

Robust Incentive-Compatible Feedback Payments

Radu Jurca and Boi Faltings

Ecole Polytechnique Fédérale de Lausanne (EPFL)
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
{radu.jurca, boi.faltings}@epfl.ch

Abstract. Online reputation mechanisms need honest feedback to function effectively. Self interested agents report the truth only when explicit rewards offset the cost of reporting and the potential gains that can be obtained from lying. Payment schemes (monetary rewards for submitted feedback) can make truth-telling rational based on the correlation between the reports of different clients.

In this paper we use the idea of automated mechanism design to construct the best (i.e., budget minimizing) incentive-compatible payments that are also robust to some degree of private information.

1 Introduction

Online buyers increasingly resort to reputation forums for obtaining information about the products or services they intend to purchase. The testimonies of previous buyers disclose hidden, *experience-related* [13], product attributes (e.g., quality, reliability, ease of use, etc.) that can only be observed after the purchase. This previously unavailable information allows buyers to take better decisions.

Quality-based differentiation of products is also beneficial for the sellers. High quality, when recognizable by the buyers, brings higher revenues. Manufacturers can therefore optimally plan the investment in their products, such that the difference between the higher revenues of a better product, and the higher cost demanded by the improved quality, is maximized. Honest reputation feedback is thus essential for establishing an efficient market.

Human users exhibit high levels of honest behavior (and truthful sharing of information) without explicit incentives. However, in a future e-commerce environment dominated by rational agents acting to maximize their revenues, reputation mechanism designers need to make sure that sharing truthful information is in the best interest of the reporter.

Two factors make this task difficult. First, feedback reporting is usually costly. Users need to understand the rating scale, must fill in feedback forms and supervise the submission of the report. All these require the time and the conscious effort of the reporters. As feedback reporting does not bring direct benefits (the information is valuable only to subsequent buyers), rational agents are better off not to report at all.

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: e.g., doctors get gifts for recommending new drugs, authors ask their friends to write positive reviews about their latest book [6, 16]. 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 Δ to offset both the cost of reporting and the gains that could be obtained through lying. Seminal work in the mechanism design literature [5, 4] shows that side payments can be designed to create the incentive for agents to report their private opinions truthfully, a property called *incentive compatibility*. The best such payment schemes have been constructed based on “proper scoring rules” [11, 7, 2], and exploit the correlation between the observations of different buyers about the same good.

Miller, Resnick and Zeckhauser [12] adapt these results to online feedback forums. A central processing facility (the reputation mechanism) “scores” every submitted feedback by comparing it with another report (called the *reference* report) about the same good. They prove the existence of general incentive-compatible payment mechanisms where the return when reporting honestly is better by at least an arbitrary margin, Δ .

Jurca and Faltings [10] use an identical setting to apply *automated mechanism design* [3, 15]. Incentive-compatible payments are computed by solving an optimization problem with the objective of minimizing the required budget. The simplicity of specifying payments through closed-form scoring rules is sacrificed for significant gains in efficiency.

Intuitively, payment mechanisms encourage truth-telling because reporters expect to get paid according to how well their feedback improves the current predictor of the reference report. Every feedback report modifies the reputation information, which acts as a predictor for future observations. The payment received by the reporter reflects the quality of the updated predictor, tested against the reference report. Assuming that the reference report is truthful, every agent naturally has the incentive to update the current reputation such that it mirrors her subjective beliefs. Thus agents report honestly, and truth-telling is a Nash equilibrium.

One key assumption behind such mechanisms 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 maximizes the expected gain.

In this paper, we investigate feedback payment mechanisms that are incentive compatible even when reporters have some private information that is unknown to the reputation mechanism. Section 2 formally describes our setting. Section 3

describes the algorithm for computing the optimal incentive compatible payments with public prior information. Section 4 exemplifies what can go wrong if the reputation mechanism does not accurately know the beliefs of the reporters, followed by an algorithm, in Section 5, for computing incentive-compatible payments that are robust against a range of private beliefs. Finally we present future work and conclude.

2 The Setting

We consider an online market where a number of rational clients (or “agents”) purchase the same product. The quality of the product remains fixed, and defines the product’s (unknown) *type*. Let Θ be the finite set of possible types, and $\theta \in \Theta$ be a member of this set. Θ is common knowledge, and we assume that all clients share a common belief¹ regarding the prior probability $Pr[\theta]$, that the product is of type θ . $\sum_{\theta \in \Theta} Pr[\theta] = 1$.

Having purchased the product, clients perceive a noisy signal about the quality (i.e., true type) of the product. Let O^i denote the random signal observed by agent i , and let $\mathcal{S} = \{s_1, s_2, \dots, s_M\}$ denote the set of possible values for O^i . The observations of different buyers are conditionally independent, given the type of the product. Let $f(s_j|\theta) = Pr[O^i = s_j|\theta]$ be the probability that a buyer observes the signal s_j when the true type of the product is θ . $f(\cdot|\cdot)$ is assumed common knowledge, and $\sum_{j=1}^M f(s_j|\theta) = 1$ for all $\theta \in \Theta$. We assume that different types generate different probability distributions for observable signals.

A central reputation mechanism asks each client to submit feedback. Let $a^i = (a_1^i, \dots, a_M^i)$ denote the reporting strategy of agent i , such that the agent will announce $a_j^i \in \mathcal{S}$ when her observed signal is s_j . The honest reporting strategy is denoted by $\bar{a} = (s_1, \dots, s_M)$, when the agent truthfully announces her observed signal.

The reputation mechanism pays clients for submitting feedback. The payment received by client i is computed by taking into account the signal announced by i , and the signal announced by another client, $r(i)$, called the reference reporter of i . Let $\tau(a_j^i, a_k^{r(i)})$ be the payment received by agent i when she announces the signal a_j^i and $r(i)$ announces the signal $a_k^{r(i)}$. The expected payment of agent i depends on the prior belief, on her observation $O^i = s_j$, and on the reporting strategies a^i and $a^{r(i)}$:

$$V(a^i, a^{r(i)}|s_j) = E_{s_k \in \mathcal{S}}[\tau(a_j^i, a_k^{r(i)})] = \sum_{k=1}^M Pr[O^{r(i)} = s_k | O^i = s_j] \tau(a_j^i, a_k^{r(i)}); \quad (1)$$

The conditional probability distribution for the signal observed by the client $r(i)$ can be computed as:

$$Pr[s_k | s_j] = \sum_{\theta \in \Theta} f(s_k|\theta) Pr[\theta | s_j]; \quad (2)$$

¹ this assumption is relaxed in Section 4




Fig. 1. Reporting feedback. Choices and Payoffs.

where $Pr[\theta|s_j]$ is the posterior probability of the type θ given the observation s_j , as given by Bayes' Law:

$$Pr[\theta|s_j] = \frac{f(s_j|\theta)Pr[\theta]}{Pr[s_j]}; \quad Pr[s_j] = \sum_{\theta \in \Theta} f(s_j|\theta)Pr[\theta]; \quad (3)$$

All agents, as well as the reputation mechanism, compute these conditional probabilities in the same way.

Reporting feedback is expensive. Let C^i be the cost incurred by agent i for acquiring and reporting the observed signal. This cost is assumed independent of the beliefs and observations of the agent. Different agents can have different reporting costs; however, the variations are sufficiently small such that the upper bound, $C = \max_i C^i$, is finite and not "too far" from individual costs.

Agents can also obtain direct benefits from lying. Let $\Delta^i(s_j, a_j^i)$ be the external benefit agent i could obtain from reporting the signal a_j^i instead of s_j . As with the reporting cost, we assume there is an upper bound $\Delta(s_j, s_k) = \max_i \Delta^i(s_j, s_k)$ for the external benefit any agent could obtain by falsely reporting the signal s_k instead of s_j . For all signals $s_j \neq s_k \in \mathcal{S}$, $\Delta(s_j, s_j) = 0$ and $\Delta(s_j, s_k) \geq 0$.

3 Optimal Incentive Compatible Feedback Payments

Let us consider an agent i that purchases the product and observes the quality signal $O^i = s_j$. When asked by the reputation mechanism to submit feedback, the agent can choose: (a) to honestly report s_j , (b) to report another signal $a_j^i \neq s_j \in \mathcal{S}$ or (c) not to report at all. Figure 1 presents the agent's expected payoff for each of these cases, given the payment scheme $\tau(\cdot, \cdot)$ and the reporting strategy $a^{r(i)}$ of the reference reporter.

Truthful reporting is a Nash equilibrium (NEQ) if agent i finds it optimal to announce the true signal when the reference reporter also reports the truth. Formally, the honest reporting strategy \bar{a} is a NEQ if and only if:

$$\begin{aligned} V(\bar{a}, \bar{a}|s_j) &\geq V(a^*, \bar{a}|s_j) + \Delta(s_j, a_j^*); \\ V(\bar{a}, \bar{a}|s_j) &\geq C; \end{aligned}$$

for all signals $s_j \in \mathcal{S}$, and all reporting strategies $a^* \neq \bar{a}$. When the inequalities are strict, honest reporting is a strict NEQ.

For any observed signal $O^i = s_j \in \mathcal{S}$, there are $M - 1$ different dishonest reporting strategies $a^* \neq \bar{a}$ that agent i can use: i.e., $a_j^* \in \mathcal{S} \setminus \{s_j\}$. Using (1) to expand the expected payment of a client, the NEQ conditions become:

$$\begin{aligned} \sum_{k=1}^M Pr[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) &> \Delta(s_j, s_h); \quad \forall s_j \neq s_h \in \mathcal{S} \\ \sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) &> C; \quad \forall s_j \neq s_h \in \mathcal{S} \end{aligned} \tag{4}$$

Any payment scheme $\tau(\cdot, \cdot)$ satisfying the conditions in (4) is incentive-compatible. [12] proves that such schemes exist.

Given the incentive-compatible payment scheme $\tau(\cdot, \cdot)$, the expected payment to an honest reporter is:

$$W = E_{s_j \in \mathcal{S}} \left[V(\bar{a}, \bar{a}|s_j) \right] = \sum_{j=1}^M Pr[s_j] \left(\sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) \right);$$

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 1

$$\begin{aligned} \min \quad & W = \sum_{j=1}^M Pr[s_j] \left(\sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) \right) \\ \text{s.t.} \quad & \sum_{k=1}^M Pr[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) > \Delta(s_j, s_h); \quad \forall s_j \neq s_h \in \mathcal{S}; \\ & \sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) > C; \quad \forall s_j \in \mathcal{S} \\ & \tau(s_j, s_k) \geq 0; \quad \forall s_j, s_k \in \mathcal{S} \end{aligned}$$

The payment scheme $\tau(\cdot, \cdot)$ solving LP 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, let us consider a simple example.

Assume that the product purchased by the clients can be either *Good* (G) or *Bad* (B) (i.e., $\Theta = \{G, B\}$). The prior belief of the clients assigns the probabilities 0.8 and 0.2 to the product being good, respectively bad (i.e., $Pr[G] =$

0.8, $Pr[B] = 0.2$). Clients perceive a binary quality signal (either *high* or *low* quality) once they purchase the product, such that $f(h|G) = 1 - f(l|G) = 0.9$ and $f(h|B) = 1 - f(l|B) = 0.2$. The probability that the next buyer experiences a high quality product is: $Pr[h] = 1 - Pr[l] = f(h|G)Pr[G] + f(h|B)Pr[B] = 0.76$.

The conditional probability distribution $Pr[O^{r(i)}|O^i]$ for the reference report follows from Bayes' law: $Pr[h|h] = 1 - Pr[l|h] = 0.86$ and $Pr[h|l] = 1 - Pr[l|l] = 0.43$. We take the reporting cost $C = 0.01$ (i.e., 1% of the normalized cost of the product) and the external benefits from lying $\Delta(l, h) = \Delta(h, l) = 0.05$. The optimization problem LP 1 is presented in Table 1, and defines the optimal payments: $\tau(h, h) = 0.083$, $\tau(l, l) = 0.15$, $\tau(h, l) = \tau(l, h) = 0$. The expected cost for the reputation mechanism (i.e., the expected payment to one agent) is 0.07 (i.e., 7% of the price of the product).

	$\tau(h, h)$	$\tau(h, l)$	$\tau(l, h)$	$\tau(l, l)$	
min	0.65	0.11	0.10	0.14	
s.t.	0.86	0.14	-0.86	-0.14	> 0.05
	0.86	0.14			> 0.01
	-0.43	-0.57	0.43	0.57	> 0.05
			0.43	0.57	> 0.01
	≥ 0	≥ 0	≥ 0	≥ 0	
Sol.	0.083	0	0	0.15	

Table 1. Example. $Pr[G] = 0.8$, $Pr[h|h] = 0.86$ and $Pr[h|l] = 0.43$.

In [10] we show how to extend LP 1 in order to compute the optimal incentive compatible payments when:

- the available budget is fixed and the margins for truth-telling are maximized;
- the reputation mechanism can use several reference reports;
- the reputation mechanism may filter out some of the reports.

All resulting optimization problems are linear, and can be solved by polynomial² time algorithms. The resulting payments may decrease the budget required by the reputation mechanism up to one order of magnitude.

4 Honest Reporting with Unknown Beliefs

The optimal incentive-compatible payments computed in Section 3 rely on the posterior beliefs (i.e., the probabilities $Pr[s_k|s_j]$) 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 θ ,

² in the size of the payment mechanism

- the conditional probabilities, $f(s_j|\theta)$, that a product of type θ generates the signal s_j ,

using Bayes' Law as shown in the Eq. (2) and (3).

However, when the reporters have different beliefs regarding the values of the reference reports, the constraints in LP 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, and the corresponding payments from Table 1. Assume an agent whose prior belief differs only slightly from that of the reputation mechanism: e.g., $Pr[G]$, the probability that the product is *Good*, is 0.82 instead of 0.8, while other values remain the same. If the agent purchases a product of *low* quality, her private belief regarding the quality observed by the next buyer is $Pr^*[h|l] = 1 - Pr^*[l|l] = 0.45$ (instead of $Pr[h|l] = 0.43$ considered by the reputation mechanism). Simple arithmetics reveals that the agent is better off by reporting *high* instead of *low* quality: the expected gain from lying is $0.45 \cdot 0.083 + 0.55 \cdot 0 + \Delta(l, h) = 0.087$, while honest reporting brings only $0.45 \cdot 0 + 0.55 \cdot 0.15 = 0.082$.

4.1 Declaration of private information

To eliminate lying incentives, Miller et al. suggest that reporters should also declare their prior beliefs before submitting feedback [12]. 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: the agent has a prior belief which assigns probability 0.82 to the product being *Good*.

If the agent truthfully declares her prior belief, the reputation mechanism computes the optimal payments: $\tau(h, h) = 0.0841$, $\tau(l, l) = 0.1598$, $\tau(h, l) = \tau(l, h) = 0$, by solving LP 1. A truthful report following a negative experience (i.e., the agent observes and declares the signal l), is rewarded by an expected revenue equal to: $0.45 \cdot 0 + 0.55 \cdot 0.1598 = 0.0879$.

The agent can, however, declare the prior belief: $\overline{Pr}[G] = 1 - \overline{Pr}[B] = 0.11$. In this case, the payment scheme computed by the reputation mechanism will be: $\tau(h, h) = 0.2792$, $\tau(l, l) = 0.1375$, $\tau(h, l) = \tau(l, h) = 0$, and the optimal strategy for the client is to declare the signal h . The client's expected feedback payment thus becomes: $0.45 \cdot 0.2792 + 0.55 \cdot 0 + \Delta(l, h) = 0.1756 > 0.0879$, where $\Delta(l, h) = 0.05$ is the external revenue the agent can obtain by falsely declaring h instead of l .

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 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 s_h that lead to the same posterior belief $(Pr[\theta|s_j])_{\theta \in \Theta}$, can be computed by solving the following system of linear equations:

$$\overline{Pr}[\theta|s_h] = \frac{f(s_h|\theta)\overline{Pr}[\theta]}{\sum_{t \in \Theta} f(s_h|t)\overline{Pr}[t]} = \frac{f(s_j|\theta)Pr[\theta]}{\sum_{t \in \Theta} f(s_j|t)Pr[t]} = Pr[\theta|s_j]; \quad \forall \theta \in \Theta; \quad (5)$$

The system has $|\Theta|$ equations and $|\Theta|+1$ variables (i.e., the probabilities $\overline{Pr}[\theta]$ and the signal s_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:

$$\begin{aligned} \max \quad & \Delta(s_j, s_h) + \sum_{k=1}^M Pr[s_k|s_j]\tau(s_h, s_k) \\ \text{s.t.} \quad & \overline{Pr}[\theta|s_h] = Pr[\theta|s_j]; \quad \forall \theta \in \Theta; \\ & \tau(\cdot, \cdot) \text{ solves LP 1 for the prior beliefs } \overline{Pr}[\theta] \end{aligned}$$

To enforce truth-telling, Prelec [14] 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 Miler et al. in [12]. Miss-reporting 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.

5 Robust Incentive Compatible Payments

In this paper we pursue an alternative solution for dealing with unknown beliefs. We 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, we 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 $\max(-\epsilon, -Pr[\theta]) \leq \varepsilon_\theta \leq \min(\epsilon, 1 - Pr[\theta])$, $\epsilon > 0$.

Replacing the private beliefs in (2) and (3), the conditional probabilities for the reference rater's signals become:

$$Pr^*[s_k|s_j] = \frac{\sum_{\theta \in \Theta} f(s_k|\theta)f(s_j|\theta)(Pr[\theta] + \varepsilon_\theta)}{\sum_{\theta \in \Theta} f(s_j|\theta)(Pr[\theta] + \varepsilon_\theta)}; \quad (6)$$

Let $Pr_m^*[s_k|s_j]$ and $Pr_M^*[s_k|s_j]$ be the minimum, respectively the maximum values of $Pr^*[s_k|s_j]$ as the variables $(\varepsilon_\theta)_{\theta \in \Theta}$ take values within the acceptable bounds. If we modify LP 1 such that the constraints on the optimal payments are satisfied for all acceptable values of $Pr^*[s_k|s_j]$, we obtain a payment scheme that 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:

$$\sum_{k=1}^M Pr^*[s_k|s_j] (\tau(s_j, s_k) - \tau(s_h, s_k)) > \Delta(s_j, s_h); \quad (7)$$

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

$$\min_{Pr^*[s_k|s_j]} \left(\sum_{k=1}^M Pr^*[s_k|s_j] (\tau(s_j, s_k) - \tau(s_h, s_k)) \right) > \Delta(s_j, s_h); \quad (8)$$

If the probabilities $Pr^*[s_k|s_j]$ were independent³, the minimum would be given by one of the combinations of extreme values: i.e., $Pr^*[s_k|s_j]$ equal either

³ they are not, because they are connected through the same variables (ε_θ)




Fig. 2. 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 2. Average CPU time (and standard deviation) for computing the optimal payment scheme with private beliefs.

$Pr_m^*[s_k|s_j]$ or $Pr_M^*[s_k|s_j]$. Therefore, by replacing every constraint (7), with 2^M linear constraints, one for every combination of extreme values of $Pr^*[s_k|s_j]$, we impose stricter condition than (8). The optimization problem defining the payment scheme is similar to LP 1, where every constraint has been replaced by 2^M linear constraints, one for every combination of extreme values of $Pr^*[s_k|s_j]$.

We evaluated experimentally the effect of private beliefs on the expected cost of the incentive-compatible mechanism. For that purpose, we generated 2000 random problems as described in Appendix A. For each problem, we took different tolerance levels to private beliefs (i.e., $\epsilon = \{0, 0.02, 0.05, 0.07, 0.1\}$) and solved the linear optimization problem that defines the robust, incentive compatible payments. We used average hardware (e.g., Pentium Centrino 1.6MHz, 1Gb RAM, WinXP) and the CPLEX⁴ linear solver. Table 2 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 2 presents the average cost of an incentive-compatible payment scheme that tolerates private beliefs. We plot 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 ϵ , 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, we intend to define non-linear algorithms that can approach the truly optimal payments.

⁴ www.ilog.com

5.1 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 $Pr[s_k|s_j]$ that parameterize LP 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 [17].

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

$$\sum_{k=1}^M \left(Pr^*[s_k|s_j] - Pr[s_k|s_j] \right)^2 \leq \varepsilon^2; \forall s_j \in \mathcal{S}; \quad (9)$$

for some positive bound ε . The incentive-compatibility constraints must enforce that for any value of the probabilities $Pr^*[\cdot|\cdot]$, honesty gives the highest payoff. Formally,

$$\begin{aligned} \min_{Pr^*[\cdot|\cdot]} & \left(\sum_{k=1}^M Pr^*[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) \right) > \Delta(s_j, s_h); \forall s_j \neq s_h \in \mathcal{S}; \\ \text{s.t.} & \sum_{k=1}^M \left(Pr^*[s_k|s_j] - Pr[s_k|s_j] \right)^2 \leq \varepsilon^2; \end{aligned}$$

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

$$\begin{aligned} \min & \left(\sum_{k=1}^M Pr^*[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) \right) = \\ & \sum_{k=1}^M Pr[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) - \varepsilon \sqrt{\sum_{k=1}^M \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right)^2}; \end{aligned}$$

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




Fig. 3. 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. Average CPU time (and standard deviation) for computing the optimal payment scheme with general private information.

CP 1

$$\begin{aligned}
 \min \quad & W = \sum_{j=1}^M Pr[s_j] \left(\sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) \right) \\
 \text{s.t.} \quad & \sum_{k=1}^M Pr[s_k|s_j] \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right) - \varepsilon \sqrt{\sum_{k=1}^M \left(\tau(s_j, s_k) - \tau(s_h, s_k) \right)^2} > \Delta(s_j, s_h); \\
 & \sum_{k=1}^M Pr[s_k|s_j] \tau(s_j, s_k) - \varepsilon \sqrt{\sum_{k=1}^M \tau(s_j, s_k)^2} > C; \\
 & \forall s_j \neq s_h \in \mathcal{S}; \quad \tau(s_j, s_k) \geq 0; \forall s_j, s_k \in \mathcal{S}
 \end{aligned}$$

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

We evaluated experimentally the cost of general private information as reflected on the expected payment to one reporter. As in the previous section, we generated 2000 random problems (details in Appendix A) and for different levels of robustness, we solved CP 1 to obtain the robust incentive compatible payments. Table 3 presents the average CPU time required to compute the payments. As expected, the values are much smaller than those of Table 2.

Figure 3 plots the average expected payment to one agent for different number of signals, and different tolerance levels to private information. Like in Figure 2, the cost of the mechanism increases exponentially with the robustness level, ε^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 verify⁵ 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 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.

6 Discussion

We have assumed in our paper that the true quality of the product (i.e., the true type) is fixed. In real settings, however, quality does change either because the technology evolves, or because initial defects get identified and corrected. Under such circumstances, it becomes even more important to design payments that are robust to a wide range of private beliefs: incremental changes in the true quality will fall within the tolerance levels of the payment scheme and do not shift reporting incentives. However, smarter payments may be capable of modeling quality change, and factor it appropriately in reporting incentives. This remains as future work.

The reporting and honesty costs enclose a powerful framework for treating the collusion between buyers and product manufacturers. Dishonest feedback from clients usually creates advantages for product manufacturers (false positive feedback is beneficial for the product’s owner, false negative feedback benefits the competitors). The collusion between clients and providers becomes interesting when the benefit obtained by providers from a false report offsets the payment returned to the client in exchange for lying. The feedback payments presented in this paper make sure that no provider can afford to buy false reports; the collusion thus becomes irrational.

The honest reporting Nash Equilibrium is unfortunately not unique. Other lying equilibria exist, and some of them generate higher expected payoffs for reporters than the truthful one. In a previous result, [8] we show that a small number of *trusted reports* (i.e., feedback reports that are true with high probability) can eliminate (or render unattractive) lying Nash equilibria. As future work, we intend to extend the presented framework to also account for multiple equilibria.

Collusion between clients remains a problem for this class of incentive compatible mechanisms. Agents can synchronize their possibly false reports in order to increase their revenue. Choosing randomly the reference report for every submitted feedback can help eliminate small coalitions: only large coalitions are rational, such that the probability of having a reference report from the same coalition is big enough. Another safeguard against reporting coalitions is to use trusted reports. In some settings [9], a small number of trusted reports can make collusion irrational.

⁵ by checking that the payments $\tau(\cdot, \cdot)$ solve the optimization problems LP 1 or CP 1

One interesting direction for future research is to design mechanisms that can better tolerate private beliefs. As discussed in Section 4, our algorithms generate payments that increase exponentially with the range of tolerated private information. However, using a combination of different techniques (e.g., the payments also depend on declared priors, priors may be discounted as they diverge from the belief of the reputation mechanism) may result in cheaper mechanism.

7 Conclusion

Honest feedback is essential for the effectiveness of online reputation mechanisms. When feedback reporters are self-interested, explicit payments can make truthful reporting rational. Most of the existing incentive-compatible payment schemes are constructed based on proper scoring rules. Lately, computational techniques based on the idea of *automated mechanism design* have made it possible to significantly decrease the cost of incentive-compatibility by computing the best payment scheme for each context.

In the current paper we extend this line of research, by studying incentive-compatible payments that are also robust to some degree of private information. We show how the smallest amount of private information (possessed by the agents, and unknown to the reputation mechanism) can disrupt the truth-telling incentives provided by traditional payment mechanisms. As a consequence, we suggest the automated design of robust payments that are incentive-compatible for a range of beliefs. The resulting optimization problems are more complex, but can still be solved efficiently for practical settings.

References

1. A. Ben-Tal and A. Nemirovski. Robust Solutions of Uncertain Linear Programs. *Op. Research Letters*, 25:1–13, 1999.
2. R. T. Clemen. Incentive contracts and strictly proper scoring rules. *Test*, 11:167–189, 2002.
3. V. Conitzer and T. Sandholm. Complexity of mechanism design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference (UAI)*, 2002.
4. J. Crémer and R. P. McLean. Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent. *Econometrica*, 53(2):345–61, 1985.
5. C. d’Aspremont and L.-A. Gard-Varet. Incentives and Incomplete Information. *Journal of Public Economics*, 11:25–45, 1979.
6. A. Harmon. Amazon Glitch Unmasks War of Reviewers. *The New York Times*, February 14, 2004.
7. S. Johnson, J. Pratt, and R. Zeckhauser. Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case. *Econometrica*, 58:873–900, 1990.
8. R. Jurca and B. Faltings. Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms. In *Internet and Network Economics*, volume 3828 of *LNCS*, pages 268 – 277. 2005.

9. R. Jurca and B. Faltings. Reputation-based Service Level Agreements for Web Services. In *Service Oriented Computing (ICSOC - 2005)*, volume 3826 of *LNCS*, pages 396 – 409. 2005.
10. R. Jurca and B. Faltings. Minimum payments that reward honest reputation feedback. In *Proceedings of the ACM Conference on Electronic Commerce*, Ann Arbor, Michigan, USA, June 11-15 2006.
11. M. Kandori and H. Matsushima. Private observation, communication and collusion. *Econometrica*, 66(3):627–652, 1998.
12. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359 –1373, 2005.
13. A. Parasuraman, V. Zeithaml, and L. Berry. A Conceptual Model of Service Quality and Its Implications for Future Research. *Journal of Marketing*, 49:41–50, 1985.
14. D. Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.
15. T. Sandholm. Automated mechanism design: A New Application Area for Search Algorithms. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming*, 2003.
16. E. White. Chatting a Singer Up the Pop Charts. *The Wall Street Journal*, October 15, 1999.
17. A. Zohar and J. Rosenschein. Robust Mechanisms for Information Elicitation. In *Proceedings of the AAAI*, 2006.

A Generating random settings

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

$$f(s_k|\theta_j) = \begin{cases} 1 - \epsilon & \text{if } k = j; \\ \epsilon/(M - 1) & \text{if } k \neq j; \end{cases}$$

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

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.