon}] & \dots & -\delta Pr[q_{M-1}|\bar{\epsilon}] \\ \vdots & \vdots & \ddots & \vdots \\ -\delta Pr[q_0|\bar{\epsilon}] & -\delta Pr[q_1|\bar{\epsilon}] & \dots & 1 - \delta Pr[q_{M-1}|\bar{\epsilon}] \\ Pr[q_0|e_0] - Pr[q_0|\bar{\epsilon}] & Pr[q_1|e_0] - Pr[q_1|\bar{\epsilon}] & \dots & Pr[q_{M-1}|e_0] - Pr[q_{M-1}|\bar{\epsilon}] \\ Pr[q_0|e_1] - Pr[q_0|\bar{\epsilon}] & Pr[q_1|e_1] - Pr[q_1|\bar{\epsilon}] & \dots & Pr[q_{M-1}|e_1] - Pr[q_{M-1}|\bar{\epsilon}] \\ \vdots & \vdots & \ddots & \vdots \\ Pr[q_0|e_{L-1}] - Pr[q_0|\bar{\epsilon}] & Pr[q_1|e_{L-1}] - Pr[q_1|\bar{\epsilon}] & \dots & Pr[q_{M-1}|e_{L-1}] - Pr[q_{M-1}|\bar{\epsilon}] \end{pmatrix}; \quad (5.25)$$

$$B = (1 - \delta) \begin{pmatrix} (p(\bar{\epsilon}) - c(\bar{\epsilon})) \\ (p(\bar{\epsilon}) - c(\bar{\epsilon})) \\ \vdots \\ (p(\bar{\epsilon}) - c(\bar{\epsilon})) \\ (c(e_0) - c(\bar{\epsilon})) / \delta \\ (c(e_1) - c(\bar{\epsilon})) / \delta \\ \vdots \\ (c(e_{L-1}) - c(\bar{\epsilon})) / \delta \end{pmatrix};$$

The dual problem corresponding to this linear program is

$$\begin{aligned} \min. \quad & B^T y \\ \text{s.t.} \quad & A^T y \geq C, y \geq 0 \end{aligned}$$

where for every constraint in the primal, there is a variable in the dual, and for every variable in the primal, there is a constraint in the dual. We will index dual variables in the following way:  $y^T = ((y_k)_{k=0, \dots, M-1}, (y^i)_{i=0, \dots, L-1, e_i \neq \bar{e}})$ , where  $y_k$ , respectively  $y^i$ , are dual variables corresponding to the primal constraints:

$$(-\delta Pr[q_0|\bar{e}], \dots, 1 - \delta Pr[q_k|\bar{e}], \dots, -\delta Pr[q_{M-1}|\bar{e}])x \leq (1 - \delta)(p(\bar{e}) - c(\bar{e}));$$

respectively,

$$(Pr[q_0|e_i] - Pr[q_0|\bar{e}], \dots, Pr[q_{M-1}|e_i] - Pr[q_{M-1}|\bar{e}])x \leq \frac{(1 - \delta)}{\delta}(c(e_i) - c(\bar{e}));$$

If the primal problem is feasible, the dual must be feasible as well. Moreover, the optimal solution of the primal corresponds to the optimal solution of the dual, and vice versa. Assume that the optimal dual solution has  $y^i = 0$  for all  $i = 0, \dots, L - 1, e_i \neq \bar{e}$ . By simplifying the constraints in the dual problem, we get:

$$y_j - \delta Pr[q_j|\bar{e}] \sum_{k=0}^{M-1} y_k \geq \delta Pr[q_j|\bar{e}]; \quad \forall j = 0, \dots, M - 1;$$

As  $\sum_{k=0}^{M-1} y_k \geq 0$  and  $Pr[q_j|\bar{e}] > 0$ ,  $y_j$  is strictly greater than 0 for all  $j = 0, \dots, M - 1$ . Therefore the corresponding primal constraints are satisfied to equality, which gives  $U_k = \bar{U}$  for all  $k = 0, \dots, M - 1$ . Unless  $\bar{U} = 0$  this is impossible.

As a consequence, at least one of the dual variables  $y^i$  has an optimal solution strictly greater than 0. Let  $i^*$  be such that  $y^{i^*} > 0$  and  $e^* = e^{i^*}$ . The corresponding primal constraint is satisfied to equality by the optimal solution and therefore:

$$\delta \sum_{q_k \in Q} (Pr[q_k|\bar{e}] - Pr[q_k|e^*])U_k = (1 - \delta)(c(e^*) - c(\bar{e}));$$

Consequently, the seller is indifferent between exerting effort  $e^*$  or  $\bar{e}$  in the first round of the PPE strategy that gives him  $\bar{U}$ . ■

## Appendix 5.D Proof of Proposition 5.4.2

*The probability that the client reports negative feedback on the equilibrium path of any pareto-optimal PPE strategy is zero.*

PROOF.

*Step 1.* Following the principle of dynamic programming (Abreu et al., 1990), the payoff profile  $V = (V_C, V_P)$  is a PPE of  $G^\infty$ , if and only if there is a strategy profile  $\sigma$  in  $G$ , and the continuation PPE payoffs profiles  $\{W(y)|y \in Y\}$  of  $G^\infty$ , such that:

- $V$  is obtained by playing  $\sigma$  in the current round, and a PPE strategy that gives  $W(y)$  as a continuation payoff, where  $y$  is the public outcome of the current round, and  $Pr[y|\sigma]$  is the probability of observing  $y$  after playing  $\sigma$ :

$$V_C = (1 - \delta)g_C(\sigma) + \delta \left( \sum_{y \in Y} Pr[y|\sigma] \cdot W_C(y) \right);$$

$$V_P = (1 - \delta)g_P(\sigma) + \delta \left( \sum_{y \in Y} Pr[y|\sigma] \cdot W_P(y) \right);$$

- no player finds it profitable to deviate from  $\sigma$ :

$$V_C \geq (1 - \delta)g_C((\sigma'_C, \sigma_P)) + \delta \left( \sum_{y \in Y} Pr[y|(\sigma'_C, \sigma_P)] \cdot W_C(y) \right); \quad \forall \sigma'_C \neq \sigma_C$$

$$V_P \geq (1 - \delta)g_P((\sigma_C, \sigma'_P)) + \delta \left( \sum_{y \in Y} Pr[y|(\sigma_C, \sigma'_P)] \cdot W_P(y) \right); \quad \forall \sigma'_P \neq \sigma_P$$

The strategy  $\sigma$  and the payoff profiles  $\{W(y)|y \in Y\}$  are said to *enforce*  $V$ .

*Step 2.* Take the PPE payoff profile  $V = (V_C, V_P)$ , such that there is no other PPE payoff profile  $V' = (V'_C, V_P)$  with  $V_C < V'_C$ . Let  $\sigma$  and  $\{W(y)|y \in Y\}$  enforce  $V$ , and assume that  $\sigma$  assigns positive probability  $\beta_0 = Pr[q_0 0|\sigma] > 0$  to the outcome  $q_0 0$ . If  $\beta_1 = Pr[q_0 1|\sigma]$  (possibly equal to 0), let us consider:

- the strategy profile  $\sigma' = (\sigma'_C, \sigma_P)$  where  $\sigma'_C$  is obtained from  $\sigma_C$  by asking the client to report 1 instead of 0 when she receives low quality (i.e.,  $q_0$ );
- the continuation payoffs  $\{W'(y)|y \in Y\}$  such that  $W'_i(q_0 1) = \beta_0 W_i(q_0 0) + \beta_1 W_i(q_0 1)$  and  $W'_i(y \neq q_0 1) = W_i(y)$  for  $i \in \{C, P\}$ . Since, the set of correlated PPE payoff profiles of  $G^\infty$  is convex, if  $W(y)$  are PPE payoff profiles, so are  $W'(y)$ .

The payoff profile  $(V'_C, V_P)$ ,  $V'_C = V_C + (1 - \delta)\beta_0 \varepsilon$  is a PPE equilibrium profile because it can be enforced by  $\sigma'$  and  $\{W'(y)|y \in Y\}$ . However, this contradicts our assumption that  $V'_C < V_C$ , so  $Pr[q_0 0|\sigma]$  must be 0. Following exactly the same argument, we can prove that  $Pr[q_1 0|\sigma] = 0$ .

*Step 3.* Taking  $V$ ,  $\sigma$  and  $\{W(y)|y \in Y\}$  from step 2, we have:

$$V_C = (1 - \delta)g_C(\sigma) + \delta \left( \sum_{y \in Y} Pr[y|\sigma] \cdot W_C(y) \right); \quad (5.26)$$

If no other PPE payoff profile  $V' = (V'_C, V_P)$  can have  $V'_C > V_C$ , it must be that the continuation payoffs  $W(y)$  satisfy the same property. (Assume otherwise that there is a PPE  $(W'_C(y), W_P(y))$  with  $W'_C(y) > W_C(y)$ . Replacing  $W'_C(y)$  in (5.26) we obtain  $V'$  that contradicts the hypothesis).

By continuing the recursion, we obtain that the client never reports 0 on the equilibrium path that enforces a payoff profile as defined in Step 2. Pareto-optimal payoff profiles clearly enter this category, hence the result of the proposition. ■

## Appendix 5.E Proof of Proposition 5.4.3

The upper bound on the percentage of false reports recorded by the reputation mechanism in any PPE equilibrium is:

$$\gamma \leq \begin{cases} \frac{(1-\alpha)(p-u)+p\rho}{p} & \text{if } p\rho \leq u(1-\alpha); \\ \frac{p\rho}{u} & \text{if } p\rho > u(1-\alpha) \end{cases}$$

PROOF. Since clients never report negative feedback along pareto-optimal equilibria, the only false reports recorded by the reputation mechanism appear when the provider delivers low quality, and the client reports positive feedback. Let  $\sigma = (\sigma_C, \sigma_P)$  be a pareto-optimal PPE strategy profile.  $\sigma$  induces a probability distribution over public histories and, therefore, over expected outcomes in each of the following transactions. Let  $\mu_t$  be the probability distribution induced by  $\sigma$  over the outcomes in round  $t$ .  $\mu_t(q_00) = \mu_t(q_10) = 0$  as proven by Proposition 5.4.2. The payoff received by the client when playing  $\sigma$  is therefore:

$$V_C(\sigma) \leq (1-\delta) \sum_{t=0}^{\infty} \delta^t (\mu_t(q_01)(-p) + \mu_t(q_11)(u-p) + \mu_t(l)0 + \mu_t(out)(u-p-p\rho));$$

where  $\mu_t(q_01) + \mu_t(q_11) + \mu_t(l) + \mu_t(out) = 1$  and  $\mu_t(q_01) + \mu_t(l) \geq (1-\alpha)\mu_t(q_11)/\alpha$ , because the probability of  $q_0$  is at least  $(1-\alpha)/\alpha$  times the probability of  $q_1$ .

When the discount factor,  $\delta$ , is the probability that the repeated interaction will stop after each transaction, the expected probability of the outcome  $q_01$  is:

$$\gamma = (1-\delta) \sum_{t=0}^{\infty} \delta^t \mu_t(q_01);$$

Since any PPE profile must give the client at least  $V_C = u - p(1 + \rho)$ , (otherwise the client is better off by resorting to the outside option),  $V_C(\sigma) \geq \underline{V}_C$ . By replacing the expression of  $V_C(\sigma)$ , and taking into account the constraints on the probability of  $q_1$  we obtain:

$$\gamma(-p) + (u-p) \cdot \min(1-\gamma, \alpha) \leq \underline{V}_C;$$

$$\gamma \leq \begin{cases} \frac{(1-\alpha)(p-u)+p\rho}{p} & \text{if } p\rho \leq u(1-\alpha); \\ \frac{p\rho}{u} & \text{if } p\rho > u(1-\alpha) \end{cases}$$

■





## Chapter 6

# Understanding Existing Online Feedback

The previous chapters have presented several results that can be used to design better online reputation mechanisms. Chapter 3 addressed signaling reputation mechanisms and discussed methods for providing honest reporting incentives to rational users. Chapter 4 described a concrete application of signaling reputation mechanism for building better quality of service monitoring system for markets of (web-)services. Chapter 5, on the other hand, addressed sanctioning reputation mechanisms, and presented methods and algorithms for designing reputation mechanisms that (i) encourage the cooperative behavior and (ii) make honest reporting rational.

While new systems could greatly benefit from these and similar research results, we will still see a great number of plain online reputation mechanisms that (i) provide little or no reporting incentives, and (ii) aggregate feedback into reputation information in trivial ways. Nevertheless, such naive mechanisms will continue to be important channel for *Word-of-mouth* regarding products, services or other types of commercial interactions, and users will most likely continue using them for taking purchasing decisions (Houser and Wooders, 2006; Melnik and Alm, 2002; Kalyanam and McIntyre, 2001; Dellarocas et al., 2006).

Recent analysis, however, raises important questions regarding the ability of existing forums to reflect the real quality of a product. In the absence of clear incentives, users with a moderate outlook will not bother to voice their opinions, which leads to an unrepresentative sample of reviews. For example, Hu et al. (2006) and Admati and Pfleiderer (2000) show that Amazon<sup>1</sup> ratings of books or CDs follow with great probability bi-modal, U-shaped distributions where most of the ratings are either very good, or very bad. Controlled experiments, on the other hand, reveal opinions on the same items that are normally distributed. Under these circumstances, using the arithmetic mean to predict quality (as most forums actually do) gives the typical user an estimator with high variance that is often false.

It is important, therefore, to better understand the interactions that take place in existing online reputation mechanisms, and come up with better ways of using the existing information. One such step has been made by Hu et al. (2006), who explain when users submit online feedback. They propose the “Brag-and-Moan Model” where users rate only if their utility of the product (drawn from a normal distribution) falls outside a median interval. The authors conclude that the model explains the empirical distribution of reports, and offers insights into smarter ways of estimating the true quality of the product.

---

<sup>1</sup><http://www.amazon.com>

In this chapter we extend this line of research, and explain further facts about the behavior of users when reporting online feedback. Using actual hotel reviews from the TripAdvisor<sup>2</sup> website, we consider two additional sources of information besides the basic numerical ratings submitted by users. The first is simple linguistic evidence from the textual review that usually accompanies the numerical ratings. We use text-mining techniques similar to Ghose et al. (2005) and Cui et al. (2006), however, we are only interested in identifying *what* aspects of the service the user is discussing, without computing the semantic orientation of the text. We find that users who comment more on the same feature are more likely to agree on a common numerical rating for that particular feature. Intuitively, lengthy comments reveal the importance of the feature to the user. Since people tend to be more knowledgeable in the aspects they consider important, users who discuss a given feature in more details might be assumed to have more *authority* in evaluating that feature.

Second we investigate the relationship between a review and the reviews that preceded it. A perusal of online reviews shows that ratings are often part of discussion threads, where one post is not necessarily independent of other posts. One may see, for example, users who make an effort to contradict, or vehemently agree with, the remarks of previous users. By analyzing the time sequence of reports, we conclude that past reviews influence the future reports, as they create some prior expectation regarding the quality of service. The subjective perception of the user is influenced by the gap between the prior expectation and the actual performance of the service (Parasuraman et al., 1985, 1988; Olshavsky and Miller, 1972; Teas, 1993) which will later reflect in the user's rating. We propose a model that captures the dependence of ratings on prior expectations, and validate it using the empirical data we collected.

Both results can be used to improve the way reputation mechanisms aggregate the information from individual reviews. Our first result can be used to determine a feature-by-feature estimate of quality, where for each feature, a different subset of reviews (i.e., those with lengthy comments of that feature) is considered. The second leads to an algorithm that outputs a more precise estimate of the real quality.

## 6.1 The Data Set

We consider real hotel reviews collected from the popular travel site TripAdvisor. TripAdvisor indexes hotels from cities across the world, along with reviews written by travelers. Users can search the site by giving the hotel's name and location (optional). The reviews for a given hotel are displayed as a list (ordered from the most recent to the oldest), with 5 reviews per page. The reviews contain:

- information about the author of the review (e.g., dates of stay, username of the reviewer, location of the reviewer);
- the overall rating (from 1, lowest, to 5, highest);
- a textual review containing a title for the review, free comments, and the main things the reviewer liked and disliked;
- numerical ratings (from 1, lowest, to 5, highest) for different features (e.g., cleanliness, service, location, etc.)

Below the name of the hotel, TripAdvisor displays the address of the hotel, general information (number of rooms, number of stars, short description, etc), the average overall rating, the TripAdvisor ranking, and an average rating for each feature. Figure 6.1 shows the page for a popular Boston hotel whose name (along with advertisements) was explicitly erased.

---

<sup>2</sup><http://www.tripadvisor.com/>

The screenshot displays the TripAdvisor interface for a Boston hotel. At the top, the TripAdvisor logo and the text "get the truth. then go.™" are visible. The page title is "Boston Hotel: Traveler Reviews". The breadcrumb trail shows the location: Home > United States > Massachusetts > Boston > Hotels > [Redacted] > Traveler Reviews. Navigation tabs for "Boston tourism" include Overview, Hotels, Flights, Attractions, Dining, Deals, Getaways, Forums, and Maps, etc. A search bar is present on the left. The hotel's name is [Redacted], and its address is 350 Stuart Street, Back Bay, Boston, MA 02116. The hotel has a 4.5-star rating and 222 rooms. A "CHECK RATES!" button shows an average price of \$421. A badge indicates it is the "Popularity Index #1" of 71 hotels in Boston. A description states it is a multi-story hotel in a 1920 landmark building. The "TripAdvisor Traveler Rating" section shows a 4.5-star rating based on 173 reviews. A "What to expect" section lists ratings for Rooms, Service, Value, Cleanliness, and Pool. A "Recommendations" section shows bars for Young singles, Couples/romantics, Business travelers, An amazing honeymoon, and Older travelers. A detailed review is provided, along with a "My ratings for this hotel are:" section with ratings for Value, Rooms, Location, Cleanliness, Check in / front desk, and Service. A "Know Boston?" section encourages users to contribute to the online guidebook. A "Free Newsletter" sign-up is also visible.

**Figure 6.1:** The TripAdvisor page displaying reviews for a popular Boston hotel. Name of hotel and advertisements were deliberately erased.

We selected three cities for this study: Boston, Sydney and Las Vegas. For each city we considered all hotels that had at least 10 reviews, and recorded all reviews. Table 6.1 presents the number of hotels considered in each city, the total number of reviews recorded for each city, and the distribution of hotels with respect to the star-rating (as available on the TripAdvisor site). Note that not all hotels have a star-rating.

For each review we recorded the overall rating, the textual review (title and body of the review) and the numerical rating on 7 features: *Rooms*(R), *Service*(S), *Cleanliness*(C), *Value*(V), *Food*(F), *Location*(L) and *Noise*(N). TripAdvisor does not require users to submit anything other than the overall rating, hence a typical review rates few additional features, regardless of the discussion in the textual comment. Only the features *Rooms*(R), *Service*(S), *Cleanliness*(C) and *Value*(V) are rated by a significant number of users. However, we also selected the features *Food*(F), *Location*(L) and *Noise*(N) because they are referred to in a significant number of textual comments. For each feature we record the numerical rating given by the user, or 0 when the rating is missing. The typical length of the textual

**Table 6.1:** A summary of the data set.

City	# Reviews	# Hotels	# of Hotels with 1,2,3,4 & 5 stars
Boston	3993	58	1+3+17+15+2
Sydney	1371	47	0+0+9+13+10
Las Vegas	5593	40	0+3+10+9+6

comment amounts to approximately 200 words. All data was collected by crawling the TripAdvisor site in September 2006.

### 6.1.1 Formal notation

We will formally refer to a review by a tuple  $(r, T)$  where:

- $r = (r_f)$  is a vector containing the ratings  $r_f \in \{0, 1, \dots, 5\}$  for the features  $f \in F = \{O, R, S, C, V, F, L, N\}$ ; note that the overall rating,  $r_O$ , is abusively recorded as the rating for the feature *Overall(O)*;
- $T$  is the textual comment that accompanies the review.

Reviews are indexed according to the variable  $i$ , such that  $(r^i, T^i)$  is the  $i^{th}$  review in our database. Since we don't record the username of the reviewer, we will also say that the  $i^{th}$  review in our data set was submitted by user  $i$ . When we need to consider only the reviews of a given hotel,  $h$ , we will use  $(r^{i(h)}, T^{i(h)})$  to denote the  $i^{th}$  review about the hotel  $h$ .

## 6.2 Evidence from Textual Comments

The free textual comments associated to online reviews are a valuable source of information for understanding the reasons behind the numerical ratings left by the reviewers. The text may, for example, reveal concrete examples of aspects that the user liked or disliked, thus justifying some of the high, respectively low ratings for certain features. The text may also offer guidelines for understanding the preferences of the reviewer, and the weights of different features when computing an overall rating.

The problem, however, is that free textual comments are difficult to read. Users are required to scroll through many reviews and read mostly repetitive information. Significant improvements would be obtained if the reviews were automatically interpreted and aggregated. Unfortunately, this seems a difficult task for computers since human users often use witty language, abbreviations, cultural specific phrases, and the figurative style.

Nevertheless, several important results use the textual comments of online reviews in an automated way. Using well established natural language techniques, reviews or parts of reviews can be classified as having a positive or negative *semantic orientation*. B. et al. (2002) classify movie reviews into positive/negative by training three different classifiers (Naive Bayes, Maximum Entropy and SVM) using classification features based on unigrams, bigrams or part-of-speech tags.

Dave et al. (2003) analyze reviews from CNet and Amazon, and surprisingly show that classification features based on unigrams or bigrams perform better than higher-order  $n$ -grams. This result is challenged by Cui et al. (2006) who look at large collections of reviews crawled from the web. They show that the size of the data set is important, and that bigger training sets allow classifiers to successfully use more complex classification features based on  $n$ -grams.

Hu and Liu (2004) also crawl the web for product reviews and automatically identify product attributes that have been discussed by reviewers. They use Wordnet to compute the semantic orientation of product evaluations and summarize user reviews by extracting positive and negative evaluations of different product features. Popescu and Etzioni (2005) analyze a similar setting, but use search engine hit-counts to identify product attributes; the semantic orientation is assigned through the *relaxation labeling technique*.

Ghose et al. (2005, 2006) analyze seller reviews from the Amazon secondary market to identify the different dimensions (e.g., delivery, packaging, customer support, etc.) of reputation. They parse the text, and tag the part-of-speech for each word. Frequent nouns, noun phrases and verbal phrases are identified as dimensions of reputation, while the corresponding *modifiers* (i.e., adjectives and adverbs) are used to derive numerical scores for each dimension. The enhanced reputation measure correlates better with the pricing information observed in the market. Pavlou and Dimoka (2006) analyze eBay reviews and find that textual comments have an important impact on reputation premiums.

Our approach is similar to the previously mentioned works, in the sense that we identify the aspects (i.e., hotel features) discussed by the users in the textual reviews. However, we do not compute the semantic orientation of the text, nor attempt to infer missing ratings.

We define the *weight*,  $w_f^i$ , of feature  $f \in F \setminus \{O\}$  in the text  $T^i$  associated with the review  $(r^i, T^i)$ , as the fraction of  $T^i$  dedicated to discussing aspects (both positive and negative) related to feature  $f$ . We propose an elementary method to approximate the values of these weights. For each feature we manually construct the word list  $L_f$  containing approximately 50 words that are most commonly associated to the feature  $f$ . The initial words were selected from reading some of the reviews, and seeing what words coincide with discussion of which features. The list was then extended by adding all thesaurus entries that were related to the initial words. Finally, we brainstormed for missing words that would normally be associated with each of the features.

Let  $L_f \cap T^i$  be the list of terms common to both  $L_f$  and  $T^i$ . Each term of  $L_f$  is counted the number of times it appears in  $T^i$ , with two exceptions:

- in cases where the user submits a title to the review, we account for the title text by appending it three times to the review text  $T^i$ . The intuitive assumption is that the user's opinion is more strongly reflected in the title, rather than in the body of the review. For example, many reviews are accurately summarized by titles such as "*Excellent service, terrible location*" or "*Bad value for money*";
- certain words that occur only once in the text are counted multiple times if their relevance to that feature is particularly strong. These were 'root' words for each feature (e.g., 'staff' is a root word for the feature *Service*), and were weighted either 2 or 3. Each feature was assigned up to 3 such root words, so almost all words are counted only once.

The list of words for the feature Rooms is given for reference in the Appendix 6.A.

The weight  $w_f^i$  is computed as:

$$w_f^i = \frac{|L_f \cap T^i|}{\sum_{f \in F \setminus \{O\}} |L_f \cap T^i|} \quad (6.1)$$

where  $|L_f \cap T^i|$  is the number of terms common to  $L_f$  and  $T^i$ . To keep a uniform notation, we also define the weight for the feature *Overall(O)* as the normalized length of the entire textual comment associated to a review:

$$w_O^i = \frac{|T^i|}{\max_i |T^i|};$$

where  $|T^i|$  is the number of character in the textual comment  $T^i$ .

The following is a TripAdvisor review for a Boston hotel (the name of the hotel is omitted): *"I'll start by saying that I'm more of a Holiday Inn person than a \*\*\* type. So I get frustrated when I pay double the room rate and get half the amenities that I'd get at a Hampton Inn or Holiday Inn. The location was definitely the main asset of this place. It was only a few blocks from the Hynes Center subway stop and it was easy to walk to some good restaurants in the Back Bay area. Boylston isn't far off at all. So I had no trouble with foregoing a rental car and taking the subway from the airport to the hotel and using the subway for any other travel. Otherwise, they make you pay for anything and everything. And when you've already dropped \$215/night on the room, that gets frustrating. The room itself was decent, about what I would expect. Staff was also average, not bad and not excellent. Again, I think you're paying for location and the ability to walk to a lot of good stuff. But I think next time I'll stay in Brookline, get more amenities, and use the subway a bit more.*

This numerical ratings associated to this review are  $r_O = 3$ ,  $r_R = 3$ ,  $r_S = 3$ ,  $r_C = 4$ ,  $r_V = 2$  for features *Overall(O)*, *Rooms(R)*, *Service(S)*, *Cleanliness(C)* and *Value(V)* respectively. The ratings for the features *Food(F)*, *Location(L)* and *Noise(N)* are absent (i.e.,  $r_F = r_L = r_N = 0$ ).

The weights  $w_f$  are computed from the following lists of common terms:

$$L_R \cap T = \{\text{room}\}; w_R = 0.066$$

$$L_S \cap T = \{3 * \text{Staff, amenities}\}; w_S = 0.267$$

$$L_C \cap T = \emptyset; w_C = 0$$

$$L_V \cap T = \{\$, \text{rate}\}; w_V = 0.133$$

$$L_F \cap T = \{\text{restaurant}\}; w_F = 0.067$$

$$L_L \cap T = \{2 * \text{center, 2 * walk, 2 * location, area}\}; w_L = 0.467$$

$$L_N \cap T = \emptyset; w_N = 0$$


The root words 'Staff' and 'Center' were tripled and doubled respectively. The overall weight of the textual review (i.e., its normalized length) is  $w_O = 0.197$ . These values account reasonably well for the weights of different features in the discussion of the reviewer.

One point to note is that some terms in the lists  $L_f$  possess an inherent semantic orientation. For example the word 'grime' (belonging to the list  $L_C$ ) would be used most often to assert the presence, and not the absence of grime. This is unavoidable, but care was taken to ensure words from both sides of the spectrum were used. For this reason, some lists such as  $L_R$  contain only nouns of objects that one would typically describe in a room (see Appendix 6.A).

The goal of this section is to analyze the influence of the weights  $w_f^i$  on the numerical ratings  $r_f^i$ .

Intuitively, users who spent a lot of their time discussing a feature  $f$  (i.e.,  $w_f^i$  is high) had something to say about their experience with regard to this feature. Obviously, feature  $f$  is important for user  $i$ . Since people tend to be more knowledgeable in the aspects they consider important, our hypothesis is that the ratings  $r_f^i$  (corresponding to high weights  $w_f^i$ ) constitute a subset of “expert” ratings for feature  $f$ .

Figure 6.2 plots the distribution of the rates  $r_C^{i(h)}$  with respect to the weights  $w_C^{i(h)}$  for the cleanliness of a Las Vegas hotel,  $h$ . Here, the high ratings are restricted to the reviews that discuss little the cleanliness. Whenever cleanliness appears in the discussion, the ratings are low. Many hotels exhibit similar rating patterns for various features. Ratings corresponding to low weights span the whole spectrum from 1 to 5, while the ratings corresponding to high weights are more grouped together.



**Figure 6.2:** The distribution of ratings against the weight of the cleanliness feature.

We therefore make the following hypothesis:

**Hypothesis 6.2.1** *The ratings  $r_f^i$  corresponding to the reviews where  $w_f^i$  is high, are more similar to each other than to the overall collection of ratings.*


To test the hypothesis, we take the entire set of reviews, and feature by feature, we compute the standard deviation of the ratings with high weights, and the standard deviation of the entire set of ratings. High weights were defined as those belonging to the upper 20% of the weight range for the corresponding feature. If Hypothesis 6.2.1 were true, the standard deviation of all ratings should be higher than the standard deviation of the ratings with high weights.

We use a standard T-test to measure the significance of the results. City by city and feature by feature, Table 6.2 presents the average standard deviation of all ratings, and the average standard deviation of ratings with high weights. Indeed, the ratings with high weights have lower standard deviation, and the results are significant at the standard 0.05 significance threshold (although for certain cities taken independently there doesn't seem to be a significant difference, the results are significant for the entire data set). Please note that only the features O,R,S,C and V were considered, since for the others (F, L, and N) we didn't have enough ratings.

Hypothesis 6.2.1 not only provides some basic understanding regarding the rating behavior of online users, it also suggests some ways of computing better quality estimates. We can, for example, construct a feature-by-feature quality estimate with much lower variance: for each feature we take the subset of

**Table 6.2:** Average standard deviation for *all* ratings, and average standard deviation for ratings with *high* weights. In square brackets, the corresponding p-values for a positive difference between the two.

City		O	R	S	C	V
Boston	all	1.189	0.998	1.144	0.935	1.123
	high	0.948	0.778	0.954	0.767	0.891
	p-val	[0.000]	[0.004]	[0.045]	[0.080]	[0.009]
Sydney	all	1.040	0.832	1.101	0.847	0.963
	high	0.801	0.618	0.691	0.690	0.798
	p-val	[0.012]	[0.023]	[0.000]	[0.377]	[0.037]
Vegas	all	1.272	1.142	1.184	1.119	1.242
	high	1.072	0.752	1.169	0.907	1.003
	p-val	[0.0185]	[0.001]	[0.918]	[0.120]	[0.126]



**Figure 6.3:** The distribution of hotels depending on the average overall rating (only reviews corresponding to high weights).


reviews that amply discuss that feature, and output as a quality estimate the average rating for this subset. Initial experiments suggest that the average feature-by-feature ratings computed in this way are different from the average ratings computed on the whole data set.

The first objection one might raise against this method is that ratings corresponding to high weights are likely to come from passionate users and are therefore likely to have extreme values. The distribution plotted Figure 6.2 supports this claim, as users who write detailed comments about the cleanliness of the hotel are mostly unhappy. Similarly, for other hotels one might expect to see the users who write a lot about a certain feature to agree more, but only on extreme ratings (ratings of 1 or 5).


This, however, does not seem to be the case for the TripAdvisor data set. As another experiment, we took all hotels from a given city, and for each hotel,  $h$ , we computed the average of all ratings  $r^{i(h)}_f$  for the feature  $f$ , where the corresponding weight,  $w_f^{i(h)}$ , was high (i.e., belongs to the upper 20% of the weight range for that feature). Figure 6.3 plot the distribution of hotels for the three cities, depending on the average of the overall ratings corresponding to high weights. For Boston (the left-most graph in Figure 6.3), the average overall rating from long reviews is almost normally distributed around 3.5, with 2 peaks around 2.5 and 4. In Sydney, on the other hand (the middle graph in Figure 6.3) the average overall rating of long reviews is almost uniformly distributed.

Similar patterns can be seen by analyzing the distribution of hotels depending on the average value





**Figure 6.4:** The distribution of hotels depending on the average value rating (only reviews corresponding to high weights).



**Figure 6.5:** The distribution of hotels depending on the average service rating (only reviews corresponding to high weights).

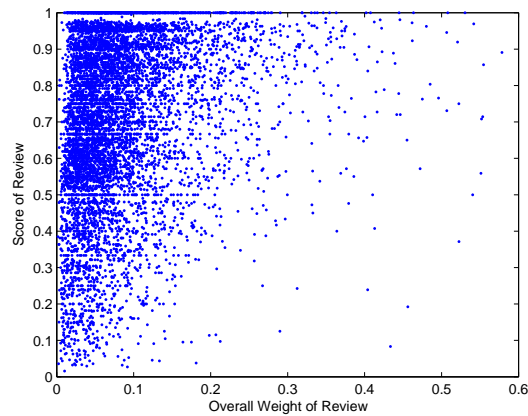
or service rating (where, again, the average was made only for those reviews discussing a lot the value, respectively the service of the hotel. Figure 6.5 presents the distribution of hotels as a function of the average rating for Service, Figure 6.4 presents the similar graphs for the value ratings. In both cases, most of the hotels have an average rating between 2 and 4, which means that most of the users discussing a lot a certain feature are not providing extreme ratings.

The second objection against building feature by feature estimators who average only the ratings corresponding to high weights, is that users who comment in more detail on a certain feature are not necessary experts of the domain. Consequently, their opinion should not count more than the opinion of other users.

This objection can partly be refuted by the following experiment. Tripadvisor lets users *vote* for the helpfulness of a review. By a simple click, a user can vote whether a particular review was useful or not. The system tracks the total number of votes received by every review, and displays below the review a footnote indicating how many of the total votes were positive.

Let the *score* of a review be the fraction of positive votes received by the review. This score can be regarded as an accurate estimator of the review's quality for the following two reasons. First, voting for a review requires almost no effort. The mouse click expressing a vote does not require authentication, and does not perturb the user when browsing for information (i.e., does not trigger a reload of the web page). As opposed to users who write a full review, the users who vote do not need the internal benefits (e.g., extreme satisfaction or dissatisfaction) to compensate the cost of reporting. Moreover, the same user can easily vote for different reviews, expressing thus a relative order between reviews.

Second, the way the information is displayed (normal font size, below the main text of a review, without any visual pointers) makes it unlikely that the score of the review has an important influence



**Figure 6.6:** The score of reviews plotted against their total length.

on the final decision of a client. This also means that there probably are few incentives to falsely cast votes, which together with the first point, lead to the conclusion that the score of a review probably reflects a representative opinion regarding the helpfulness of the review.

Figure 6.6 plots the score of the reviews in our data set against their weight for the feature *Overall* (which is actually proportional to the total length of the review). Every point in the figure represents a review. One can see that the long reviews tend to have high scores, so they were generally helpful to the other users. However, the inverse is not true: more helpful reviews are not necessarily longer.

Testing the hypothesis that high weights are indicators of *expert* opinions, requires, however, further experiments. There is no data available on TripAdvisor today which might reveal the correlation between the weight and the quality of the rating for a certain feature. However, one should devise controlled experiments where human users assess the quality of the ratings. This remains a direction for future work.

### 6.2.1 Correlation between Reporting Effort and Transactional Risk

Another factor behind the feedback bias observed on Amazon, is the relative low value of the items purchased in the market. Books and CD's are generally seen as very cheap, so the risk associated with buying a boring book or a bad CD is quite low. Feedback from previous clients decreases the information asymmetry of future buyers, and therefore their risk of choosing the wrong item. However, if the transaction poses little risk in the first place, the contribution of a feedback report to the further decrease of risk is so small, that it does not compensate the effort of reporting<sup>3</sup>.

Feedback reporting is not rational in such environments, since neither the reporter, nor the community benefit from the costly action of reporting. Therefore, the feedback that gets submitted is probably a consequence of the strong emotional response of some users who strongly like or disliked the product. This internal emotional motivation can explain the bimodal, u-shaped distribution of ratings observed on Amazon (Hu et al., 2006).

On Amazon, the correlation between reporting incentives and perceived risk is difficult to analyze,

<sup>3</sup>We thank Vincent Schickel-Zuber for suggesting this idea.

since most books and CDs have comparable prices, and involve comparable risks. On TripAdvisor, on the other hand, some hotels can be an order of magnitude more expensive than others. Moreover, an inappropriate hotel may ruin a vacation, a business meeting or a romantic weekend, so the risk of choosing a bad hotel probably exceeds by far the amount paid for the room. The right feedback in this context is very valuable to future users, as it can make the difference between a memorable trip and a dreadful one. This added value to the future travelers should motivate more users to contribute with feedback.


The data collected from TripAdvisor cannot reveal a positive correlation between the risk posed by a hotel and the motivation to submit feedback, since there is no way of estimating the actual percentage of users who left feedback for a particular hotel. However, the TripAdvisor data set can be used to study the correlation between the risk associated to a hotel and the effort spent by previous users in describing their experience with that hotel. Intuitively, we expect that the reviewers of high risk hotels spend more effort in submitting their reviews, as a consequence of feeling stronger motivated to share their experience.

Before presenting the results, let us explain our choices for measuring risk and effort. The effort spent in writing a review is measured by the length of the textual comment accompanying a review, and this choice is simpler to argument. The TripAdvisor feedback submission form is the same for everybody, so the difference between a fast and a careful review is given by the level of detail given in the textual comment. It is true that (i) some users can convey more information in shorter text, and (ii) shorter, concise English is harder to write than a long, sloppy text. Nevertheless, we expect that on the average a longer textual comment contains more information than a short one, and therefore signals more effort.

We will use as a measure for risk the official star-rating of the hotels. The intuition behind this choice is the following. A traveler who books a room in a one or two stars hotel probably expects as little as a decent bed, a relatively clean room, and a minimum of assistance from the hotel's staff. Given the high competition and the strict hygiene legislation in Boston, Las Vegas or Sydney (the cities chosen for our study) any hotel is likely to meet these basic requirements, keeping thus the traveler happy. Therefore, the risk taken by the traveler in choosing a low-end hotel is minimum. A four or five stars hotel, on the other hand, exposes a traveler to a much higher risk. The traveler chooses to pay more because she seeks certain qualities and services that were not offered by the cheaper alternative hotels. A business person, for example, might need reliable communication facilities, a concierge who can recommend good restaurants and arrange local transportation, or a fitness center to relax after a stressful day. These facilities are arguably essential to the success of the trip, and therefore, the business person has a lot to lose if she doesn't obtain them from the chosen hotel.

There are two important reasons for choosing the star-rating as a measure for risk, instead of the more straight-forward information on price. First, the price information available on the TripAdvisor is not reliable. TripAdvisor is not a booking agency, and the price they quote is only informative, obtained by averaging prices quoted by their booking partners. The pricing structure of hotel rooms is very complex today (prices depend on the season, on the type of room, but also on the occupancy of the hotel), and is often subject to special offers and discounts. Therefore, the average displayed by TripAdvisor is probably a very poor estimate of what the users really paid for the room. A quick manual scan of some reviews suffices to reveal important differences between the price quoted by the users in their textual comments, and the average price displayed by TripAdvisor.

Star rating, on the other hand, are fixed, and much easier to verify. Tourist offices, or Yellow Pages directories have detailed lists of the hotels available in a certain city, together with their official star ratings. Eventual errors in the TripAdvisor database are easy to spot by the users, and therefore, will probably be corrected. Again, a manual crosschecking with the information published by booking sites like Expedia or Travelocity will convince you that the star rating recorded by TripAdvisor is probably



**Figure 6.7:** The average normalized length of all textual comments, as a function of the star rating.

correct.

The second reason for choosing the star rating instead of the price for assessing risk, is that the risk associated to choosing a hotel is relative to the other alternatives available to the traveler. Consider Alice, who needs to travel to a remote location where there is only one available hotel. There is no risk for Alice in choosing the wrong hotel, since there is no choice at all. On the other hand, when Alice travels to a city like Boston where there is plenty of choice, she does incur a significant risk for choosing a particular hotel since many other hotels could have offered her similar or better service for lower price.

Therefore, assessing the risk of a hotel based on its prices requires a normalization of prices respective to the other hotels available in a given location. Hotels in the center of Boston are significantly more expensive than hotels in Las Vegas; paying 200 dollars for a hotel in Boston is by far less risky that paying the same amount for a hotel in Las Vegas. Since we do not have neither accurate price information, nor precise location information, the corresponding normalization is impossible to do.

The star rating, on the other hand, provides a better relative ranking of hotels, since the distribution of hotels depending on the star rating tends to be the same for all locations. The risk of a four-star hotel is probably the same, whether in Boston, or in Las Vegas, because both cities are assumed to offer a comparable choice of lower-stars hotels.


To study the correlation between the effort spent for writing a review and the risk associated to a hotel, we conducted the following experiment. For each hotel  $h$ , we computed the average normalized length of the comments present in the reviews submitted about the hotel  $h$ . We then grouped all hotels depending on their star rating, and computed the mean, respectively the standard deviation, of the average review length. Figure 6.7 plots the two values as a function of the number of stars. Indeed, higher-rated hotels receive, on the average, longer reviews, which supports our hypothesis that users spend more effort in reviewing the 'riskier' hotels. However, due to the large variance of the review length, we cannot conclude statistically that reviews about  $x$ -star hotels are significantly<sup>4</sup> longer than the reviews submitted about  $x + 1$ -star hotels, where  $x$  can be one, two, three or four. Nevertheless, when hotels are split into two risk groups:

<sup>4</sup>The T-tests were conducted at a 5% significance level

- the *low-risk* hotels have one, two, or three stars;
- the *high-risk* hotels have four or five stars;

there is a significant increase in the length of reviews submitted about the high-risk hotels as compared to the length of the reviews about low-risk hotels. The p-value for the corresponding T-test is  $1.8 \cdot 10^{-4}$ .

Visually, the difference between the length of reviews received by high-risk and low-risk hotels can be better seen in Figure 6.8. Here, we only considered the three-longest reviews submitted about every hotel. The plot displays the average and the standard deviation of the 3 longest reviews about all hotels in a certain star category. Clearly, the most detailed reviews of four and five star hotels are significantly longer than the most detailed reviews submitted about the less-than-four-stars hotels.

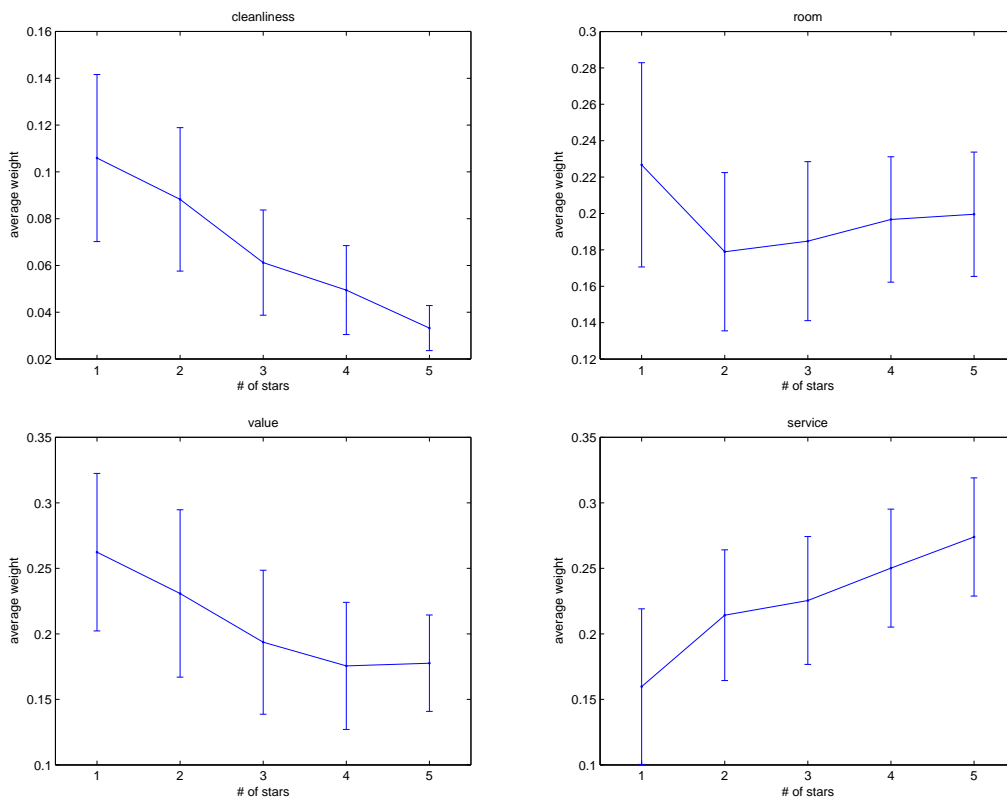


**Figure 6.8:** The average normalized length of the 3 longest textual comments, as a function of the star rating.

It is also interesting to consider how the reviews address different aspects of quality depending on the star rating of the hotels. Figure 6.9 plots the average weight of the features *Cleanliness*, *Rooms*, *Value* and *Service* for hotels with different numbers of stars. The cleanliness, for example, is mostly discussed for the low-end hotels. Four or five star hotels are expected to be impeccably clean, and probably are very clean. The cleanliness is definitely not a risk factor for high-end hotel, hence reviewers spend little time discussing it. For low-end hotels, on the other hand, the cleanliness can be a major decision factor, and there is a significantly higher risk of choosing a dirty hotel. Hence the increased fraction of the comments addresses the cleanliness.

The same trend is observable for the fraction of the text discussing the value of the hotel, although overall, the value is discussed much more than the cleanliness. High-end hotels apparently have a very well established tradeoff between price and quality, which makes this feature a low risk factor. For low-end hotels, on the other side, there can be large variations between the value of different hotels, hence the presence of this feature in the reviews. The service, on the other hand, becomes increasingly important for high-end hotels. As argued previously, the travelers who chose fancier hotels do so because they need or like certain services that are not available in the cheaper hotels. The main risk associated to choosing four or five star hotels comes from not getting the services you want, hence naturally, the reviews for these hotels go into more detail regarding the services offered by the hotel.

The quality of the room, on the other hand does not seem to vary a lot depending on the star rating of the hotel. The rooms tend to be discussed more for one star hotels, however, the difference is not



**Figure 6.9:** The fraction of the comment taken by different features, as a function of the star rating.

statistically significant.

### 6.3 The Influence of Past Ratings

Two important assumptions are generally made about reviews submitted to online forums. The first is that ratings truthfully reflect the quality observed by the users; the second is that reviews are independent from one another. While anecdotal evidence (Harmon, 2004; White, 1999) challenges the first assumption<sup>5</sup>, in this section, we address the second.

A perusal of online reviews shows that reviews are often part of discussion threads, where users make an effort to contradict, or vehemently agree with the remarks of previous users. Consider, for example, the following review:

*"I don't understand the negative reviews... the hotel was a little dark, but that was the style. It was very artsy. Yes it was close to the freeway, but in my opinion the sound of an occasional loud car is better than hearing the "ding ding" of slot machines all night! The staff on-hand is FABULOUS. The waitresses are great (and \*\*\* does not deserve the bad review she got, she was 100% attentive to us!), the bartenders are friendly and professional at the same time..."*

<sup>5</sup>Part of Amazon reviews were recognized as strategic posts by book authors or competitors

Here, the user was disturbed by previous negative reports, addressed these concerns, and set about trying to correct them. Not surprisingly, his ratings were considerably higher than the average ratings up to this point.

It seems that TripAdvisor users regularly read the reports submitted by previous users before booking a hotel, or before writing a review. Past reviews create some prior expectation regarding the quality of service, and this expectation has an influence on the submitted review. We believe this observation holds for most online forums. The subjective perception of quality is directly proportional to how well the actual experience meets the prior expectation, a fact confirmed by an important line of econometric and marketing research (Parasuraman et al., 1985, 1988; Olshavsky and Miller, 1972; Teas, 1993).

The correlation between the reviews has also been confirmed by recent research on the dynamics of online review forums (Forman et al., 2006).

### 6.3.1 Prior Expectations

We define the prior expectation of user  $i$  regarding the feature  $f$ , as the average of the previously available ratings on the feature  $f$ <sup>6</sup>:

$$e_f(i) = \frac{\sum_{j < i, r_f^j \neq 0} r_f^j}{\sum_{j < i, r_f^j \neq 0} 1}$$

As a first hypothesis, we assert that the rating  $r_f^i$  is a function of the prior expectation  $e_f(i)$ :

**Hypothesis 6.3.1** *For a given hotel and feature, given the reviews  $i$  and  $j$  such that  $e_f(i)$  is high and  $e_f(j)$  is low, the rating  $r_f^j$  exceeds the rating  $r_f^i$ .*

We define *high* and *low* expectations as those that are above, respectively below a certain cutoff value  $\theta$ . The set of reviews preceded by high, respectively low expectations are defined as follows:

$$R_f^{high} = \{r_f^i | e_f(i) > \theta\}$$

$$R_f^{low} = \{r_f^i | e_f(i) < \theta\}$$

These sets are specific for each (hotel, feature) pair, and in our experiments we took  $\theta = 4$ . This rather high value is close to the average rating across all features across all hotels, and is justified by the fact that our data set contains mostly high quality hotels.

For each city, we take all hotels and compute the average ratings in the sets  $R_f^{high}$  and  $R_f^{low}$  (see Table 6.3). The average rating amongst reviews following low prior expectations is significantly higher than the average rating following high expectations.

There are two ways to interpret the function  $e_f(i)$ :

---

<sup>6</sup>if no previous ratings were assigned for feature  $f$ ,  $e_f(i)$  is assigned a default value of 4.

**Table 6.3:** Average ratings for reviews preceded by low (first value in the cell) and high (second value in the cell) expectations. The P-values for a positive difference are given within square brackets.

City	O	R	S	C	V
Boston	3.953	4.045	3.985	4.252	3.946
	3.364	3.590	3.485	3.641	3.242
	[0.011]	[0.028]	[0.0086]	[0.0168]	[0.0034]
Sydney	4.284	4.358	4.064	4.530	4.428
	3.756	3.537	3.436	3.918	3.495
	[0.000]	[0.000]	[0.035]	[0.009]	[0.000]
Las Vegas	3.494	3.674	3.713	3.689	3.580
	3.140	3.530	2.952	3.530	3.351
	[0.190]	[0.529]	[0.007]	[0.529]	[0.253]

- The expected value for feature  $f$  obtained by user  $i$  before his experience with the service, acquired by reading reports submitted by past users. In this case, an overly high value for  $e_f(i)$  would drive the user to submit a negative report (or vice versa), stemming from the difference between the actual value of the service, and the inflated expectation of this value acquired before his experience.
- The expected value of feature  $f$  for all subsequent visitors of the site, if user  $i$  were not to submit a report. In this case, the motivation for a negative report following an overly high value of  $e_f$  is different: user  $i$  seeks to *correct* the expectation of future visitors to the site. Unlike the interpretation above, this does not require the user to derive an *a priori* expectation for the value of  $f$ .

Note that neither interpretation implies that the average up to report  $i$  is inversely related to the rating at report  $i$ . There might exist a measure of influence exerted by past reports that pushes the user behind report  $i$  to submit ratings which to some extent conforms with past reports: a low value for  $e_f(i)$  can influence user  $i$  to submit a low rating for feature  $f$  because, for example, he fears that submitting a high rating will make him out to be a person with low standards<sup>7</sup>. This, at first, appears to contradict Hypothesis 6.3.1. However, this conformity rating cannot continue indefinitely: once the set of reports project a sufficiently deflated estimate for  $v_f$ , future reviewers with comparatively positive impressions will seek to correct this misconception.

### 6.3.2 Impact of Textual Comments on Quality Expectation

Further insight into the rating behavior of TripAdvisor users can be obtained by analyzing the relationship between the weights  $w_f$  and the values  $e_f(i)$ . In particular, we examine the following hypothesis:

**Hypothesis 6.3.2** *When a large proportion of the text of a review discusses a certain feature, the difference between the rating for that feature and the average rating up to that point tends to be large.*

<sup>7</sup>The idea that negative reports can encourage further negative reporting has been suggested before Khopkar and Resnick (2005)



The intuition behind this claim is that when the user is adamant about voicing his opinion regarding a certain feature, his opinion differs from the collective opinion of previous postings. This relies on the characteristic of reputation systems as feedback forums where a user is interested in projecting his opinion, with particular strength if this opinion differs from what he perceives to be the general opinion.

To test Hypothesis 6.3.2 we measure the average absolute difference between the expectation  $e_f(i)$  and the rating  $r_f^i$  when the weight  $w_f^i$  is high, respectively low. Weights are classified high or low by comparing them with certain cutoff values:  $w_f^i$  is low if smaller than 0.1, while  $w_f^i$  is high if greater than  $\theta_f$ . Different cutoff values were used for different features:  $\theta_R = 0.4$ ,  $\theta_S = 0.4$ ,  $\theta_C = 0.2$ , and  $\theta_V = 0.7$ . *Cleanliness* has a lower cutoff since it is a feature rarely discussed; *Value* has a high cutoff for the opposite reason. Results are presented in Table 6.4.

**Table 6.4:** Average of  $|r_f^i - e_f(i)|$  when weights are high (first value in the cell) and low (second value in the cell) with P-values for the difference in sq. brackets.

City	R	S	C	V
Boston	1.058	1.208	1.728	1.356
	0.701	0.838	0.760	0.917
	[0.022]	[0.063]	[0.000]	[0.218]
Sydney	1.048	1.351	1.218	1.318
	0.752	0.759	0.767	0.908
	[0.179]	[0.009]	[0.165]	[0.495]
Las Vegas	1.184	1.378	1.472	1.642
	0.772	0.834	0.808	1.043
	[0.071]	[0.020]	[0.006]	[0.076]

This demonstrates that when weights are unusually high, users tend to express an opinion that does not conform to the net average of previous ratings. As we might expect, for a feature that rarely was a high weight in the discussion, (e.g., cleanliness) the difference is particularly large. Even though the difference in the feature *Value* is quite large for Sydney, the P-value is high. This is because only few reviews discussed value heavily. The reason could be cultural or because there was less of a reason to discuss this feature.

### 6.3.3 Reporting Incentives

Previous models suggest that users who are not highly opinionated will not choose to voice their opinions (Hu et al., 2006). In this section, we extend this model to account for the influence of expectations. The motivation for submitting feedback is not only due to extreme opinions, but also to the difference between the current reputation (i.e., the prior expectation of the user) and the actual experience.

Such a rating model produces ratings that most of the time deviate from the current average rating. The ratings that confirm the prior expectation will rarely be submitted. We test on our data set the proportion of ratings that attempt to “correct” the current estimate. We define a deviant rating as one that deviates from the current expectation by at least some threshold  $\theta$ , i.e.,  $|r_f^i - e_f(i)| \geq \theta$ . For each of the three considered cities, the following tables, show the proportion of deviant ratings for  $\theta = 0.5$  and  $\theta = 1$ .

**Table 6.5:** Proportion of deviant ratings with  $\theta = 0.5$ 

City	O	R	S	C	V
Boston	0.696	0.619	0.676	0.604	0.684
Sydney	0.645	0.615	0.672	0.614	0.675
Las Vegas	0.721	0.641	0.694	0.662	0.724

**Table 6.6:** Proportion of deviant ratings with  $\theta = 1$ 

City	O	R	S	C	V
Boston	0.420	0.397	0.429	0.317	0.446
Sydney	0.360	0.367	0.442	0.336	0.489
Las Vegas	0.510	0.421	0.483	0.390	0.472

The above results suggest that a large proportion of users (close to one half, even for the high threshold value  $\theta = 1$ ) deviate from the prior average. This reinforces the idea that users are more likely to submit a report when they believe they have something distinctive to add to the current stream of opinions for some feature. Such conclusions are in total agreement with prior evidence that the distribution of reports often follows bi-modal, U-shaped distributions.

## 6.4 Modelling the Behavior of Raters

To account for the observations described in the previous sections, we propose a model for the behavior of the users when submitting online reviews. For a given hotel, we make the assumption that the quality experienced by the users is normally distributed around some value  $v_f$ , which represents the “objective” quality offered by the hotel on the feature  $f$ . The rating submitted by user  $i$  on feature  $f$  is:

$$\hat{r}_f^i = \delta_f v_f^i + (1 - \delta_f) \cdot \text{sign}(v_f^i - e_f(i)) \left[ c + d(v_f^i, e_f(i) | w_f^i) \right] \quad (6.2)$$

where:

- $v_f^i$  is the (unknown) quality actually experienced by the user.  $v_f^i$  is assumed normally distributed around some value  $v_f$ ;
- $\delta_f \in [0, 1]$  can be seen as a measure of the bias when reporting feedback. High values reflect the fact that users rate objectively, without being influenced by prior expectations. The value of  $\delta_f$  may depend on various factors; we fix one value for each feature  $f$ ;
- $c$  is a constant between 1 and 5;
- $w_f^i$  is the weight of feature  $f$  in the textual comment of review  $i$ , computed according to Eq. (6.1);
- $d(v_f^i, e_f(i) | w_f^i)$  is a distance function between the expectation and the observation of user  $i$ . The distance function satisfies the following properties:

$$- d(y, z | w) \geq 0 \text{ for all } y, z \in [0, 5], w \in [0, 1];$$

- $|d(y, z|w)| < |d(z, x|w)|$  if  $|y - z| < |z - x|$ ;
- $|d(y, z|w_1)| < |d(y, z|w_2)|$  if  $w_1 < w_2$ ;
- $c + d(v_f, e_f(i)|w_f^i) \in [1, 5]$ ;

The second term of Eq. (6.2) encodes the bias of the rating. The higher the distance between the true observation  $v_f^i$  and the function  $e_f$ , the higher the bias.

### 6.4.1 Model Validation

We use the data set of TripAdvisor reviews to validate the behavior model presented above. We split for convenience the rating values in three ranges: bad ( $B = \{1, 2\}$ ), indifferent ( $I = \{3, 4\}$ ), and good ( $G = \{5\}$ ), and perform the following two tests:

- First, we will use our model to predict the ratings that have extremal values. For every hotel, we take the sequence of reports, and whenever we encounter a rating that is either good or bad (but not indifferent) we try to predict it using Eq. (6.2)
- Second, instead of predicting the value of extremal ratings, we try to classify them as either good or bad. For every hotel we take the sequence of reports, and for each report (regardless of its value) we classify it as being good or bad

However, to perform these tests, we need to estimate the objective value,  $v_f$ , that is the average of the true quality observations,  $v_f^i$ . The algorithm we are using is based on the intuition that the amount of conformity rating is minimized. In other words, the value  $v_f$  should be such that as often as possible, bad ratings follow expectations above  $v_f$  and good ratings follow expectations below  $v_f$ .

Formally, we define the sets:

$$\begin{aligned}\Gamma_1 &= \{i | e_f(i) < v_f \text{ and } r_f^i \in B\}; \\ \Gamma_2 &= \{i | e_f(i) > v_f \text{ and } r_f^i \in G\};\end{aligned}$$

that correspond to irregularities where even though the expectation at point  $i$  is lower than the delivered value, the rating is poor, and vice versa. We define  $v_f$  as the value that minimizes the union of the two sets:

$$v_f = \arg \min_{v_f} |\Gamma_1 \cup \Gamma_2| \quad (6.3)$$

In Eq. (6.2) we replace  $v_f^i$  by the value  $v_f$  computed in Eq. (6.3), and use the following distance function:

$$d(v_f, e_f(i)|w_f^i) = \frac{|v_f - e_f(i)|}{v_f - e_f(i)} \sqrt{|v_f^2 - e_f(i)^2|} \cdot (1 + 2w_f^i);$$

The constant  $c \in I$  was set to  $\min\{\max\{e_f(i), 3\}, 4\}$ . The values for  $\delta_f$  were fixed at  $\{0.7, 0.7, 0.8, 0.7, 0.6\}$  for the features  $\{Overall, Rooms, Service, Cleanliness, Value\}$  respectively. The weights are computed as described in Section 6.2.

As a first experiment, we take the sets of “extremal” ratings  $\{r_f^i | r_f^i \notin I\}$  for each hotel and feature. For every such rating,  $r_f^i$ , we try to estimate it by computing  $\hat{r}_f^i$  using Eq. (6.2). We compare this

estimator with the one obtained by simply averaging the ratings over all hotels and features: i.e.,

$$\bar{r}_f = \frac{\sum_{j, r_f^j \neq 0} r_f^j}{\sum_{j, r_f^j \neq 0} 1};$$

Table 6.7 presents the ratio between the root mean square error (RMSE) when using  $\hat{r}_f^i$  and  $\bar{r}_f$  to estimate the actual ratings. In all cases the estimate produced by our model is better than the simple average.

**Table 6.7:** Average of  $\frac{RMSE(\hat{r}_f)}{RMSE(\bar{r}_f)}$

City	O	R	S	C	V
Boston	0.987	0.849	0.879	0.776	0.913
Sydney	0.927	0.817	0.826	0.720	0.681
Las Vegas	0.952	0.870	0.881	0.947	0.904

As a second experiment, we try to distinguish the sets  $B_f = \{i | r_f^i \in B\}$  and  $G_f = \{i | r_f^i \in G\}$  of bad, respectively good ratings on the feature  $f$ . For example, we compute the set  $B_f$  using the following classifier (called  $\sigma$ ):

$$r_f^i \in B_f (\sigma_f(i) = 1) \Leftrightarrow \hat{r}_f^i \leq 4;$$

Tables 6.8, 6.9 and 6.10 present the Precision( $p$ ), Recall( $r$ ) and  $s = \frac{2pr}{p+r}$  for classifier  $\sigma$ , and compares it with a naive majority classifier,  $\tau$ ,  $\tau_f(i) = 1 \Leftrightarrow |B_f| \geq |G_f|$ :

**Table 6.8:** Precision(p), Recall(r),  $s = \frac{2pr}{p+r}$  while spotting poor ratings for Boston

		O	R	S	C	V
$\sigma$	p	0.678	0.670	0.573	0.545	0.610
	r	0.626	0.659	0.619	0.612	0.694
	s	0.651	0.665	0.595	0.577	0.609
$\tau$	p	0.684	0.706	0.647	0.611	0.633
	r	0.597	0.541	0.410	0.383	0.562
	s	0.638	0.613	0.502	0.471	0.595

We see that recall is always higher for  $\sigma$  and precision is usually slightly worse. For the  $s$  metric  $\sigma$  tends to add a 1-20% improvement over  $\tau$ , much higher in some cases for hotels in Sydney. This is likely because Sydney reviews are more positive than those of the American cities and cases where the number of bad reviews exceeded the number of good ones are rare. Replacing the test algorithm with

**Table 6.9:** Precision(p), Recall(r),  $s = \frac{2pr}{p+r}$  while spotting poor ratings for Las Vegas

		O	R	S	C	V
$\sigma$	p	0.654	0.748	0.592	0.712	0.583
	r	0.608	0.536	0.791	0.474	0.610
	s	0.630	0.624	0.677	0.569	0.596
$\tau$	p	0.685	0.761	0.621	0.748	0.606
	r	0.542	0.505	0.767	0.445	0.441
	s	0.605	0.607	0.670	0.558	0.511

**Table 6.10:** Precision(p), Recall(r),  $s = \frac{2pr}{p+r}$  while spotting poor ratings for Sydney

		O	R	S	C	V
$\sigma$	p	0.650	0.463	0.544	0.550	0.580
	r	0.234	0.378	0.571	0.169	0.592
	s	0.343	0.452	0.557	0.259	0.586
$\tau$	p	0.562	0.615	0.600	0.500	0.600
	r	0.054	0.098	0.101	0.015	0.175
	s	0.098	0.168	0.172	0.030	0.271

one that plays a 1 with probability equal to the proportion of bad reviews improves its results for this city, but it is still outperformed by around 80%.

## 6.5 Summary of Results

The main goal of this chapter is push forward our understanding of the factors that (i) drive a user to submit feedback, and (ii) bias a user in the rating she provides to the reputation mechanism. For that we use two additional sources of information besides the vector of numerical ratings: first we look at the textual comments that accompany the reviews, and second we consider the reports that have been previously submitted by other users.

Using simple natural language processing algorithms, we were able to establish a correlation between the weight of a certain feature in the textual comment accompanying the review, and the noise present in the numerical rating. Specifically, it seems that users who discuss amply a certain feature are likely to agree on a common rating. This observation allows the construction of feature-by-feature estimators of quality that have a lower variance, and are hopefully less noisy. Initial experiments suggest that longer reviews tend to be more helpful to the other users, backing up the argument that reputation estimators should weigh more the corresponding ratings. Nevertheless, further evidence is required to support the intuition that at a feature level, high weight ratings are also more accurate, and therefore deserve higher priority when computing estimates of quality.

Using the same natural language processing of the textual comments associated to reviews, we were able to establish a correlation between the risk associated to a hotel and the effort spent in submitting the review. For the reasons detailed in Section 6.2.1 we assume that hotels with higher number of stars

present a higher risk for the travelers, in terms of taking a bad decision. The average length of the reviews submitted about high risk hotels is significantly bigger than the average length of low-end hotel reviews, meaning that users are willing to spend more effort when they perceive a higher risk of taking a bad decision. An immediate extension of this observation is that users will also be more motivated to submit feedback about high-risk transactions, however, we did not have the proper data to validate this assumption.

Second, we emphasize the dependence of ratings on previous reports. Previous reports create an expectation of quality which affects the subjective perception of the user. We validate two facts about the hotel reviews we collected from TripAdvisor: First, the ratings following low expectations (where the expectation is computed as the average of the previous reports) are likely to be higher than the ratings following high expectations. Intuitively, the perception of quality (and consequently the rating) depends on how well the actual experience of the user meets her expectation. Second, we include evidence from the textual comments, and find that when users devote a large fraction of the text to discussing a certain feature, they are likely to motivate a divergent rating (i.e., a rating that does not conform to the prior expectation). Intuitively, this supports the hypothesis that review forums act as discussion groups where users are keen on presenting and motivating their own opinion.

We have captured the empirical evidence in a behavior model that predicts the ratings submitted by the users. The final rating depends, as expected, on the true observation, and on the gap between the observation and the prior expectation. The gap tends to have a bigger influence when an important fraction of the textual comment is dedicated to discussing a certain feature. The proposed model was validated on the empirical data and provides better estimates than the simple average.

One assumption that we make is about the existence of an objective quality value  $v_f$  for the feature  $f$ . This is rarely true, especially over large spans of time. Other explanations might account for the correlation of ratings with past reports. For example, if  $e_f(i)$  reflects the true value of  $f$  at a point in time, the difference in the ratings following high and low expectations can be explained by hotel revenue models that are maximized when the value is modified accordingly. However, the idea that variation in ratings is not primarily a function of variation in value turns out to be a useful one. Our approach to approximate this elusive 'objective value' is by no means perfect, but conforms neatly to the idea behind the model.

A natural direction for future work is to examine concrete applications of our results. Significant improvements of quality estimates are likely to be obtained by incorporating all empirical evidence about rating behavior. Exactly how different factors affect the decisions of the users is not clear. The answer might depend on the particular application, context and culture.

## Appendix 6.A List of words, $L_R$ , associated to the feature Rooms

All words serve as prefixes: *room, space, interior, decor, ambiance, atmosphere, comfort, bath, toilet, bed, building, wall, window, private, temperature, sheet, linen, pillow, hot, water, cold, water, shower, lobby, furniture, carpet, air, condition, mattress, layout, design, mirror, ceiling, lighting, lamp, sofa, chair, dresser, wardrobe, closet*

## Chapter 7

# Conclusions

The internet is moving rapidly towards an interactive milieu where online communities and economies gain importance over their traditional counterparts. While this shift creates opportunities and benefits that have already improved our day-to-day life, it also brings a whole new set of problems. For example, the lack of physical interaction that characterizes most electronic transactions, leaves the systems much more susceptible to fraud and deception. According to the US Federal Trade Commission approximately one half of all consumer fraud reports in 2005 were related to the Internet.

Reputation mechanisms offer a novel and effective way of ensuring the necessary level of trust which is essential to the functioning of any market. They collect information about the history (i.e., past transactions) of market participants and make public their *reputation*. Prospective partners guide their decisions by considering reputation information, and thus make more informative choices.

Online reputation mechanisms enjoy huge success. They are present in most e-commerce sites available today, and are seriously taken into consideration by human users. Numerous empirical studies prove the existence of *reputation premiums* which allow service providers with good reputation to charge higher prices.

The economical value of online reputation raises questions regarding the trustworthiness of mechanisms themselves. Existing systems were conceived with the assumption that users will share feedback honestly. However, we have recently seen increasing evidence that some users strategically manipulate their reports. For some sites, the extent of *spam* feedback called for immediate solutions to eliminate suspicious reports. Known approaches include ad-hoc semi-manual, semi-automatic filters that provide satisfactory results only temporarily.

This thesis presents a more systematic approach to making reputation mechanisms trustworthy. Assuming that the agents participating in the market are rational (i.e., they are interested in maximizing their utility) I investigate different methods for making the reputation mechanisms truthful, such that users find in their best interest to report the truth.

The design of honest reporting incentives depends on the function and implementation of the reputation mechanism, therefore I address separately the two important roles of reputation. The first role of the reputation mechanism is to inform future users on the hidden quality attributes of the product or service they intend to purchase. The mechanism acts as a signaling device which differentiates the products or service providers depending on their innate attributes and capabilities. Chapter 3 formally describes the signaling reputation mechanisms, and then proceeds to discussing reporting incentives.

Previous work showed that explicit payments can encourage truthful reporting in signaling reputation mechanisms. The basic idea is that every reporter gets rewarded for submitting feedback, but different reporters may receive different amounts. The payment received by a reporter depends on the value of the submitted report, but also on the feedback submitted by some other user about the same product or service. When carefully designed, these feedback payments support an honest Nash equilibrium, where no agent can benefit by lying as long as all other reporters tell the truth.

I bring several novel results to the design of incentive compatible feedback payments for signaling reputation mechanisms:

1. I use the idea of automated mechanism design (Conitzer and Sandholm, 2002) to define the incentive-compatible payments that are also optimal. The payments are defined by a linear optimization problem which minimizes the expected payment to an honest reporter, and hence the expected cost to the reputation mechanism. The use of automated mechanism design does not pose unpractical computation burdens, and can decrease the budget required by the reputation mechanism by up to 70%;
2. The payments can be further decreased by computing the payment received by a reporter based on the feedback of *several* other agents. It has been formally proven that every supplementary reference report considered by the payment scheme reduces the expected cost of the mechanism. Fortunately, the biggest reductions in cost are obtained from the first few supplementary reference reports, which means that the design problem can remain simple enough to be practical;
3. Further reductions in cost can be obtained by filtering out some reports. The reputation mechanism will randomly throw away some feedback. However, the probability of discarding a report depends on the reports of other agents. When designed together, a payment and a filtering mechanism can be very efficient, and up to an order of magnitude cheaper;
4. I describe an algorithm for making the feedback payments robust to noise and private information. When the agents detain private information about a product, the payments proposed by the reputation mechanism are not necessarily incentive compatible. I give an example of how lying becomes profitable, and I describe a method for correcting the incentives as long as the amount of private information is bounded;
5. I address the problem of collusion, and show that general incentive compatible payments are vulnerable to collusion. If several agents coordinate their reporting strategies they can manipulate the reputation information without suffering payment losses. The vulnerability to collusion comes from the fact that general payment schemes have several Nash equilibria, some of which being more attractive than the honest equilibrium. I add to the design problem constraints to ensure that honest reporting is the unique or the pareto-optimal equilibrium. I describe several positive and negative results for different collusion scenarios.

Another direct contribution of this thesis is to identify novel applications for reputation mechanisms where a reward scheme brings provable properties about the reliability of reputation information. One such example is quality of service (QoS) monitoring in decentralized markets of (web-)services. By using the feedback reported by the clients, the QoS monitoring process can be made more precise, reliable and cheaper.

The second role of reputation mechanisms is to encourage cooperative behavior. In many markets the seller must exert costly effort in order to satisfy the request of the buyer. However, the buyer usually pays first and remains vulnerable to cheating; having received the payment, the seller no longer exerts the required effort. The main idea is that present feedback determines the future reputation of an agent, and implicitly affects the future revenues accessible to the agent. When carefully designed,



reputation mechanisms can make it such that the momentary gain obtained by cheating is offset by the future losses caused by a bad reputation. Hence cheating is appropriately sanctioned by the reputation mechanism, and this encourages every participant in the market to behave cooperatively.

Sanctioning reputation mechanisms can be designed in many ways, depending (i) on the granularity of feedback requested from the users, (ii) on the algorithms for aggregating feedback into reputation information, or (iii) on the extent and form of reputation information dissemination. Dellarocas (2005) discusses binary reputation mechanisms (where seller, for example, can cooperate or cheat, and buyers can receive high or low quality) and investigates the effect these design decisions on the efficiency of the mechanism.

The novel contribution of Chapter 5 is to extend Dellarocas' results to general settings, where the seller can choose between several effort levels, and buyers can observe several quality levels. The following results are worth mentioning:

1. A reputation mechanism where the reputation has a binary value (i.e., can be either good or bad) can be equally efficient to other mechanisms where the reputation is infinitely finer grained;
2. Efficient reputation mechanism may only consider a finite number of past reports, and should not consider the entire history of feedback;
3. A mechanism which accepts a finite number of different feedback values can be just as efficient as a mechanism who allows infinitely fine-grained feedback. More detailed feedback can increase the efficiency of the mechanism only to the extent it allows a better grouping of signals into a finite number of classes.

The second contribution of Chapter 5 is to discuss a mechanism for encouraging the submission of honest feedback. CONFESS works by comparing the feedback submitted by the buyer to the feedback implicitly submitted by the seller. After every transaction, the seller is allowed to acknowledge failures and reimburse the affected buyer. If, however, the seller does not reimburse the buyer, and the buyer submits negated feedback, the reputation mechanism concludes that one of the agents is lying, and punishes them both. This simple mechanism supports an equilibrium where all sellers cooperate, and all buyers report the truth. Moreover, it allows the buyers to build a reputation for always reporting the truth, which in the end, can be proven to limit the amount of false information received by the reputation mechanism in any pareto-optimal equilibrium.

Last but not least, the thesis discusses methods for improving existing online reputation mechanisms. Since we will most likely continue to have reputation mechanisms that do not address the reporting incentives of the participants, it is important to understand what are the factors that drive users to submit feedback and bias their ratings. Some recent works proposed several models for explaining why and how human users provide feedback. In Chapter 6 I extend this line of research by experimentally analyzing a database of hotel reviews from the TripAdvisor web site. The main results are:

1. The users who amply discuss a certain feature are more likely to agree on a common rating for that feature. These users are not outliers, and their reviews are generally regarded as more helpful by the other users.
2. The effort spent in writing a review is correlated to the risk (in terms of making a bad choice) associated to a hotel. High-end hotels have more risk associated to them since travelers pay more without knowing if they will get the desired service. Human raters apparently feel motivated to decrease the decision risk of future users, hence they spend more effort rating the high-end hotels.
3. The rating expressed by a reviewer is strongly biased by the reviews submitted by previous users. The information already available in the forum creates a prior expectation of quality, and changes

the user's subjective perception of quality. The gap between the expectation and the actual quality is reflected in the rating submitted to the site.

4. A great proportion of users submit ratings that are significantly different from the average of previous ratings. This leads to the conclusion that human users are more likely to voice their opinion when they can bring something different to the discussion, and can contribute with new information.

## 7.1 Directions for Future Work

The results of this thesis constitute an important theoretical foundation for a new generation of reputation mechanisms that are more practical, and bound to function under realistic conditions. I see several interesting directions that have a great potential to improve online reputation mechanisms.

### 7.1.1 From "lists of reviews" to designed reputation mechanisms

Current reputation mechanisms usually require users to submit feedback that (a) numerically rates the quality of the service on several dimensions, and (b) comments further aspects of the service that were not captured by numerical ratings. Reputation is then displayed as an average score of numerical ratings, followed by a list of the textual comments submitted by individual users. I see several problems with such systems.

First, the interpretation of reputation scores is difficult, and requires both an understanding of the rating scale, and contextual information about the market. On eBay, for example, feedback can be negative, neutral or positive, and the score of a user is displayed as the number (respectively the percentage) of positive feedback reports. However, most sellers have many positive reports and very few negative ones. To really assess the risk of being cheated, buyers must compare the profile of the chosen seller with that of sellers who trade similar products.

Moreover, scores are *inflated* by time, as all sellers continuously accumulate feedback. The trustworthiness of a seller that has one thousand positive ratings today, is not the same as the trustworthiness of a seller that will have one thousand positive reports next month. The effort invested in assessing the reputation today, is therefore not reusable the next time the same user transacts on eBay.

Second, the overall scores only contain partial information. Human users heavily use textual comments to disclose important information about sellers or items. On Amazon's secondary market, for example, it is not uncommon to discover that a book seller with good ratings has really slow delivery. For someone pressed to buy a birthday present, the overall rating is misleading. On the other hand, scrolling through a list of repetitive textual comments is time consuming and prone to missing out important information.

Third, reputation scores are sometimes aggregated in the wrong way. Available ratings are usually submitted by biased users that were either very happy, or very unhappy with the product. Using the arithmetic average to aggregate reports gives a poor estimator with high variance. Finally, available feedback reports do not necessarily reflect the truth.

I believe that a new generation of reputation mechanisms must be carefully *designed* to be as effective as possible. A reputation mechanism should clearly specify the semantics of reputation information (i.e., what is reputation information and how does it change with new feedback), and provide clear guidelines for taking decisions based on reputation information. These aspects influence one another and should

therefore be designed in equilibrium.

In a first phase, the transition to the new systems can be made by leveraging on natural language techniques to automate the extraction of information from textual comments. Initial results suggest that simple algorithms can output valuable data that improves the default aggregation of numerical scores. One can carry on this track to obtain better methods for extracting reputation information from available ratings.

Based on improved information, the mechanism can then compute and recommend optimal behavior strategies that are contingent to reputation information. In equilibrium, the truster will find it optimal to adopt the behavior recommended by the reputation mechanism, and the trustee will have incentives to do as expected.

Finally, reputation mechanisms should be augmented with incentive schemes that elicit honest feedback and prevent collusion. The schemes from Chapters 3 and 5 should be extended to more practical contexts where ratings are multi-dimensional, quality observations are subjective, and rewards can be based on something else than money (e.g., access to resources, social prestige, etc).

### 7.1.2 Signaling and sanctioning reputation mechanisms

In real markets both roles of reputation information are present in the same time: clients must choose providers with different abilities, who can all cheat or cooperate to some extent. While pure signaling, and pure sanctioning reputation mechanisms are well understood, the design of mechanisms that concomitantly fulfill both roles remains a theoretical challenge. One reason for that is an apparent conflict regarding the use of information: signaling reputation mechanisms become more efficient as they dispose of more information, while sanctioning mechanisms, on the contrary, are more efficient when they consider fewer reports (Dellarocas, 2005).

### 7.1.3 Factoring human behavior into reputation mechanism design

Pure theoretical analysis often leads to negative or impractical results. Take for example feedback payment schemes that are incentive-compatible. They make honest reporting a Nash equilibrium, but also create other lying equilibria which are sometimes more attractive than the honest one. The resulting mechanism is vulnerable to collusion, as groups of agents can coordinate their behavior on the more attractive lying equilibrium.

However, when a fraction of the reporters are assumed honest, it is possible to deter lying coalitions by making them unstable. The assumption that some reporters will unconditionally report the truth is justified by both empirical studies, and common-sense intuition. Some users are altruistic and obtain comfort from being honest. Other users blindly follow the recommendations of the mechanism designer, and do not have the skills to discover the higher paying lying strategies. Yet another category of users cannot coordinate their actions with those of a coalition.

I believe there is a great potential to factor facts about human behavior into the game-theoretic design of reputation mechanisms. Humans are not entirely rational; habit, social norms and moral values constrain their behavior towards actions that are usually more altruistic and cooperative than game-theory predicts. These new constraints should allow the design of better mechanisms.

The large literature of psychological, sociological and economical experimental research can provide numerous insights about the behavior of humans in online feedback reporting settings. These insights can then be transformed into real models by analyzing the statistical patterns of behavior extracted

from existing databases. I expect significant improvements for factoring these assumptions in the design of online reputation mechanisms.

#### 7.1.4 Mechanisms for social networks and P2P systems

One important application of reputation mechanisms is in the context of social networks and P2P systems, where no central entity can host the reputation mechanism. Feedback is distributed in the network and reflects the direct interactions among peers.

The design of reputation mechanism is significantly harder in such environments due to the new technological constraints. Communication channels are imperfect, and messages between two peers must usually be relayed through several other peers. Peers continuously join and leave the system, which often triggers reconfigurations in the network. Feedback reports are likely to get lost, the reputation of a peer can be different in different parts of the network, and can fluctuate heavily with time.

Reputation mechanisms must therefore be *robust* and computationally *efficient*. Robustness implies that peers should not be able to exploit temporary fluctuations of information. Efficiency is required to reduce the total overhead of reputation management. I believe it is possible to design a reputation mechanism with asymmetric feedback transmission. Negative feedback should be widely broadcasted so that cheaters are effectively punished. Positive feedback, on the other hand, can be implicitly assumed in order to reduce the communication overhead. Nevertheless, the problems associated with changing or faking online identities (Cheng and Friedman, 2005) should be carefully considered.

#### 7.1.5 Reputation mechanisms translated to other domains

The same design principles used for constructing better reputation mechanisms can be applied to other domains and improve current solutions. One example about Quality of Service monitoring systems was described in Chapter 4.

A future generation of **recommender systems** can also be designed to account for the strategic interests of the participants. Users can be incentivised to submit honest ratings by reward mechanisms similar to those used for signaling reputation mechanisms. Moreover, providers can be effectively deterred from manipulating the recommendations of the system by filtering techniques initially designed for reputation mechanisms. I believe that merging recommender and reputation mechanisms will result in more robust and precise systems which will improve the market efficiency.

Yet another example, novel **advertising mechanisms** can be constructed based on social networks regulated by reputation mechanisms. Online advertisements are targeted today based on contextual information about the user (e.g., keywords of interest, profile based on search and browsing activity). However, people tend to rely more on friends for getting recommendations about the items they intend to purchase. I believe there is a great potential for passing on advertisements through the social network of a user. The role of the reputation mechanism in this context is to prevent abuse, and promote the spread of responsible advertisements.

In general, any decentralized information system can employ a reputation mechanism to deter cheating, improve the quality of information, and increase the participation of rational users. I believe that IT systems which consider the incentives of individual users and make the desired behavior an equilibrium have a great potential to improve the trust and security of online applications.

# Bibliography

- Abdul-Rahman, A., Hailes, S., 2000. Supporting Trust in Virtual Communities. In: Proceedings Hawaii International Conference on System Sciences. Maui, Hawaii.
- Aberer, K., Despotovic, Z., 2001. Managing Trust in a Peer-2-Peer Information System. In: Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM).
- Abreu, D., Sen, A., 1991. Virtual Implementation in Nash Equilibria. *Econometrica* 59, 997–1022.
- Abreu, P., Pearce, D., Stacchetti, E., 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58 (5), 1041 – 1063.
- Admati, A., Pfleiderer, P., 2000. Noisytalk.com: Broadcasting opinions in a noisy environment. Working Paper 1670R, Stanford University.
- Akerlof, G. A., 1970. The market for 'lemons': Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics* 84 (3), 488–500.
- Alunkal, B., Veljkovic, I., Laszewski, G., Amin, K., 2003. Reputation-Based Grid Resource Selection. In: Proceedings of AGridM.
- Anderson, C., 2006. *The Long Tail: Why the Future of Business is Selling Less of More?* Hyperion.
- Andrieux, A., Czajkowski, K., Dan, A., Keahey, K., Ludwig, H., Pruyne, J., Rofrano, J., Tuecke, S., Xu, M., 2005. Web Services Agreement Specification (WS-Agreement), Version 2005/09. [http://www.ggf.org/Public\\_Comment\\_Docs/Documents/Oct2005/WS-AgreementSpecificationDraft050920.pdf](http://www.ggf.org/Public_Comment_Docs/Documents/Oct2005/WS-AgreementSpecificationDraft050920.pdf).
- Aumann, R., Sorin, S., 1989. Cooperation and Bounded Recall. *Games and Economic Behavior* 1, 5–39.
- Avery, C., Resnick, P., Zeckhauser, R., 1999. The Market for Evaluations. *American Economics Review* 89 (3), 564–584.
- Axelrod, R., 1984. *The Evolution of Cooperation*. Basic Books, New York.
- B., P., Lee, L., Vaithyanathan, S., 2002. Thumbs up? sentiment classification using machine learning techniques. In: Proceedings of the EMNLP-02, the Conference on Empirical Methods in Natural Language Processing.
- Ba, S., Pavlou, P. A., 2002. Evidence of the Effect of Trust Building Technology in Electronic Markets: Price Premiums and Buyer Behavior. *MIS Quarterly* 26, 243–268.
- Bacharach, M., 2002. How Human Trusters Assess Trustworthiness in Quasi-Virtual Contexts. In: Proceedings of the AAMAS Workshop on Trust Deception and Fraud. Bologna, Italy.

- Bajari, P., Hortacsu, A., 2004. Economic Insights from Internet Auctions: A Survey. *Journal of Economic Literature* 42 (2), 457–486.
- Barber, S., Kim, J., 2001. Belief Revision Process Based on Trust: Agents Evaluating Reputation of Information Sources. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 73–82.
- Barbon, F., Traverso, P., Pistore, M., Trainotti, M., 2006. Run-Time Monitoring of Instances and Classes of Web-Service Compositions. In: *Proceedings of ICWS 2006*.
- Berg, J., Dickhaut, J., McCabe, K., 1995. Trust, Reciprocity and Social History. *Games and Economic Behavior* 10 (1), 122–42.
- Bernheim, B. D., Ray, D., 1989. Collective Dynamic Consistency in Repeated Games. *Games and Economic Behavior* 1, 295–326.
- Beth, T., Borchering, M., Klein, B., 1994. Valuation of Trust in Open Networks. In: *Proceedings of the European Symposium on Research in Computer Security (ESORICS)*. Springer-Verlag, Brighton, UK, pp. 3–18.
- Bianculli, D., Ghezzi, C., 2007. Monitoring conversational web services. In: *Proceedings of the 2nd International Workshop on Service-Oriented Software Engineering (IW-SOSWE'07)*.
- Birk, A., 2000. Boosting Cooperation by Evolving Trust. *Applied Artificial Intelligence* 14, 769–784.
- Birk, A., 2001. Learning to Trust. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 133–144.
- Biswas, A., Sen, S., Debnath, S., 2000. Limiting Deception in a Group of Social Agents. *Applied Artificial Intelligence* 14, 785–797.
- Bolton, G., Katok, E., Ockenfels, A., 2004. How Effective Are Electronic Reputation Mechanisms? An Experimental Investigation. *Management Science* 50 (11), 1587–1602.
- Braynov, S., Sandholm, T., 2002. Incentive Compatible Mechanism for Trust Revelation. In: *Proceedings of the AAMAS*. Bologna, Italy.
- Buchegger, S., Le Boudec, J.-Y., 2003. The Effect of Rumour Spreading in Reputation Systems for Mobile Ad-hoc Networks. In: *WiOpt '03: Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*. Sophia-Antipolis, France.
- Buchegger, S., Le Boudec, J.-Y., 2005. Self-Policing Mobile Ad-Hoc Networks by Reputation. *IEEE Communication Magazine* 43 (7), 101–107.
- Buskens, V., Barrera, D., 2005. Third-party effects on trust in an embedded investment game. Paper presented at the annual meeting of the American Sociological Association, Philadelphia, [http://www.allacademic.com/meta/p18354\\_index.html](http://www.allacademic.com/meta/p18354_index.html).
- Buttayan, L., Hubaux, J.-P., 1999. Toward a Formal Model of Fair Exchange - a Game Theoretic Approach. Tech. rep.
- Castelfranchi, C., Falcone, R., 2000. Trust and Control: A Dialectic Link. *Applied Artificial Intelligence* 14, 799–823.
- Chan, J., 2000. On the Non-Existence of Reputation Effects in Two-Person Infinitely-Repeated Games. <http://www.econ.jhu.edu/People/Chan/reputation.pdf>.

- Chen, K.-Y., Hogg, T., Wozny, N., 2004. Experimental Study of Market Reputation Mechanisms. In: Proceedings of the ACM Conference on Electronic Commerce (EC'04). pp. 234–235.
- Cheng, A., Friedman, E., 2005. Sybilproof reputation mechanisms. In: Proceeding of the Workshop on Economics of Peer-to-Peer Systems (P2PECON). pp. 128–132.
- Chevalier, J., Mayzlin, D., 2006. The Effect of Word of Mouth on Sales: Online Book Reviews. *Journal of Marketing Research*. Forthcoming.
- Clemen, R. T., 2002. Incentive contracts and strictly proper scoring rules. *Test* 11, 167–189.
- Conitzer, V., Sandholm, T., 2002. Complexity of mechanism design. In: Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI).
- Conitzer, V., Sandholm, T., 2003a. Applications of Automated Mechanism Design. In: Proceedings of the the UAI-03 Bayesian Modeling Applications Workshop.
- Conitzer, V., Sandholm, T., 2003b. Automated Mechanism Design with a Structured Outcome Space.
- Conitzer, V., Sandholm, T., 2004. An Algorithm for Automatically Designing Deterministic Mechanisms without Payments. In: Proceedings of the AAMAS-04.
- Conitzer, V., Sandholm, T., 2007. Incremental Mechanism Design. In: Proceedings of the IJCAI.
- Conte, R., Paolucci, M., 2002. Reputation in Artificial Societies. *Social Beliefs for Social Order*. Kluwer, Boston (MA).
- Cox, J. C., Deck, C. A., 2005. On the Nature of Reciprocal Motives. *Economic Inquiry* 43 (3), 623–635.
- Cr mer, J., McLean, R. P., 1985. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica* 53 (2), 345–61.
- Cripps, M., Mailath, G., Samuelson, L., 2004. Imperfect Monitoring and Impermanent Reputations. *Econometrica* 72, 407–432.
- Cripps, M. W., Dekel, E., Pesendorfer, W., 2005. Reputation with Equal Discounting in Repeated Games with Strictly Conflicting Interests.
- Cripps, M. W., Thomas, J. P., 1997. Reputation and Perfection in Repeated Common Interest Games. *Games and Economic Behavior* 18, 141–158.
- Cui, H., Mittal, V., Datar, M., 2006. Comparative Experiments on Sentiment Classification for Online Product Reviews. In: Proceedings of AAAI.
- Dan, A., Davis, D., Kearney, R., Keller, A., King, R. P., Kuebler, D., Ludwig, H., Polan, M., Spreitzer, M., Youssef, A., 2004. Web services on demand: WSLA-driven automated management. *IBM Systems Journal* 43 (1), 136–158.
- d'Aspremont, C., Gard-Varet, L.-A., 1979. Incentives and Incomplete Information. *Journal of Public Economics* 11, 25–45.
- Dave, K., Lawrence, S., Pennock, D., 2003. Mining the peanut gallery: opinion extraction and semantic classification of product reviews. In: Proceedings of the 12th International Conference on the World Wide Web (WWW03).
- Dellarocas, C., 2000. Immunizing Online Reputation Reporting Systems Against Unfair Ratings and Discriminatory Behaviour. In: Proceedings of the 2nd ACM conference on Electronic Commerce. Minneapolis, MN.

- Dellarocas, C., 2002. Goodwill Hunting: An Economically Efficient Online Feedback. In: Padget, J., et al. (Eds.), *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*. Vol. LNCS 2531. Springer Verlag, pp. 238–252.
- Dellarocas, C., 2004. Information Society or Information Economy? A combined perspective on the digital era. Idea Book Publishing, Ch. *Building Trust On-Line: The Design of Robust Reputation Mechanisms for Online Trading Communities*, pp. 95–113.
- Dellarocas, C., 2005. Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard. *Information Systems Research* 16 (2), 209–230.
- Dellarocas, C., 2006a. *Handbook on Economics and Information Systems*. Elsevier Publishing, Ch. *Reputation Mechanisms*, pp. 629–660.
- Dellarocas, C., 2006b. Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms. *Management Science* 52 (10), 1577–1593.
- Dellarocas, C., Awad, N., Zhang, X., 2006. Exploring the Value of Online Product Ratings in Revenue Forecasting: The Case of Motion Pictures. Working paper.
- Dellarocas, C., Wood, C., 2006. The Sound of Silence in Online Feedback: Estimating Trading Risks in The Presence of Reporting Bias, under review at *Management Science*.
- Deora, V., Shao, J., Gray, W., Fiddian, J., 2003. A Quality of Service Management Framework Based on User Expectations. In: *Proceedings of ICSOC*.
- Despotovic, Z., 2005. Building Trust-aware P2P Systems: From Trust and Reputation Management to Decentralized E-Commerce Applications. Ph.D. thesis, Ecole Polytechnique Fédérale de Lausanne.
- Despotovic, Z., Aberer, K., 2004. A Probabilistic Approach to Predict Peers' Performance in P2P Networks. In: *Eighth International Workshop on Cooperative Information Agents, CIA 2004*. Erfurt, Germany.
- Deutsch, M., 1962. *Nebraska Symposium on Motivation*. Nebraska University Press, Ch. *Cooperation and trust: Some theoretical notes*.
- Dewally, M., Ederington, L. H., 2006. Reputation, Certification, Warranties, and Information as Remedies for Seller-Buyer Information Asymmetries: Lessons from the Online Comic Book Market. *The Journal of Business* 79 (2).
- Dewan, S., Hsu, V., 2004. Adverse Selection in Electronic Markets: Evidence from Online Stamp Auctions. *Journal of Industrial Economics* 52 (4), 497–516.
- Diekmann, Wyder, 2002. Vertrauen und Reputationseffekte bei Internet-Auktionen. *Kölner Zeitschrift für Soziologie und Sozialpsychologie* 54, 674–693.
- Dimitrakos, T., 2003. A Service-Oriented Trust Management Framework. In: Falcone, R., Barber, R., Korba, L., Singh, M. (Eds.), *Trust, Reputation, and Security: Theories and Practice*. Vol. LNAI 2631. Springer-Verlag, Berlin Heidelberg, pp. 53–72.
- Eaton, D. H., 2002. Valuing information: Evidence from guitar auctions on ebay.  
URL [citeseer.ist.psu.edu/eaton02valuing.html](http://citeseer.ist.psu.edu/eaton02valuing.html)
- Elliott, C., February 7, 2006. Hotel Reviews Online: In Bed With Hope, Half-Truths and Hype. *The New York Times*.



- Falcone, R., Castelfranchi, C., 2001. The Socio-cognitive Dynamics of Trust: Does Trust create Trust. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 55–72.
- Farrell, J., Maskin, E., 1989. Renegotiation in Repeated Games. *Games and Economic Behavior* 1, 327–360.
- Fehr, E., Gächter, S., 2000. Fairness and Retaliation: The Economics of Reciprocity. *Journal of Economic Perspectives* 14 (3), 159–181.
- Forman, C., Ghose, A., Wiesenfeld, B., July 2006. A Multi-Level Examination of the Impact of Social Identities on Economic Transactions in Electronic Markets, available at SSRN: <http://ssrn.com/abstract=918978>.
- Friedman, E., Parkes, D., 2003. Pricing WiFi at Starbucks - Issues in online mechanism design. In: *Proceedings of EC'03*. pp. 240–241.
- Friedman, E., Resnick, P., 2001. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy* 10(2), 173–199.
- Fudenberg, D., Levine, D., 1989. Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica* 57, 759–778.
- Fudenberg, D., Levine, D., 1992. Maintaining a Reputation when Strategies are Imperfectly Observed. *Review of Economic Studies* 59 (3), 561–579.
- Fudenberg, D., Levine, D., 1994. Efficiency and Observability with Long-Run and Short-Run Players. *Journal of Economic Theory* 62, 103–135.
- Fudenberg, D., Levine, D., Maskin, E., 1994. The Folk Theorem with Imperfect Public Information. *Econometrica* 62 (5), 997–1039.
- Fudenberg, D., Maskin, E., 1989. The Folk Theorem in Repeated Games with Discounting or Incomplete Information. *Econometrica* 54 (3), 533–554.
- Gambetta, D., 1988. *Can We Trust Trust?* Department of Sociology, University of Oxford, Ch. Trust: Making and Breaking Cooperative Relations, pp. 213–237.  
URL <http://www.sociology.ox.ac.uk/papers/gambetta213-237.pdf>
- Ghose, A., Ipeirotis, P., Sundararajan, A., 2005. Reputation Premiums in Electronic Peer-to-Peer Markets: Analyzing Textual Feedback and Network Structure. In: *Third Workshop on Economics of Peer-to-Peer Systems, (P2PECON)*.
- Ghose, A., Ipeirotis, P., Sundararajan, A., 2006. The Dimensions of Reputation in electronic Markets. Working Paper CeDER-06-02, New York University.
- Gibbard, A., 1973. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 587–601.
- Goitein, S. D., 1973. *Letters of Medieval Jewish Traders*. Princeton University Press.
- Goldreich, O., 1998. Secure multi-party computation. Working paper. Available at <http://www.wisdom.weizmann.ac.il/~oded/pp.html>.
- Greif, A., 1989. Reputation and Coalitions in Medieval Trade: Evidence on the Maghribi Traders. *The Journal of Economic History* XLIX (4), 857–882.
- Greif, A., March 2002. Institutions and Impersonal Exchange: From Communal to Individual Responsibility. *Journal of Institutional and Theoretical Economics (JITE)* 158 (1), 168–204.

- Guo, M., Conitzer, V., 2007. Worst-Case Optimal Redistribution of VCG Payments. In: Proceedings of EC'07. pp. 30–39.
- Hajiaghayi, M., Kleinberg, R., Sandholm, T., 2007. Automated Online Mechanism Design and Prophet Inequalities. In: Proceedings of AAAI'07.
- Halberstadt, A., Mui, L., 2002. A Computational Model of Trust and Reputation. In: Proceedings of the Goddard/JPL Workshop on Radical Agents Concepts. NASA Goddard Space Flight Center.
- Harmon, A., February 14, 2004. Amazon Glitch Unmasks War of Reviewers. The New York Times.
- Hennig-Thurau, T., Gwinner, K. P., Walsh, G., Gremler, D. D., 2004. Electronic Word-of-Mouth via Consumer-Opinion Platforms: What Motivates Consumers to Articulate Themselves on the Internet? *Journal of Interacting Marketing* 18 (1), 38–52.
- Holmström, B., 1982. Moral Hazard in Teams. *Bell Journal of Economics* 13, 324–340.
- Houser, D., Wooders, J., 2006. Reputation in Auctions: Theory and Evidence from eBay. *Journal of Economics and Management Strategy* 15, 353–369.
- Hu, M., Liu, B., 2004. Mining and summarizing customer reviews. In: Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD04).
- Hu, N., Pavlou, P., Zhang, J., 2006. Can Online Reviews Reveal a Product's True Quality? In: Proceedings of ACM Conference on Electronic Commerce (EC 06).
- Ismail, R., Jøsang, A., 2002. The Beta Reputation System. In: Proceedings of the 15th Bled Conf. on E-Commerce.
- Jackson, M., 2001. A crash course in implementation theory. *Social Choice and Welfare* 18 (4), 655–708.
- Jackson, M., 2003. *Encyclopedia of Life Support Systems*. EOLSS Publishers, Oxford UK, Ch. Mechanism Theory.
- Jackson, M. O., 1991. Bayesian Implementation. *Econometrica* 59, 461–477.
- Jin, G. Z., Kato, A., 2004. Blind Trust Online: Experimental Evidence from Baseball Cards. University of Maryland. Working Paper.
- Johnson, S., Pratt, J., Zeckhauser, R., 1990. Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case. *Econometrica* 58, 873–900.
- Jøsang, A., Lo Presti, S., 2004. Analysing the Relationship between Risk and Trust. In: Jensen, C., Poslad, S., Dimitrakos, T. (Eds.), *Trust Management*. Vol. LNCS 2995. Springer-Verlag, Berlin Heidelberg, pp. 135–145.
- Jurca, R., Faltings, B., 2003. An Incentive-Compatible Reputation Mechanism. In: Proceedings of the IEEE Conference on E-Commerce (CEC'03). Newport Beach, CA, USA, pp. 285–292.
- Jurca, R., Faltings, B., June 11–15 2006. Minimum Payments that Reward Honest Reputation Feedback. In: Proceedings of the ACM Conference on Electronic Commerce (EC'06). Ann Arbor, Michigan, USA, pp. 190–199.
- Kalepu, S., Krishnaswamy, S., Loke, S., 2003. Verity; A QoS Metric for Selecting Web Services and Providers. In: Proceedings of WISEW.
- Kalyanam, K., McIntyre, S., 2001. Return on reputation in online auction market. Working Paper 02/03-10-WP, Leavey School of Business, Santa Clara University.

- Kamvar, S., Schlosser, M., Garcia-Molina, H., 2003. The EigenTrust Algorithm for Reputation Management in P2P Networks. In: Proceedings of the World Wide Web Conference.
- Kandori, M., Matsushima, H., 1998. Private observation, communication and collusion. *Econometrica* 66 (3), 627–652.
- Kauffman, R. J. and Wood, C., 2000. Running up the Bid: Modeling Seller Opportunism in Internet Auctions. In: Proceedings of the Americas Conference on Information Systems. pp. 929–935.
- Keates, N., June 1, 2007. Deconstructing TripAdvisor. *The Wall Street Journal*, page W1.
- Keller, A., Ludwig, H., 2002. Defining and monitoring service-level agreements for dynamic e-business. In: Proceedings of the 16th Conference on Systems Administration.
- Keser, C., 2003. Experimental games for the design of reputation management systems. *IBM Systems Journal* 42 (3), 498–506.
- Khopkar, L., Resnick, P., 2005. Self-Selection, Slipping, Salvaging, Slacking, and Stoning: the Impacts of Negative Feedback at eBay. In: Proceedings of ACM Conference on Electronic Commerce (EC 05).
- Klein, T. J., Lambertz, C., Spagnolo, G., Stahl, K. O., 2006. Last Minute Feedback. University of Mannheim, Working Paper.
- Kramer, R., 2001. Trust Rules for Trust Dilemmas: How Decision Makers Think and Act in the Shadow of Doubt. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 9–26.
- Kreps, D. M., Milgrom, P., Roberts, J., Wilson, R., 1982. Rational Cooperation in the Finitely Repeated Prisoner's Dilemma. *Journal of Economic Theory* 27, 245–252.
- Kreps, D. M., Wilson, R., 1982. Reputation and Imperfect Information. *Journal of Economic Theory* 27, 253–279.
- Kuwabara, K., 2003. Decomposing Reputation Effects: Sanctioning or Signaling? Working paper.
- Lamsal, P., October 2001. Understanding Trust and Security. Online-Reference.  
URL <http://www.cs.helsinki.fi/u/lamsal/asgn/trust/UnderstandingTrustAndSecurity.pdf>
- Lave, L., August 1962. An Empirical Approach to the Prisoners' Dilemma Game. *The Quarterly Journal of Economics* 76 (3), 424–436.
- Li, S., Balachandran, K., 2000. Collusion proof transfer payment schemes with multiple agents. *Review of Quantitative Finance and Accounting* 15, 217–233.
- Liu, Y., Ngu, A., Yeng, L., 2004. QoS Computation and Policing in Dynamic Web Service Selection. In: Proceedings of WWW.
- Livingston, J., 2002. How valuable is a good reputation? A sample selection model of internet auctions. University of Maryland. Available on-line at <http://www.wam.umd.edu/~kth/reputation1.pdf>.
- Lucking-Reiley, D., Bryan, D., Prasad, N., Reeves, D., 2000. Pennies from eBay: The Determinants of Price in Online Auctions. *Econometric society world congress 2000 contributed papers*, Econometric Society.
- Ludwig, H., Dan, A., Kearney, R., 2004. Cremona: An architecture and library for creation and monitoring of WS-Agreements. In: ICSOC '04: Proceedings of the 2nd international conference on Service oriented computing. ACM Press, New York, NY, USA, pp. 65–74.

- Ma, C., 1988. Unique implementation of incentive contracts with many agents. *Review of Economic Studies*, 555–572.
- Mahbub, K., Spanoudakis, G., 2004. A framework for requirements monitoring of service based systems. In: *Proceedings of ICSOC*.
- Mailath, G., Samuelson, L., 2006. *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.
- Manchala, D., 1998. Trust Metrics, Models and Protocols for Electronic Commerce Transactions. In: *Proceedings of the 18th International Conference on Distributed Computing Systems*. IEEE Computer Society, pp. 312–321.
- Marsh, S., 1994. Formalising Trust as a Computational Concept. Ph.D. thesis, Department of Mathematics and Computer Science, University of Stirling.
- Marsh, S., Dibben, M. R., 2005. Trust Management. Vol. 3477 of *Lecture Notes in Computer Science*. Springer Berlin, Ch. Trust, Untrust, Distrust and Mistrust An Exploration of the Dark(er) Side, pp. 17–33.  
URL [http://dx.doi.org/10.1007/11429760\\_2](http://dx.doi.org/10.1007/11429760_2)
- Maskin, E., 1999. Nash Equilibrium and Welfare Optimality. *Review of Economic Studies* 66, 23–28.
- Maskin, E., Sjöström, T., 2002. *Handbook of Social Choice and Welfare*. Vol. 1. Elsevier, Ch. Implementation Theory, pp. 237–288.
- Matsushima, H., 1988. A New Approach to the Implementation Problem. *Journal of Economic Theory* 45, 128–144.
- Maximilien, E. M., Singh, M. P., 2004. Toward Autonomic Web Services Trust and Selection. In: *Proceedings of ICSOC*.
- McCabe, K. A., Rigdon, M. L., Smith, V. L., 2003. Positive Reciprocity and Intentions in the Trust Game. *Journal of Economic Behavior and Organization* 52 (2), 267–275.
- McDonald, C. G., Slawson Jr., V., 2002. Reputation in an Internet Auction Market. *Economic Inquiry* 40 (4), 533–650.
- McIlraith, S. A., Martin, D. L., 2003. Bringing semantics to web services. *IEEE Intelligent Systems* 18 (1), 90–93.
- McKnight, D. H., Choudhury, V., Kacmar, C., 2002. Developing and Validating Trust Measures for E-Commerce: An Integrative Typology. *Information Systems Research* 13 (3), 334–359.
- McKnight, H., Chervany, N., 2001. Trust and Distrust: One Bite at a Time. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 27–54.
- Melnik, M., Alm, J., 2002. Does a seller's reputation matter? evidence from ebay auctions. *Journal of Industrial Economics* 50 (3), 337–350.
- Milgrom, P., Roberts, J., 1982. Predation, Reputation and Entry Deterrence. *Journal of Economic Theory* 27, 280–312.
- Miller, N., Resnick, P., Zeckhauser, R., 2005. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51, 1359–1373.

- Milosevic, Z., Dromey, G., 2002. On expressing and monitoring behaviour in contracts. In: Proceedings of EDOC. Lausanne, Switzerland.
- Moore, J., Repullo, R., 2005. A characterization of virtual Bayesian implementation. *Games and Economic Behavior* 50, 312–331.
- Mui, L., 2003. Computational Models of Trust and Reputation: Agents, Evolutionary Games, and Social Networks. Ph.D. thesis, Massachusetts Institute of Technology.
- Mui, L., Halberstadt, A., Mohtashemi, M., 2002a. Notions of Reputation in Multi-Agents Systems: A Review. In: Proceedings of the AAMAS. Bologna, Italy.
- Mui, L., Mohtashemi, M., Halberstadt, A., 2002b. A Computational Model of Trust and Reputation. In: Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS).
- Nash, J., 1950. Equilibrium Points in N-person Games. *Proceedings of the National Academy of Sciences* 36, 48–49.
- Nowak, M. A., Sigmund, K., 1994. The Alternating Prisoner's Dilemma. *Journal of Theoretical Biology* 168, 219–226.
- Olshavsky, R., Miller, J., February 1972. Consumer Expectations, Product Performance and Perceived Product Quality. *Journal of Marketing Research* 9, 19–21.
- Osborne, M., Rubinstein, A., 1997. *A Course in Game Theory*. MIT Press.
- Page, L., Brin, S., Motwani, R., Winograd, T., 1998. The PageRank Citation Ranking: Bringing Order to the Web. Tech. rep., Stanford Digital Library Technologies Project.
- Palfrey, T., Srivastava, S., 1991. Nash-implementation using Undominated Strategies. *Econometrica* 59, 479–501.
- Papaioannou, T. G., Stamoulis, G. D., 2005a. An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems. In: Proceedings of IEEE/ACM CCGRID 2005.
- Papaioannou, T. G., Stamoulis, G. D., 2005b. Optimizing an Incentives' Mechanism for Truthful Feedback in Virtual Communities. In: Proceedings of AAMAS (Workshop on Agents and Peer-to-Peer Computing). Utrecht, The Netherlands.
- Papaioannou, T. G., Stamoulis, G. D., 2006. Enforcing Truthful-Rating Equilibria in Electronic Marketplaces. In: Proceedings of the IEEE ICDCS Workshop on Incentive-Based Computing. Lisbon, Portugal.
- Parasuraman, A., Zeithaml, V., Berry, L., 1985. A Conceptual Model of Service Quality and Its Implications for Future Research. *Journal of Marketing* 49, 41–50.
- Parasuraman, A., Zeithaml, V., Berry, L., 1988. SERVQUAL: A Multiple-Item Scale for Measuring Consumer Perceptions of Service Quality. *Journal of Retailing* 64, 12–40.
- Parkes, D., 2007. *Algorithmic Game Theory*. Cambridge University Press, Ch. Online Mechanisms, pp. 411–439.
- Pavlou, P., Dimoka, A., 2006. The Nature and Role of Feedback Text Comments in Online Marketplaces: Implications for Trust Building, Price Premiums, and Seller Differentiation. *Information Systems Research* 17 (4), 392–414.

- Pistore, M., Barbon, F., Bertoli, P., Shaparau, D., Traverso, P., 2004. Planning and monitoring web service composition. In: Workshop on Planning and Scheduling for Web and Grid Services (held in conjunction with The 14th International Conference on Automated Planning and Scheduling). URL <http://www.isi.edu/ikcap/icaps04-workshop/final/pistore.pdf>
- Popescu, A., Etzioni, O., 2005. Extracting product features and opinions from reviews. In: Proceedings of the Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing.
- Povey, D., 1999. Developing Electronic Trust Policies Using a Risk Management Model. In: Proceedings of the Secure Networking - CQRE (Secure)99, International Exhibition and Congress. Vol. LNCS 1740. Springer, Germany, pp. 1–16.
- Prelec, D., 2004. A bayesian truth serum for subjective data. *Science* 306 (5695), 462–466.
- Rapoport, A., Chammah, A., 1965. Prisoner's Dilemma: A Study in Conflict and Cooperation. University of Michigan Press, Ann Arbor, Michigan.
- Reichling, F., 2004. Effects of Reputation Mechanisms on Fraud Prevention in eBay Auctions. Stanford University, Working Paper.
- Resnick, P., Zeckhauser, R., 2002. Trust Among Strangers in Electronic Transactions: Empirical Analysis of eBay's Reputation System. In: Baye, M. (Ed.), *The Economics of the Internet and E-Commerce*. Vol. 11 of *Advances in Applied Microeconomics*. Elsevier Science, Amsterdam.
- Resnick, P., Zeckhauser, R., Swanson, J., Lockwood, K., 2006. The Value of Reputation on eBay: A Controlled Experiment. *Experimental Economics* 9 (2), 79–101.
- Richardson, M., Agrawal, R., Domingos, P., 2003. Trust Management for the Semantic Web. In: Proceedings of the Second International Semantic Web Conference. Sanibel Island, FL, pp. 351–368.
- Robinson, W. N., 2003. Monitoring web service requirements. In: RE'03: Proceedings of the 11th IEEE International Conference on Requirements Engineering. IEEE Computer Society, Washington, DC, USA, p. 65.
- Roth, A., December 1988. Laboratory Experimentation in Economics: A Methodological Overview. *The Economic Journal* 98 (393), 974–1031.
- Sabater, J., Sierra, C., 2001. REGRET: A reputation Model for Gregarious Societies. In: Proceedings of the 4th Workshop on Deception, Fraud and Trust in Agent Societies.
- Sahai, A., Machiraju, V., Sayal, M., van Moorsel, A. P. A., Casati, F., 2002. Automated SLA monitoring for web services. In: DSOM. Vol. 2506 of *Lecture Notes in Computer Science*. Springer, pp. 28–41.
- Sandholm, T., 2003. Automated mechanism design: A New Application Area for Search Algorithms. In: Proceedings of the International Conference on Principles and Practice of Constraint Programming.
- Sandholm, T., Conitzer, V., Boutilier, C., 2007. Automated Design of Multistage Mechanisms. In: Proceedings of IJCAI'07.
- Sandholm, T., Ferrandon, V., 2000. Safe Exchange Planner. In: Proceedings of the International Conference on Multi-Agent Systems.
- Sandholm, T., Lesser, V., 1995. Equilibrium Analysis of the Possibilities of Unenforced Exchange in Multiagent Systems. In: Proceedings of IJCAI. pp. 694–701.
- Sandholm, T., Wang, X., 2002. (Im)possibility of Safe Exchange Mechanism Design. In: Proceedings of AAAI.

- Savage, L. J., 1971. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association* 66 (336), 783–801.
- Schillo, M., Funk, P., Rovatsos, M., 2000. Using Trust for Detecting Deceitful Agents in Artificial Societies. *Applied Artificial Intelligence* 14, 825–848.
- Schmidt, K. M., 1993. Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests. *Econometrica* 61, 325–351.
- Sen, S., Sajja, N., 2002. Robustness of Reputation-based Trust: Boolean Case. In: *Proceedings of the AAMAS*. Bologna, Italy.
- Shapiro, C., 1983. Premiums for High Quality Products as Returns to Reputations. *The Quarterly Journal of Economics* 98 (4), 659–680.
- Sherali, H., Shetty, C., 1980. *Optimization with Disjunctive Constraints*. Springer-Verlag.
- Singh, M. P., Huhns, M. N., 2005. *Service-Oriented Computing*. Wiley.
- Srivatsa, M., Xiong, L., Liu, L., 2005. TrustGuard: Countering Vulnerabilities in Reputation Management for Decentralized Networks. In: *Proceedings of the World Wide Web Conference*. Japan.
- Sztompka, P., 1999. *Trust: A Sociological Theory*. Cambridge University Press.
- Talwar, A., Jurca, R., Faltings, B., June 11–15 2007. Understanding User Behavior in Online Feedback Reporting. In: *Proceedings of the ACM Conference on Electronic Commerce (EC'07)*. San Diego, USA, pp. 134–142.
- Teacy, L., Patel, J., Jennings, N., Luck, M., 2005. Coping with Inaccurate Reputation Sources: Experimental Analysis of a Probabilistic Trust Model. In: *Proceedings of AAMAS*. Utrecht, The Netherlands.
- Teas, R., 1993. Expectations, Performance Evaluation, and Consumers' Perceptions of Quality. *Journal of Marketing* 57, 18–34.
- von Ahn, L., Dabbish, L., 2004. Labeling Images with a Computer Game. In: *Proceedings of ACM CHI*.
- Vu, L.-H., Hauswirth, M., Aberer, K., 2005. QoS-based Service Selection and Ranking with Trust and Reputation Management. In: *Proceedings of the International Conference on Cooperative Information Systems (CoopIS 2005)*.
- Walsh, K., Siner, E., 2005. Fighting Peer-to-Peer SPAM and Decoys with Object Reputation. In: *Proceedings of P2PECON*. Philadelphia, USA.
- Wang, X., Vitvar, T., Kerrigan, M., Toma, I., Dec. 2006. A QoS-aware selection model for semantic web services. In: *4th International Conference on Service Oriented Computing (ICSOC 2006)*. Chicago, USA.
- Wedekind, C., Milinski, M., 1996. Human cooperation in the simultaneous and the alternating Prisoner's Dilemma: Pavlov versus Generous Tit-for-Tat. *Proceedings of the National Academy of Science* 93, 2686–2689.
- Wehrli, S., 2005. Alles bestens, gerne wieder. Reputation und Reziprozität in Online Auktionen. Master's thesis, University of Bern.
- Whitby, A., Jøsang, A., Indulska, J., 2004. Filtering out Unfair Ratings in Bayesian Reputation Systems. In: *Proceedings of the 7th Intl. Workshop on Trust in Agent Societies*.

- White, E., October 15, 1999. Chatting a Singer Up the Pop Charts. *The Wall Street Journal*.
- Witkowski, M., Artikis, A., Pitt, J., 2001. Experiments in building Experiential Trust in a Society of Objective-Trust Based Agents. In: Falcone, R., Singh, M., Tan, Y.-H. (Eds.), *Trust in Cyber-societies*. Vol. LNAI 2246. Springer-Verlag, Berlin Heidelberg, pp. 111–132.
- Xiong, L., Liu, L., 2004. PeerTrust: Supporting Reputation-Based Trust for Peer-to-Peer Electronic Communities. *IEEE Transactions on Knowledge and Data Engineering*. Special issue on Peer to Peer Based Data Management 16 (7), 843–857.
- Xu, L., Jeusfeld, M. A., 2003. Pro-active Monitoring of Electronic Contracts. *Lecture Notes in Computer Science* 2681, 584–600.
- Yu, B., Singh, M., 2000. A Social Mechanism of Reputation Management in Electronic Communities. In: *Proceedings of the Forth International Workshop on Cooperative Information Agents*. pp. 154–165.
- Yu, B., Singh, M., 2002. An Evidential Model of Distributed Reputation Management. In: *Proceedings of the AAMAS*. Bologna, Italy.
- Yu, B., Singh, M., 2003. Detecting Deception in Reputation Management. In: *Proceedings of the AAMAS*. Melbourne, Australia.
- Zacharia, G., Moukas, A., Maes, P., 1999. Collaborative Reputation Mechanisms in Electronic Marketplaces. In: *Proceedings of the 32nd Hawaii International Conference on System Sciences (HICSS)*.
- Zeng, L., Benatallah, B., Dumas, M., Kalagnanam, J., Sheng, Q. Z., 2003. Quality driven web services composition. In: *WWW*. pp. 411–421.
- Zeng, L., Benatallah, B., Ngu, A. H. H., Dumas, M., Kalagnanam, J., Chang, H., 2004. QoS-aware middleware for web services composition. *IEEE Trans. Software Eng.* 30 (5), 311–327.
- Zohar, A., Rosenschein, J., 2006. Robust Mechanisms for Information Elicitation. In: *Proceedings of the AAAI*.



# Radu Jurca

Ecole Polytechnique Fédérale de Lausanne (EPFL)  
School of Computer and Communication Sciences  
CH - 1015 Lausanne  
Switzerland

Phone: +41 21 693 6679;  
Fax: +41 21 693 5225;  
radu.jurca@epfl.ch;  
<http://liawww.epfl.ch/People/jurca>

## Research Interests

---

My research focuses on the design of feedback and reputation mechanisms, social networks and other online systems where the information is shared by individual participants, without being verified by a trusted third party. My main goal is to create trustworthy mechanisms where rational, self-interested users participate honestly and adopt the socially desired behavior. I hope to contribute in this way to a new generation of efficient information systems, that rely as little as possible on central control and coordination.

## Education

---

- PhD**, Ecole Polytechnique Fédérale de Lausanne (EPFL), expected Oct. 2007  
*Artificial Intelligence Lab, Thesis Adviser: Prof. Boi Faltings*
- Pre-doctoral School**, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2001 – 2002  
*Computer Science, first in my class*
- MS**, Universitatea Politehnica Timisoara (Romania), 1996 – 2001  
*Computer Science, second in my class, GPA: 98.4%*

## Work Experience

---

- Ecole Polytechnique Fédérale de Lausanne**, Switzerland 2002 – 2007  
Research and Teaching Assistant
- Celeus IT**, Timisoara (RO), 2001 – 2002  
Founder and Project Manager  
*designed and implemented an Information Management System for chemical laboratories*
- Dataproducts SRL**, Timisoara (RO), 2000 – 2001  
Software Engineer  
*designed and implemented an XML based UI rendering engine for printer drivers, used by Hitachi Koki Imaging Solutions Inc. in their PostScript and PCL 5 printer drivers for Windows NT/2000*
- Romanian Debate Association**, Timisoara (RO) 1998 – 1999  
Regional Manager  
*project management, coordination, evaluation and reporting, primary accounting*

## Publications

---

### Journal papers

- Obtaining Reliable Feedback for Sanctioning Reputation Mechanisms.  
Radu Jurca and Boi Faltings  
*Journal of Artificial Intelligence Research (JAIR)*, 29, pp. 391–419, 2007.

Truthful Rewards for Adverse Selection Feedback Mechanisms.

Radu Jurca and Boi Faltings

*under review*, 2007.

Reporting Incentives and Biases in Online Review Forums.

Radu Jurca and Boi Faltings

*under review*, 2007.

### **Conference and workshop papers**

Collusion Resistant, Incentive Compatible Feedback Payments.

Radu Jurca and Boi Faltings

In *Proceedings of the ACM Conference on Electronic Commerce (EC'07)*, San Diego, USA, June 11–15, 2007.

Understanding User Behavior in Online Feedback Reporting.

Arjun Talwar, Radu Jurca, and Boi Faltings

In *Proceedings of the ACM Conference on Electronic Commerce (EC'07)*, San Diego, USA, June 11–15, 2007.

Reliable QoS Monitoring Based on Client Feedback.

Radu Jurca, Walter Binder, and Boi Faltings

In *Proceedings of the 16th International World Wide Web Conference (WWW'07)*, pages 1003 – 1011, Banff, Canada, May 8–12, 2007.

Robust Incentive-Compatible Feedback Payments.

Radu Jurca and Boi Faltings

In M. Fasli and O. Shehory, editors, *Trust, Reputation and Security: Theories and Practice*, volume LNAI 4452, pages 204 – 218. Springer-Verlag, Berlin Heidelberg, 2007.

Minimum Payments that Reward Honest Reputation Feedback.

Radu Jurca and Boi Faltings

In *Proceedings of the ACM Conference on Electronic Commerce (EC'06)*, pages 190 – 199, Ann Arbor, Michigan, USA, June 11–15, 2006.

Using CHI-Scores to Reward Honest Feedback from Repeated Interactions.

Radu Jurca and Boi Faltings

In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*, pages 1233–1240, Hakodate, Japan, May 8–12, 2006.

Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms.

Radu Jurca and Boi Faltings

In *Internet and Network Economics (WINE'05)*, volume 3828 of *LNCS*, pages 268 – 277. Springer-Verlag, 2005.

Reputation-based Service Level Agreements for Web Services.

Radu Jurca and Boi Faltings

In *Service Oriented Computing (ICSOC'05)*, volume 3826 of *LNCS*, pages 396 – 409. Springer-Verlag, 2005.

Reputation-based Pricing of P2P Services.

Radu Jurca and Boi Faltings

In *Proceedings of the Workshop on Economics of P2P Systems (P2PECON'05)*, Philadelphia, USA, 2005.

Eliminating Undesired Equilibrium Points from Incentive Compatible Reputation Mechanisms.

Radu Jurca and Boi Faltings

In *Proceedings of the Seventh International Workshop on Agent Mediated Electronic Commerce (AMEC VII 2005)*, Utrecht, The Netherlands, 2005.

“CONFESS”. Eliciting Honest Feedback without Independent Verification Authorities.

Radu Jurca and Boi Faltings

In *Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI 2004)*, New York, USA, July 19, 2004.

“CONFESS”. An Incentive Compatible Reputation Mechanism for the Online Hotel Booking Industry.

Radu Jurca and Boi Faltings

In *Proceedings of the IEEE Conference on E-Commerce (CEC'04)*, pages 205 – 212, San Diego, CA, USA, 2004.

Eliciting Truthful Feedback for Binary Reputation Mechanisms.

Radu Jurca and Boi Faltings

In *Proceedings of the International Conference on Web Intelligence (WI'04)*, pages 214 – 220, Beijing, China, 2004.

An Incentive-Compatible Reputation Mechanism.

Radu Jurca and Boi Faltings

In *Proceedings of the IEEE Conference on E-Commerce (CEC'03)*, pages 285 – 292, Newport Beach, CA, USA, 2003.

Towards Incentive-Compatible Reputation Management.

Radu Jurca and Boi Faltings

In R. Falcone, R. Barber, L. Korba, and M. Singh, editors, *Trust, Reputation and Security: Theories and Practice*, volume LNAI 2631, pages 138 – 147. Springer-Verlag, Berlin Heidelberg, 2003.

### Presentations

Designing Incentive-compatible Reputation Mechanisms. *International General Online Research Conference*, Leipzig, March 28, 2007

The Price of Truth: Minimum Payments that Make Truthful-telling Rational. *Graduate Student Association Pizza Talk*, EPFL, July 21, 2006.

Obtaining Honest Feedback from Self-interested Agents. *Workshop on Trust-based Networks and Robustness in Organizations*, ETHZ, Zurich, March 13–17, 2006.

Trust, Reputation and Incentives in Online Environments. *DIP internal workshop*, EPFL, January 27, 2005.

## Patents

---

System and Method for Monitoring Quality of Service.

Radu Jurca and Boi Faltings

US Patent Pending, 2006

## Teaching Experience

---

Intelligent Agents (Master level course) 2003, 2004, 2005, 2006

*Teaching Assistant with Professor Boi Faltings  
designed exercises, graded assignments and exam*

I have directly supervised:

- 2 EPFL Master projects; Students: *Romain Revol, Laurent Grangier*;
- 4 EPFL semester projects; Students: *Cédric Meichtry, Swalp Rastogi, Aurelien Frossard, Rui Guo*;
- 2 summer internships; Student: *Arjun Talwar, Florent Garcin*;

## Awards

---

*Dean's Scholarship*, Universitatea Politehnica Timisoara 1998, 1999, 2000  
(awarded to the best 1% students of the CS department)

*Student of the Year*, Computer Science Dept., Universitatea Politehnica Timisoara 2000  
(awarded to one student every year)

*Elected Member in the Board of the Computer Science Dept.* 1997 – 2001

---

## Service and Professional Membership

---

reviewer for:

- International Journal of Web and Grid Services (IJWGS),
- Journal of Artificial Intelligence Research (JAIR),
- Web Intelligence and Agent Systems (WIAS),
- IEEE Communication Letters,
- The International Conference on Electronic Commerce (ICEC) 2007,
- Autonomous Agents and Multiagent Systems (AAMAS) 2008,
- The World Wide Web Conference (WWW) 2008,

external reviewer for:

- The ACM Conference on Electronic Commerce (EC) 2007;

member of IEEE;