

Computational Approaches for Studying Dietary Behaviors with Digital Traces

Présentée le 14 octobre 2022

Faculté informatique et communications
Laboratoire de science des données
Programme doctoral en informatique et communications

pour l'obtention du grade de Docteur ès Sciences

par

Kristina GLIGORIC

Acceptée sur proposition du jury

Prof. M. Jaggi, président du jury
Prof. R. West, directeur de thèse
Prof. M. De Choudhury, rapporteuse
Prof. T. Althoff, rapporteur
Prof. D. Gatica-Perez, rapporteur

Acknowledgements

The last five years were a period of intense personal and professional growth. It was wonderful to have the chance to spend the last five years doing what I love; learning, growing, teaching, mentoring, and being taught and mentored. What a privilege! I have a deep sense of gratitude towards dlab, EPFL, and many, many people who supported me on the way. I will start with Bob.

I am so grateful to have such an amazing advisor. I admire Bob's creativity, enthusiasm, attention to detail, rigor, and clarity of thought. During the last five years, Bob was exactly the advisor I needed to grow and learn. Bob also always pushed me to be more ambitious and selflessly opened the doors for me, which I am grateful for and do not take for granted.

I would like to thank my thesis committee members, Munmun, Tim, Daniel, and Martin, for all their time, inspiring discussions, and their supportive and valuable feedback.

I am fortunate that during my Ph.D. I had the opportunity to work with amazing researchers who are all extraordinarily successful, sharp, and talented but kind and supportive. First, I am grateful to Emre, Ryen, and Eric. The opportunity to work closely together, your constant input, a flurry of enthusiasm, ideas, and constructive feedback had an immeasurable impact on my Ph.D. work and shaped the content of this thesis. I am grateful to Arnaud for guiding us through this line of work with his domain knowledge and constant support and encouragement. Thanks, Microsoft Swiss Joint Research Center, for funding this project and providing valuable opportunities to get feedback and network. I would also like to thank Scarlet for all the support throughout.

Special thanks go to Ashton, a great mentor and collaborator on many fun projects throughout this journey. I remember how in the first semester of my PhD, if on any occasion I got silent and overwhelmed, Ashton would notice and ask, "*Kristina, what do you think?*". Ashton's mentoring style is something I will take away when I work with junior collaborators. I am grateful to all the collaborators from whom I learned so much by working together on various projects over the last five years, thank you, Anita, Aris, Justyna, Andrea, Jake, George, Giovanni, Marcel, Dan, Kevin, Florian, and Markus!

I am grateful to Evgeniy and Shailesh for their amazing mentorship and for having hosted me during my internship at Google; the summer in California was one of the highlights of my PhD. I am grateful to Krishna, who hosted me for a research internship as an undergrad student. This experience motivated me to pursue a PhD, and Krishna greatly impacted my career by encouraging and supporting me in the process. Thank you.

Acknowledgements

I am very grateful to all the dlab members. I am lucky to have been part of such a friendly and supportive environment. I loved our dmondays, rounds of favorite ice-creams, dSmas dinners, and Secret Santas. Thank you, Akhil, Bhargav, Marija, Martin, Valentin, Ramtin, Debjit, Alberto, Andreas, and Kiran. Manoel, thank you for all the great collaborations and everything we learned together growing as researchers. Thanks, Maxime, for being a great collaborator and friend I could always trust and rely on, for all your feedback, the chats about research, and for being the best office mate. Thank you, Candice, for all your support and patience.

I am grateful to have had an opportunity to work with an army of interns, undergrad, and master students who supported my research and from whom I learned so much. Special thanks to Robin, Irena, and Gorjan, who contributed to the data collection and analysis in the contributions described in this thesis. Collaboration with you was integral for making progress on this research. I would also like to thank Nils Rinaldi, Aurore Nembrini, and Philippe Vollichard for their help in obtaining and anonymizing the campus-wide purchase logs. I am also grateful to Jonas and Kiran for their help with data engineering, and to Digital Epidemiology Lab for making the MyFoodRepo dataset publicly available via the AICrowd Food Recognition Challenge.

I am thankful to everyone who made the last five years in Lausanne amazing. Thanks, Manoel and Jessica, for being such great friends and couple-mates throughout this adventure. All the trips, hikes, snow adventures, and nights out made it infinitely more pleasurable. Thanks, Panayiotis, Stella, and Panos, for all the great moments together, partying like Greeks do. I am also grateful to my girlfriends in Lausanne, Beatriz, Sana, and Jessica; our gatherings were a much-needed secret oasis of feminine energy and support. Thanks, Michele, David, Gael, Marija², for all the fun we had hanging out in Lausanne and around!

I am also grateful for amazing fellow PhD students from my cohort, Panos, Sena, Stefan, Cey, Okan, and Novak—you were a great support, especially in the most critical early days of this journey, and the time we spent together was always a lot of fun. Thanks to my special furry four-legged best friend from whom I learned some valuable life lessons, including “*You can never regret taking a nap*” and “*Tail up, fake it 'till you make it!*”.

I am grateful to my family, friends from Serbia, and from abroad, for their support throughout my life and education. Thank you all for always being there to spend time together whenever possible during the last five years!

Finally, all the love and gratitude I have in me goes to Tiziano, who has bravely taken on so many roles through this amazing adventure together. My partner, my best friend, my favorite person to talk to about research, and my favorite person—period. My number one fan. You ground me, but also encourage and push me to be more ambitious. I am fortunate to have you, who went through the same experience simultaneously, so we could be there for each other. You are a bigger support than I could ever deserve.

Lausanne, September 20, 2022

Abstract

Human nutrition and dietary habits shape our health, daily life, societies, the environment, and life on earth in general. However, it remains challenging to understand and attempt to change dietary behaviors using traditional methods due to measurement and causal identification challenges.

In this thesis, we contribute computational and causal approaches leveraging large-scale passively sensed digital traces to shed new light on our dietary behaviors and derive novel scientific insights. We study dietary behaviors in two types of contexts: campus-wide and worldwide. Based on digital traces capturing behaviors of tens of thousands of people on campus and millions of internet users, we develop observational study designs that enable the isolation of causal effects of implicit behavioral interventions, including forming social ties, being exposed to the food choice of others, and stay-at-home interventions. The thesis is organized into three parts.

The first part presents studies based on situated on-campus food purchase logs. In the first study, we show how, when a person acquires a new eating partner on campus, the healthiness of their food choice shifts significantly in the direction of their new eating partner's dietary patterns. In the second study, we identify purchasing mimicry, i.e., copying the food choices of others in the purchasing queue, as a specific behavioral mechanism for how similarities in dietary behaviors between individuals occur on campus.

The second part of the thesis leverages online information-seeking traces (Google search query logs). Studying worldwide dietary behaviors, we identify and describe global shifts in dietary interests during the first wave of the COVID-19 pandemic, larger and longer-lasting than the typical changes during the end-of-year holidays in Western countries.

In the third part, we critically investigate the limits to how much computational approaches can reveal about dietary behaviors in the general population. We contribute a framework for reasoning about biases of digital traces and present a case study of food consumption in Switzerland. The bias estimates derived in the case study imply that researchers should aim to establish evidence of validity before relying on social media and tracking apps—two frequently used digital traces—as proxies for true food consumption in the general population.

Acknowledgements

The novel scientific findings and methodological advances presented here contribute to the existing knowledge about human dietary behaviors and inform the design of future food systems, policies, and behavioral interventions.

Keywords: *data science, human behaviors, causal inference, diet, health, sustainability, campus, COVID-19, validity, search logs, transaction logs*

Abstract (Italian)

L'alimentazione e le abitudini alimentari modellano la nostra salute, la vita quotidiana, le società, l'ambiente e la vita sulla terra in generale. Tuttavia, rimane difficile comprendere e modificare i comportamenti alimentari utilizzando metodi tradizionali a causa delle sfide di misurazione e identificazione.

In questa tesi, utilizzando approcci computazionali e causali e sfruttando tracce digitali ottenute passivamente su larga scala, contribuiamo a gettare nuova luce sui nostri comportamenti alimentari. In particolare, studiamo i comportamenti alimentari in due tipi di contesti: a livello di campus ed a livello globale. Usando tracce digitali che catturano i comportamenti di decine di migliaia di persone nel campus e milioni di utenti su Internet, sviluppiamo studi osservazionali per isolare gli effetti causali di interventi impliciti come la formazione di legami sociali, l'esposizione alla scelta alimentare di altre persone e le restrizioni di mobilità.

La tesi è organizzata in tre parti. La prima parte presenta due studi basati sui logs di acquisto nei ristoranti del campus. Nel primo studio, mostriamo come, quando un soggetto inizia a frequentare una nuova persona durante i pasti, la sua scelta alimentare si sposta in modo significativo nella direzione delle abitudini del suo nuovo partner. Nel secondo studio, identifichiamo il mimetismo d'acquisto. Esso rappresenta la copia delle scelte alimentari degli altri, identificato come il meccanismo che causa queste somiglianze nei comportamenti.

La seconda parte della tesi si concentra sulle tracce digitali relative alla ricerca di informazioni online (log delle query di ricerca di Google). Studiando i comportamenti alimentari a livello globale, identifichiamo e descriviamo i cambiamenti negli interessi alimentari durante la prima fase della pandemia di COVID-19. Questi cambiamenti risultano più marcati e duraturi anche rispetto all'impennata tipicamente osservabile durante le festività di fine anno nei paesi occidentali.

Nella terza parte, indaghiamo criticamente le limitazioni di questi approcci computazionali applicati ai comportamenti alimentari nella popolazione generale. Forniamo un quadro per ragionare sui bias delle tracce digitali e presentiamo un caso di studio sul consumo di cibo in Svizzera. Questo studio rivela che i ricercatori dovrebbero ambire a analizzare la validità dei dati nel contesto del proprio studio prima di fare affidamento sui social media e sulle app di tracciamento—due tracce digitali utilizzate di frequente—per misurare il vero consumo di cibo nella popolazione generale.

Acknowledgements

Le nuove scoperte scientifiche e progressi metodologici qui presentati contribuiscono alle conoscenze esistenti sui comportamenti dietetici umani e informano la progettazione di futuri sistemi alimentari, politiche e interventi comportamentali.

Parole chiave: *data science, comportamenti umani, inferenza causale, dieta, salute, sostenibilità, campus, COVID-19, validity, search logs, registri delle transazioni*

Contents

Acknowledgements	i
Abstract (English)	iii
Abstract (Italian)	v
List of Figures	x
List of Tables	xiii
I Introduction and background	1
1 Introduction	3
1.1 Motivation	3
1.2 Contributions and thesis overview	7
2 Background	14
2.1 Digital trace data as a passive sensor for diets and nutrition	14
2.2 Biases of digital trace data	15
2.3 Campus-wide context: Dietary behaviors on-campus	17
2.4 Worldwide context: Population-level information seeking behaviors	21
II Studying campus-wide dietary behaviors with purchase traces	23
3 Food consumption on campus: Dataset and opportunities	25
3.1 Introduction	25
3.2 Dataset	27
3.3 Research opportunities leveraging purchase logs	36
3.4 Discussion	40
4 Social tie formation and food consumption on campus	43
4.1 Introduction	43
4.2 Materials and methods	45
4.3 Results	53
	vii

Contents

4.4 Discussion	61
5 Purchasing mimicry in food consumption on campus	65
5.1 Introduction	65
5.2 Materials and methods	69
5.3 Results	74
5.4 Discussion	82
III Studying worldwide dietary behaviors with information seeking traces	85
6 COVID-19-induced shifts in dietary interests	87
6.1 Introduction	87
6.2 Materials and methods	89
6.3 Results	95
6.4 Discussion	106
IV Validity of studying dietary behaviors with digital traces	110
7 Validity of dietary digital traces	112
7.1 Introduction	112
7.2 Materials and methods	115
7.3 Results	122
7.4 Discussion	128
V Discussion and conclusion	137
8 Discussion	139
8.1 Summary of contributions	139
8.2 Significance	141
8.3 Advantages and limitations of studying diets locally vs. globally	142
8.4 Future opportunities	144
8.5 Connections with further thesis work	149
9 Conclusion	152
A Purchasing mimicry in food consumption on campus	154
A.1 Supplementary information and robustness checks	154
A.2 Supplementary analyses	154
B COVID-19-induced shifts in dietary interests	159
B.1 Supplementary information and robustness checks	159
B.2 Supplementary analyses	160

Bibliography	209
Curriculum Vitae	210

List of Figures

1.1	A hypothetical causal diagram.	5
1.2	Outline of the thesis contributions and the corresponding chapters.	7
3.1	Sustainability challenge participants.	28
3.2	The yearly heartbeat of the campus.	29
3.3	The daily heartbeat of the campus.	30
3.4	Shop co-occurrences.	31
3.5	Food item co-occurrences.	31
3.6	The supply and demand for vegetarian meals.	32
3.7	Vegetarian meals across subpopulations.	32
3.8	Beers, energy drinks, and vending machine purchases across subpopulations.	33
3.9	Sustainability challenge subpopulation vs. entire population.	34
3.10	Exam period vs. lectures.	36
3.11	Lectures vs. breaks.	37
4.1	Distribution of purchases across item categories, across a campus.	45
4.2	Annual distribution of food purchases and annual distribution of detected onsets of social ties.	46
4.3	Study design diagrams.	48
4.4	The matched incident user study design with active comparators.	49
4.5	Daily fraction of purchases annotated as potentially unhealthy, tracked over five years.	50
4.6	Distribution (before matching) of propensity to start eating with a healthy-eating partner and the importance of most indicative features for predicting treatment assignment.	52
4.7	Histogram of the number of high confidence indicators of eating together and histogram of the partner's pre-treatment healthiness score.	53
4.8	Linear regression effect estimates.	54
4.9	Sensitivity analysis.	57
4.10	Dose-response relationship.	58
4.11	Post-treatment increase in healthiness score, stratified by pre-treatment healthiness score of focal person	59

4.12	Estimated difference-in-differences effects of co-eating onset with healthy- vs. unhealthy-eating partner.	60
5.1	Study design.	68
5.2	Directed acyclic graphs (DAGs) encoding the assumptions about the causal relationship between variables.	71
5.3	Daily histogram and the estimated risk differences.	75
5.4	The estimated risk differences, depending on the estimated status.	76
5.5	Separately for lunch, breakfast, and afternoon or evening snack, the estimated risk differences.	78
5.6	Dose-response.	80
5.7	Sensitivity analysis.	81
6.1	Global shifts in dietary interests.	94
6.2	Study design.	95
6.3	Temporal evolution of dietary interests.	97
6.4	Spearman rank correlation coefficient between mobility decrease and interest volume.	100
6.5	The short-term effect of mobility decrease on interest in accessing food.	101
6.6	The short-term effect of the shock of mobility decrease on interest in categories of food entities.	103
6.7	Number of weeks until food interest goes back to normal.	105
7.1	Perceived tastiness of tweeted food vs. actually consumed and tracked food.	113
7.2	Example images in food type “pizza”.	119
7.3	Estimation of image rank with simulated sampling of duels.	120
7.4	Comparison of food type frequency among Twitter food vs. MyFoodRepo food.	122
7.5	Histogram of crowdsourced participants’ preferences.	123
7.6	Relative rank (with respect to estimated latent scores) of median tweeted food among consumed and tracked foods.	125
7.7	Correlation matrix between estimated qualities among 2,400 competing images.	126
7.8	Percentile comparisons: for percentiles among Twitter food, the corresponding percentile among MyFoodRepo food.	127
7.9	Illustration of the biases between the true food consumption, tracked food, and tweeted food.	131
8.1	Outline of the further thesis contributions.	149
A.1	The estimated risk difference across the matched situations depending on the individuals’ status, age, and gender.	157
A.2	Separately for lunch, breakfast, and afternoon or evening snack, the estimated risk ratio for the different food item additions.	158
A.3	The histogram of the social tie strength and the risk ratio estimate within the subset of matched situations.	158

List of Figures

B.1	Mobility in 18 studied countries.	165
B.2	Pearson's correlation coefficient between mobility and interest volume	166
B.3	Short-term effects estimated with a linear model.	167
B.4	Short-term effects estimated with a constant model.	168
B.5	Example of the quadratic model fit.	169
B.6	Example of the linear model fit.	170
B.7	Example of the constant model fit.	171
B.8	Estimating the impact of the bandwidth on the fitted coefficients for constant, linear, quadratic, and cubic model.	172
B.9	Short-term effects on the share of interest, estimated with a quadratic model. .	173
B.10	Short-term effects across food categories, grouped by country.	174
B.11	Short-term effects estimated with a quadratic model, in the second wave. . . .	175
B.12	Long-term effect of mobility decrease on food interests.	176

List of Tables

4.1	Pre-treatment covariate balance.	53
4.2	Contingency table counting number of pairs of matched focal persons in each condition.	55
5.1	Summary of food addition items.	70
5.2	Contingency table.	73
5.3	Matched contingency table.	75
7.1	Tags most distinctive of MyFoodRepo food and tags most distinctive of Twitter food.	129
A.1	Gender co-eating matrix.	156
A.2	Status co-eating matrix.	156
A.3	Age co-eating matrix.	156
B.1	Summary of the 28 food entity categories.	161
B.3	Entity-level Spearman rank correlation between interest and mobility.	163
B.4	Main quadratic model: fitted coefficient alpha and R^2 statistic.	164
B.5	The impact of model order.	165

Introduction and background

Part I

1 Introduction

1.1 Motivation

The importance of human dietary behaviors can hardly be overstated. Our diets shape our health and well-being, the fabric of our culture and societies, as well as life on earth in general. A healthy diet is essential for our health [93, 131], and eating healthy protects us against chronic non-communicable diseases, such as heart disease, diabetes, or cancer. A good diet also allows people with chronic illnesses to manage their condition and avoid complications [265].

At the same time, food plays a central role in our daily activities, socialization, and social life. According to recent estimates, Americans aged fifteen or older daily spend, on average, an hour consuming foods and beverages as their primary activity [160]. Meal preparation is an expression of beauty and creativity. Across cultures, sharing a meal expresses hospitality and welcome. Beyond daily activities, throughout history and across various civilizations, food shapes and disrupts the fabric of society. For instance, there is a widely held assumption that, during the Neolithic Age, the surplus of food triggered the First Agricultural Revolution [71], i.e., a transition from a lifestyle of hunting and gathering to one of agriculture and settlement, which, in turn, provided the basis for the rise of the first civilizations. A similar transformative effect is linked to the introduction of the potato in Europe, which promoted economic development and population growth that eventually fueled the Industrial Revolution of the 19th-century [239].

Food is equally prominent in culture, art, and religion. In all major world religions, offering food is a common ritual. In the Old Testament, the concept of gaining consciousness is described through the act of biting food. In the New Testament, bread and wine are associated with the body and blood [228]. Throughout history, Western European and American paintings prominently featured meals—in particular, the foods that were aspirational to the commissioners, aesthetically pleasing, or that encoded important aspects of the cultural, religious, or political context. For instance, since antiquity, lobsters, crayfish, and crabs have appeared on the table in banquet paintings, symbolizing rebirth and resurrection [387].

Finally, beyond ourselves and the societies we build, the food we consume shapes life on earth in a most general sense. Ever since the Industrial Revolution, the human diet has implied a pervasive and ever-growing environmental degradation. Today, around 50% of our planet's ice- and desert-free land is used for agriculture, and half of all agriculturally used land is dedicated to animals [292]. If this trend persists, the comfort zone for humanity and ecosystems to thrive will have to be compromised [47]. Worldwide food consumption, a major driver of climate change, is responsible for around 26% of all human-made greenhouse gas emissions, with meat (and beef in particular) being the biggest driver [279].

To summarize, human nutrition and dietary habits shape our health, daily life, our societies, and the environment. In this context, understanding behaviors related to food has been undertaken through the lenses of various disciplines, including humanities, social sciences, natural sciences, medicine, and, most recently, computer science. However, despite the broad recognition of the crucial role of human dietary behaviors, it remains challenging to understand and attempt to change them. In what follows, we highlight three major sets of challenges related to **measurement**, **identification**, and **feasibility of interventions**.

First, **measurement challenges** relate to the fact that, traditionally, measurement—i.e., collecting data about human dietary behaviors in terms of who consumes what, when, and where—has been expensive and rare [315]. In this respect, conventional data collection methods available to researchers typically rely on surveying and self-reporting, both expensive and prone to biases [220].

Second, **identification challenges** stem from the fact that the most important questions about human dietary behaviors are inherently causal. Understanding causal relationships is crucial to making good decisions. For example, how does the consumption of red meat influence our hypertension risk? An analysis attempting to address this question might compare the prevalence of hypertension among groups with varying levels of red meat consumption. However, while it may be valuable to measure and describe the correlation between the two phenomena—i.e., red meat consumption and hypertension—it would remain unclear how much of the difference in the prevalence of hypertension can be attributed to red meat consumption specifically, as opposed to other associated dietary behaviors, such as a general preference for processed food [212], as illustrated in Figure 1.1. A person might ask: “If I cut down my weekly red meat intake by a specific amount, how much will my disease risk decrease?” or “Should I quit the occasional weekend steak, or not?” The causal and counterfactual nature of these questions is revealed through the usage of phrases “*if I...*”, “*should I...*” [271].

Indeed, nutrition is a topic elusive for a good research design [181]. Along with the difficulties of accurately assessing what people eat, it is equally complicated to tease apart true causal pathways. How can we pick out the effect of one food from the preparation method and all the other foods people eat? How do we separate the impact of the food from the factors that determine food choice? The gold standard way to estimate such effects is randomized




Figure 1.1: A hypothetical causal diagram capturing the relationship between red meat consumption frequency, hypertension, and preference for processed food. High red meat consumption frequency may cause hypertension, while a preference for processed food in general may cause both high red meat consumption frequency and hypertension.

controlled experimentation, which assigns a treatment to some people and not to others. However, in most dietary behaviors, randomization is neither feasible nor ethical [299]. One cannot intervene and change people’s consumption, at least not for prolonged periods of time. For instance, it would be unethical to prescribe to a group of people to consume unhealthy food. Additionally, nutrition research has faced a reproducibility and accuracy crisis [340]. For example, associations with cancer risk have been reported for most common food ingredients, with numerous studies describing implausibly large effects [322].

Third, **feasibility challenges** relate to the practicality of modifying human dietary behaviors. Since food behaviors are deeply personal, it is difficult to intervene and change people’s habits. Food is simultaneously an expression of culture and a means of celebration or personal comfort. Hence, it hits a special nerve when we are told we should change what and how we should eat. Modifying our dietary behavioral patterns is challenging, even when the change is meant to fight climate change or protect our health, or that of our loved ones [397]. Unfortunately, such a modification is even more challenging when it has to occur on a cross-cultural scale [247]. Despite the reports on small-scale examples of successful interventions, relevant large-scale evidence remains scarce. Overall, the problem of healthy and sustainable dietary behaviors is linked to a wide variety of factors, including culture, habits, access to affordable and nutritious food, values, social status, economics, and all aspects of agricultural systems [47].

Yet, despite the challenges outlined above, we are now at a crossroads, facing new opportunities that might allow us to understand better human behaviors related to food. **The central premise of this thesis is that novel computational approaches powered by digital traces and the new causal science can offer an un-tapped potential to understand and causally explain human dietary behaviors.** This premise is underpinned by several considerations.

First, **digitalization** has led to major improvements in **measurement** and the types of data that can be collected. Digitalization has enabled a transition from a world where behavioral data is rare and expensive to a world where behavioral data is abundant. In today’s digital

age, the behaviors of billions of individuals are recorded and stored for further analysis, as part of regular operations of digital products and services. Such data, typically a byproduct of digital platforms, is therefore referred to as digital traces. Since these digital traces are “big”, it has become possible to study rare events, detect small effects, and uncover heterogeneity among subpopulations. Digital data is “always-on”, and data streams are constantly flowing. The data is also continuously collected, thus enabling research on unexpected events and the generation of real-time measurements without delays. Another important characteristic of digital data that distinguishes it from traditional data is that the former is often “passively collected” or “nonreactive”, meaning that the act of data collection does not impact or interfere with people’s behaviors [315].

Second, recent **causal science** advances have equipped us with paradigms that enable **identification** and reasoning about causes and effects. Today, it has become possible to extract scientific insights from data not created for scientific research purposes. Causal revolution [271] allowed formalizing the notions of causality, encoding variables, expressing relationships, identifying relevant factors, estimating causal effects, and reasoning about the counterfactuals. Observational studies relying on such causal approaches allow us to approximate experiments we cannot conduct. This progress in causal approaches has been fueled by contributions from various fields, including but not limited to computer science, statistics, and econometrics. Quasi-experimental methods and structural models have led to a credibility revolution in economics and policy-making [20]. Neyman-Rubin’s causal model [305], also known as the potential outcomes framework, is widely adopted in economics, political science, and legal studies [206]. Computer science has contributed the language of causal diagrams, which serve as a bias analysis tool allowing identification and enabling precise communication [271]. Such causal approaches provide a rigorous foundation for researchers to make causal estimates from non-experimental digital traces, climbing up the ladders of causality from describing associations to considering interventions and counterfactuals.

Third, considering the challenges that our societies are facing, urgent action is needed to promote sustainability, avert the climate crisis, and improve the health of our growing populations. There is a need to provide better causal explanations of human dietary behaviors and to ensure **feasibility** of dietary interventions. To that end, **United Nations sustainable development goals** [223] provide a potential framework to guide the policies and interventions such that they are feasible and globally coordinated. Two of the seventeen goals are closely related to dietary behaviors, namely, “to take urgent action to combat climate change and its impacts” (Goal 13) and “to ensure healthy lives and promote well-being for all at all ages” (Goal 3). The former goal, as mentioned above, concerns environmental and sustainability challenges. The latter goal is linked to the world’s prevalent diseases, such as heart disease and diabetes, both of which are, to a large extent, associated with dietary habits.

The fusion of the three factors—availability of large-scale passively sensed behavioral traces, the advances in causal science, and the urgency to act to ensure sustainability and health—sets

the stage for the contributions of this thesis. In the next section, we summarize the main contributions and outline the structure of the thesis.

1.2 Contributions and thesis overview

To better understand and causally explain human dietary behaviors, we design and conduct observational studies leveraging large-scale passively sensed datasets that capture behaviors of a large number of persons for prolonged periods. We demonstrate how we can develop study designs that enable the identification of causal effects of interest and how in doing so, we enrich and refine existing knowledge about human dietary behaviors.

The thesis is organized into five parts. Part I contains the Introduction and the Background. The main contributions of the thesis are presented in Parts II, III, and IV, while Part V contains the Discussion and the Conclusion.

Parts II and III present scientific findings derived by monitoring dietary behaviors in two distinct types of contexts. In Part II, we study dietary behaviors with millions of food purchase logs, collected *in a situated campus-wide context*. In Part III, we study dietary behaviors with information seeking logs, *worldwide*, studying behaviors of millions of internet users. To derive meaningful insights from passively sensed behavioral traces, we adopt a causal inference

Part II: Studying **campus-wide** dietary behaviors with **purchase traces**

Food consumption on campus: Dataset and opportunities (Chapter 3)

Shaping the Nutritional Environment via Dietary Behavior Analysis of Food Sales Logs: Case Studies, Opportunities, and a Call to Action.
Kristina Gligorić, Robin Zbinden, Arnaud Chiolero, Emre Kiciman, Ryan White, Eric Horvitz, and Robert West.
Working paper.

Social tie formation and food consumption on campus (Chapter 4)

Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study.
Kristina Gligorić, Ryan White, Emre Kiciman, Eric Horvitz, Arnaud Chiolero, and Robert West.
ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), 2021.

Purchasing mimicry in food consumption on campus (Chapter 5)

Food Purchasing Mimicry in an On-Campus Nutritional Environment.
Kristina Gligorić, Arnaud Chiolero, Emre Kiciman, Ryan White, Eric Horvitz, and Robert West.
Working paper.

Digital traces

On-campus
purchase logs

On-campus
purchase logs

On-campus
purchase logs

Methods

Descriptive
statistical analyses

Observational study:
Matched incident
user design

Observational study:
Matched design

Part III: Studying **worldwide** dietary behaviors with **information seeking traces**

COVID-19-induced shifts in dietary interests (Chapter 6)

Population-Scale Dietary Interests During the COVID-19 Pandemic.
Kristina Gligorić, Arnaud Chiolero, Emre Kiciman, Ryan White, and Robert West.
Nature Communications, 13(1073), 2022.

Information seeking logs:
Google search queries

Observational study:
Regression
discontinuity design

Part IV: **Validity** of studying dietary behaviors with **digital traces**

Validity of dietary digital traces (Chapter 7)

Biased Bytes: On the Validity of Estimating Food Consumption from Digital Traces.
Kristina Gligorić, Irena Đorđević, and Robert West.
ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), 2022.

Social media images +
Food tracking images

Experimental study:
Estimation from
pairwise comparisons

Figure 1.2: Outline of the main thesis contributions and the corresponding chapters. For each chapter, we summarize the leveraged digital trace data and the main methods.

approach that aims to minimize the impact of biasing factors, while being transparent about the assumptions and the limitations. The variety of studied contexts (ranging from *local* to *global*) and leveraged traces (including *purchase logs*, *search logs*, and *social media posts*) demonstrates the generalizability of the proposed approaches and highlights the breadth and depth of the scientific insights that can be derived.

Specifically, in Part II, we present studies modeling dietary behaviors with anonymized food purchase transaction logs made over an 8-year period on the EPFL university campus. Studying diets in campus environments has broad implications for the campus community and the general population. In Switzerland alone, over a million people eat in cafeterias in school or at work every day, making these establishments key players for healthy nutrition and a focus of Swiss policy-makers [288]. Chapter 3 introduces the dataset and identifies research opportunities. Since social norms have received considerable attention from researchers due to their influence on diets and the potential for interventions, in Chapters 4 and 5, we isolate the causal effect of implicit naturally-occurring social interventions—social tie formation and food choices of others we eat with. In the studies monitoring behaviors in an on-campus setting, special focus is placed on isolating the causal effects by controlling for confounding factors, such as, for instance, habits before the tie formation (cf. Chapter 4) and the environmental factors (cf. Chapter 5) which can lead to biases if unaccounted for.

In Part III, we turn to a different kind of digital trace data—aggregated information seeking logs. Leveraging anonymized traces recorded as part of everyday web browsing, we can scale up the computational approaches beyond the campus environment, from campus-wide to worldwide. We study dietary behaviors globally, across a dozen of countries. Focusing on the disruption of COVID-19-induced mobility restrictions, we isolate the effect of the population-wide interventions on dietary interests, as revealed via Google search volumes (Chapter 6). As in Part II, specific focus is placed on isolating the causal effects of the interventions by controlling for confounding factors such as seasonal fluctuations and trends (cf. Chapter 6) which can skew the estimates if left unaccounted. We develop a rigorous study design relying on regression discontinuity that isolates the impact of COVID-19-induced shocks. The derived scientific insights we contribute in this study have immediate implications and can inform policy-making, the design of food systems, and behavioral interventions.

Nonetheless, there is no such a thing as a free lunch—digital traces are not without their limitations. Unlike directly collected “designed data”, e.g., data collected via surveys, digital traces are not collected for scientific research purposes. Hence, digital traces are admittedly imperfect. Most notably, they tend to be incomplete, inaccessible, non-representative, drifting in time, algorithmically confounded, error-prone, and sensitive [214, 315]. Therefore, in Part IV, we take a step back and rigorously investigate the bounds to how much digital traces can tell us about the actual phenomenon of interest: true offline behaviors of the general population and their determinants. To that end, we contribute a novel crowdsourcing framework for estimating biases and perform a case study of food consumption in Switzerland (Chapter 7).

The outline of the thesis contributions with the original publication serving as the basis for each chapter is depicted in Figure 1.2. The remainder of the thesis is structured as follows. In Chapter 2, we present the background and set the scope for the contributions. Chapters 3–7 describe the contributions. We end with a discussion (Chapter 8) and a conclusion (Chapter 9). Next, we summarize each chapter’s main scientific contributions and novelties.

1.2.1 Studying campus-wide dietary behaviors with purchase traces (Part II)

Food consumption on campus: Dataset and opportunities (Chapter 3)

Adapted from:

Shaping the Nutritional Environment via Dietary Behavior Analysis of Food Sales Logs: Case Studies, Opportunities, and a Call to Action. Kristina Gligorić, Robin Zbinden, Arnaud Chiolero, Emre Kıcıman, Ryen White, Eric Horvitz, and Robert West. *Working paper.*

Modern campuses evolved to be complete ecosystems where people spend large fractions of their time and consume foods regularly. Despite the implications of the foods consumed in such situated contexts for health, performance, and the environment, it remains challenging to provide good measurements of consumed food. We aim to bridge this gap by introducing a dataset of passively sensed purchase traces made on the EPFL university campus as part of regular operations.

In Chapter 3, we introduce the dataset, present descriptive statistical analyses, and identify opportunities leveraging the introduced dataset. For instance, descriptive analyses reveal that, on the studied campus, purchases reflecting potentially harmful dietary behaviors (alcohol, energy drinks, and vending machines) are most prevalent among students, Ph.D. students, younger subpopulations, and males. We find that the academic schedules drive food consumption on-campus, both at the yearly level (lecture season vs. exam season) and at the daily level (lectures vs. breaks).

The following two chapters address specific research questions related to social interactions on campus in detail.

Social tie formation and food consumption on campus (Chapter 4)

Adapted from [151]:

Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study. Kristina Gligorić, Ryen White, Emre Kıcıman, Eric Horvitz, Arnaud Chiolero, and Robert West. *ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), 2021 (Best Paper Honorable Mention Award).*

Chapter 1. Introduction

Social influence has long been theorized to be a key determinant of nutrition. However, it has been difficult to quantify the postulated role of social influence on nutrition using traditional methods such as surveys, due to the typically small scale and short duration of studies. To overcome these limitations, in Chapter 4, we leverage logs of food purchases made on the EPFL university campus. In a longitudinal observational study, we ask: How is a person's food choice affected by eating with someone else whose own food choice is healthy vs. unhealthy?

To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design: we identify focal users who at first do not have any regular eating partners but then start eating with a fixed partner regularly, and we match focal users into comparison pairs such that paired users are nearly identical with respect to covariates measured before acquiring the partner, where the two focal users' new eating partners diverge in the healthiness of their respective food choice. We contribute to the rich literature on social influences with the following scientific findings:

1. When a focal person acquires a new eating partner, the healthiness of the focal user's food choice shifts significantly in the direction of their new eating partner's dietary patterns. Focal persons acquiring a healthy-eating partner change their habits significantly more in the direction of healthy foods than focal persons acquiring an unhealthy-eating partner. We quantify the robustness of this finding in a sensitivity analysis, and we provide further evidence by observing a dose-response relationship between the difference in exposures and the difference in effects.
2. Further identifying foods whose purchase frequency is impacted significantly by the eating partner's healthiness of food choice reveals that focal persons who start eating with healthy-eating partners show an increase in the purchase of coffee and lunch meals, items generally purchased in large numbers, with the strongest effect. On the other hand, items purchased at higher rates by the matched counterparts loosely form a cluster of potentially unhealthy items that should not be eaten in large quantities.

Beyond the main scientific findings, this contribution also demonstrates the utility of passively sensed food purchase logs for deriving insights, with the potential of informing the design of public health interventions and food offerings, especially on university campuses.

Purchasing mimicry in food consumption on campus (Chapter 5)

Adapted from:

Food Purchasing Mimicry in an On-Campus Nutritional Environment. Kristina Gligorić, Arnaud Chiolero, Emre Kiciman, Ryen White, Eric Horvitz, and Robert West. Working paper.

The following chapter, Chapter 5, studies social influence on campus in more detail, at the meal-level.

We leverage the sequential queue nature of cafeterias and the fact that we can monitor many persons in many situations. We consider a large number of situations where a dyad of partner (early decision-maker, i.e., the person who goes first in the purchasing queue) and focal user (late decision-maker, i.e., the person who goes second in the purchasing queue) are adjacent, and both make a purchase. Identifying the partner's impact on the focal user, we find evidence in favor of a specific behavioral mechanism for how dietary similarities between individuals occur—purchasing mimicry. We contribute the following scientific findings:

1. We find significant mimicry of partners' purchases affecting all food types and diminishing once the ordering of the purchasing queue is randomized.
2. Purchasing mimicry is present across age, gender, and status subpopulations on campus, but strongest for students and youngest persons.
3. We find that the mimicry diminishes as the proximity in the purchasing queue (measured in seconds between transactions) decreases, thus exhibiting a dose-response relationship where the smallest distances in the purchasing queue correspond to the largest effect estimates.

The results of this study elucidate the behavioral mechanism of purchasing mimicry and have further implications for understanding dietary behaviors among on-campus subpopulations. Our findings imply that modifying the availability of supplementary food items sold on campuses (such as fruits and desserts) can be leveraged to increase or reduce the intake of specific foods and nutrients.

1.2.2 Studying worldwide dietary behaviors with information seeking traces (Part III)

COVID-19-induced shifts in dietary interests (Chapter 6)

Adapted from [146]:

Population-Scale Dietary Interests During the COVID-19 Pandemic. Kristina Gligorić, Arnaud Chiolero, Emre Kiciman, Ryen White, and Robert West. Nature Communications, 13(1073), 2022.

Moving from a campus-wide to a worldwide context, in Chapter 6, we study diets with information-seeking traces. Motivated by the fact that the SARS-CoV-2 virus has altered people's lives around the world, in Chapter 6, we study population-wide shifts in dietary interests in 18 countries in 2020, as revealed through time series of Google search volumes. Methodologically, drawing meaningful conclusions from the longitudinal Google search volume time series is challenging due to the presence of trends and seasonalities.

Chapter 1. Introduction

We overcome these hurdles via quasi-experimental time-series analyses, isolating the effect of the 2020 discontinuity in mobility patterns on food interests and going beyond simple correlations by accounting for 2019 baseline trends.

Our main contributions are the following findings:

1. During the first wave of the COVID-19 pandemic, there was an overall surge in food interest, larger and longer-lasting than the surge during typical end-of-year holidays in Western countries.
2. The shock of decreased mobility manifested as a drastic increase in interest in consuming food at home and a corresponding decrease in consuming food outside of home.
3. The increased food interest is not uniform across types of food. The largest (up to three-fold) increases occurred for calorie-dense carbohydrate-based foods such as pastries, bakery products, bread, and pies.

The identified shifts in interests, many of which persisted for months and some of which continued past our observation period, represent a potential danger for public health. Thus, the observed shifts in dietary interests have the potential to globally affect food consumption and health outcomes. This contribution can inform governmental and organizational decisions regarding measures to mitigate the effects of the COVID-19 pandemic on diet and nutrition worldwide.

1.2.3 Validity of studying dietary behaviors with digital traces (Part IV)

Validity of dietary digital traces (Chapter 7)

Adapted from [148]:

Biased Bytes: On the Validity of Estimating Food Consumption from Digital Traces. Kristina Gligorić, Irena Đorđević, and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW), 2022.

Lastly, digital traces such as purchase logs and information-seeking logs studied in the previous chapters lend themselves as potential proxies for food consumption. However, the validity of such methods has not been established. It remains unclear to what extent digital traces reflect real food consumption.

In Chapter 7, we aim to bridge this gap by quantifying the link between dietary behaviors as captured via social media and via food-tracking applications, at a population scale. We investigate the link between online and offline dietary behaviors by studying food images, as

measured via two platforms: Twitter and MyFoodRepo food tracking app, in Switzerland. The main contributions are the following:

1. A novel crowdsourcing framework for estimating biases.
2. A case study of food consumption in Switzerland applying the bias estimation framework. Controlling for location, period, and food types, we contrast an extensive set of tweeted food images with images of consumed and tracked food. We find that:
 - Food type distributions among social media foods vs. among consumed and tracked foods diverge; e.g., bread is 2.5 times more frequent among consumed and tracked foods than on Twitter, whereas cake is 12 times more frequent on Twitter.
 - Studying food types individually, Twitter still provides a biased view of food consumption as measured via food tracking. Tweeted food is, on average, across food types, perceived as more caloric, less healthy, less likely to have been consumed at home, and tastier compared to actually consumed and tracked food. For example, on average, across food types, a median-tasty Twitter dish is among the top 26% tastiest MyFoodRepo dishes, and a median-caloric Twitter dish is among the top 34% most caloric MyFoodRepo dishes.

While social media traces can be a reasonable proxy of tracked consumption for certain food types, we find that, overall, food shared on social media and consumed and tracked food significantly diverge from each other. The fact that there is a divergence between food consumption measured via the two platforms—food tracking and social media—implies that at least one of the two is not a faithful representation of the true food consumption in the general Swiss population. We conclude that researchers should thus be attentive and try to establish evidence of validity before using digital traces as a proxy for true food consumption in the general population.

Measuring biases in digital traces is the first step towards correcting them and drawing valid conclusions despite their presence. Through a case study of the Twitter and MyFoodRepo platforms in Switzerland, in Chapter 7, we contribute grounding and first insights regarding the validity of dietary traces.

In the next chapter, we provide background and scope for the outlined contributions.

2 Background

In this chapter, we provide a broad background and set the scope for the contributions of the thesis. We start by discussing related work monitoring dietary behaviors with digital traces (Section 2.1), focusing on the opportunities and recent advancements. We then provide background about the known limits of such approaches (Section 2.2) by reviewing related work studying biases of digital traces. In the remainder of the chapter, we provide background for the two contexts in which we study diets: campus-wide dietary behaviors and the corresponding traces (Section 2.3), and worldwide information-seeking behaviors and the corresponding traces (Section 2.4).

2.1 Digital trace data as a passive sensor for diets and nutrition

2.1.1 Estimating food consumption from passively-sensed digital traces

Analyzing nutrition behaviors via digital traces has been an active area of research. Previous work leveraged various kinds of digital trace data including search engine logs [376, 392], purchase logs [11, 12, 51, 101, 187, 200], online recipes [297, 314, 360, 362, 378, 379], reviewing platforms and websites [65, 82, 163, 388], crowdsourcing platforms [105, 177], geo-location signals [309], wastewater-based signals [103], and ubiquitous devices and wearables [5, 34]. Another active area of research has been focused on improving methods for monitoring food consumption relying on mobile phones [77, 78] and wearable devices to recognize when eating activities occur [356].

2.1.2 Estimating food consumption from social media digital traces

In particular, given the popularity and prevalence of food-related content, studying diets through social media posts has been an active research area. Instagram [133, 261, 275, 332] and Twitter [3, 88, 241, 242, 244, 395] have emerged as particularly promising platforms. Researchers have studied specific dietary issues and harmful behaviors. In work with important implications for the health and well-being of vulnerable populations, researchers studied re-

ports of eating disorders [56, 57, 86, 268], dietary choices, nutritional challenges in food deserts (places with poor access to healthy and affordable food) [88], and obesity patterns in online behaviors [243]. Related work has also studied eating disorder support online communities, quantifying and predicting disease severity and recovery [57, 86].

2.1.3 Leveraging user-generated food content for machine learning tasks and NLP

Researchers have been utilizing user-generated food content to train and develop machine learning models. Current AI applications that use online food images include mining food photos to perform segmentation [262], recognize food [17, 45, 313, 403], learn food and recipe embeddings [316], and perform calorie [257] and nutrient estimation [128].

Previous work has also studied the language of food—the impact of the dish vocabulary [252] and its historical origins and socio-cultural dimensions [126, 368]—through the language of food advertisements [126], menus [364, 369], and reviews [196, 282].

Knowledge gaps and relation to thesis work. Re-purposing inadvertent and indirect proxies and data sources for studying diets is a promising research direction. The contributions of this thesis demonstrate the potential of various kinds of digital traces—including purchase logs, search logs, and social media posts—for deriving novel insights about human behaviors.

2.2 Biases of digital trace data

While large-scale digital traces are promising for monitoring and modeling nutrition, little is known about how these passively sensed behavioral signals can be used for understanding the factors that govern food consumption. For instance, although social media has emerged as a rich data source, food shared or discussed on Instagram and Twitter might not be representative of food that people actually consume. Researchers have compared population-scale statistics extracted from tweet text with public health statistics regarding the prevalence of obesity and diabetes [3, 243, 314]. However, the content of posted images and the foods themselves have not been contrasted with actually consumed foods to date.

Similarly, while food shared on social media might not be representative of consumed food, more distant proxies make it even harder to determine validity. For example, do recipe searches on search engines correspond to eating the food? Does reading an online recipe imply that the food was prepared and consumed? It is unknown to what extent such proxies imply food consumption, and it is unclear whether studies of food consumption via such digital traces truly measure the quantities intended to be measured. Moreover, if the food people consume is systematically different from food shared online, models trained and evaluated on online datasets might not generalize to real-world scenarios. In the next section, we provide background regarding the biases linked to studying diets with digital traces.

2.2.1 Construct validity

The goal of measurements using behavioral trace data is to extract meaning from raw data that most often was not collected with the extraction of scientific insight in mind. Data-driven research has thus been criticized for asking questions that appear to be opportunistically answerable with the data at hand, overlooking different types of biases [156]. Lazer et al. [214] argue that digital traces need to be linked to known constructs before we can use the data to answer scientific questions. Thus, the key challenge of studying digital data is determining whether measurements accurately capture the construct one would ideally want to examine. For example, if one is measuring physical activity based on mobile phone location traces, how consequential is the omission of stationary activities such as treadmill or yoga [214]? If one is tracking influenza with Web search logs of symptoms, how consequential are searches from persons not experiencing any symptoms [215]?

The mismatch between the theoretical understanding of a concept and its operationalization, known as the issue of construct validity, can have harmful consequences [380]. In particular, when data that allows for measurement (e.g., arrest records) does not properly match the actual social construct that the measurement is intended to capture (e.g., a criminal act), measurements can replicate, mask, or exacerbate existing social issues [76].

Related work has thus aimed to establish the validity of studying human behaviors with Web and social media traces. Example studies include studying the validity of screening depression [202], location traces [176], inferring political approval [326], sentiment analysis [273], or using Twitter's APIs [254]. De Choudhury et al. [87] have studied seeking and sharing health information online by comparing search engines and social media. Researchers have also studied decisions around whether to post content online [2], political, racial, and gender biases in Web systems [161, 207], and how Web systems influence offline user behavior [14].

2.2.2 Error frameworks

Further related work includes studies that issue calls to carefully scrutinize the use of social media data against biases and provide practical implications to aid researchers in performing their data-driven studies. Sen et al. [327] proposed a total error framework for digital traces of human behavior on online platforms, Olteanu et al. [264] identified a variety of challenges in the practices of social media use for research, and Hofman et al. [174] advocated for measuring the extent to which causal estimates made in one domain transfer to another domain.

Knowledge gaps and relation to thesis work. Whereas related work [264, 306, 327] aims to put the biases into a unified framework cutting through different domains, in Chapter 7, we specifically establish the bounds to the validity of estimating food consumption from digital traces.

2.3 Campus-wide context: Dietary behaviors on-campus

Next, we provide a broad background regarding on-campus dietary behaviors and we position our contributions described in Chapters 3-5 relative to gaps in the previous work.

2.3.1 Food consumption in campus environments

A rich body of work examined the determinants of food consumption in campus environments. The dominant factors impacting consumption include price, value for money, healthfulness, and taste [37, 304]. Factors that represent barriers to healthy eating are time constraints, snacking, high-calorie foods, stress, high prices of healthy food, and easy access to junk food. In contrast, enablers to healthy behavior are food knowledge and education, meal planning, involvement in food preparation, and being physically active [338]. In another study by Deliens et al. [94], beyond individual factors, individuals reported being influenced by their social networks, physical environment (e.g., availability and accessibility, appeal and prices of food products), and macro environment (including impacts of media and advertising). Furthermore, the relationships between determinants and university students' eating behavior are moderated by university characteristics, such as residency, student societies, university lifestyle, and exams [94, 95].

Previous work focused specifically on the role of the academic schedule and exams, demonstrating how cafeteria snack purchases become less healthy with each passing week of the semester, implying an increased demand for unhealthy foods as the college semester progresses, and in particular, at the very end of the semester [386]. Lastly, large-scale passively sensed signals have been harnessed in university campus environments to measure determinants of well-being [383] and performance [384], outside of nutrition [26, 227, 260, 324]. Recent studies point towards the feasibility and the potential of leveraging behavioral traces for campus-centric applications [84, 253, 312, 350].

On-campus dietary interventions that aim to improve food availability, accessibility, prices, and promotions through policies received considerable attention from researchers [96, 259, 303, 304]. Potentially effective interventions include reduced pricing and price manipulation [351], increased availability [134, 209, 302] and variety of fresh, seasonal, local, and healthy foods [175, 233, 396], optimized product placement of healthier items [393], positional interventions [50, 135], and providing more nutrition information [114, 116, 232]. Nutrition education is another potentially powerful tool in health campaigns to promote healthy eating patterns on campus [401].

Related studies performed interventions that targeted dietary behaviors, for instance, by altering product placement and naming to increase store sales of healthy foods [366, 367, 375]. Recent field experiments [23, 365] exploring the impact of renaming vegetarian dishes on menus found that such linguistic nudges have the potential to encourage individual action. Removing terms such as “meat-free”, which highlights the lack of meat in the dish, and

replacing them with words such as “fieldgrown” or “garden” increased the proportion of people choosing the target vegetarian dishes.

Knowledge gaps and relation to thesis work. Understanding all the determinants of food consumption and intervention effectiveness is challenging since the key factors, such as availability and attitudes, change over time, interact, and do not necessarily generalize across different campuses. There is a need for unified frameworks to better monitor and understand food consumption in campus environments worldwide.

2.3.2 Harmful dietary behaviors in campus environments

Next, we provide more background on the specific potentially harmful dietary behaviors described in Chapter 3: energy drinks, alcohol, and vending machine food. Emerging evidence has linked energy drink consumption with several negative health consequences, such as risk-seeking behaviors, poor mental health, adverse cardiovascular effects, and metabolic, renal, or dental conditions [328]. Nonetheless, previous work studying energy drink use in university students [52, 343] found that energy drink consumption is a common practice among university students. In particular, tailored health promotion strategies and interventions are needed to address misconceptions about energy drinks and alcohol mixing.

Research has established elevated levels of alcohol consumption among young adults. Research suggests that students today drink more, with increasing emphasis on binge drinking and drunkenness than among earlier generations [85]. It remains what factors contribute to alcohol consumption among university students, and unclear how to approach this issue [125].

Previous studies found that products sold in university vending machines tend to be nutritionally poor [152]. Vending machines typically market and sell less healthy food and beverages to university students. There is a need for healthier vending machines in a university setting [373], especially given the strong correlation between the availability of vending machine items and the corresponding purchases [152].

Knowledge gaps and relation to thesis work. However, despite the importance of these potentially harmful dietary behaviors, precise and generalizable estimates are lacking. For instance, how many energy drinks do students consume globally in university environments? What is the effect of energy drink consumption on academic performance, and how do academic schedules impact the consumption of energy drinks? In Chapter 3, we perform statistical analyses of on-campus purchase logs to shine a light on such consequential behaviors and reveal the axes along which they vary.

2.3.3 Social influence and diet on campus

Social influence on dietary habits is an active area of research [172, 334]. Food consumption is influenced by eating with others [170], and the food choices of others, including people one

does not know, have been observed to influence food choices, even when not consciously recognized [67, 294]. Previous research aimed to understand the governing psychological mechanisms, including the seeking of dish uniformity driven by the goal of regret minimization or the seeking of dish variety driven by self-presentation [22, 89, 256].

Although the underlying mechanisms are not fully understood, uniformity seeking is observed across a range of studies. For example, it is observed that the quantity dimension is used to communicate gender identity, and the food-type dimension to ingratiate the co-eater's preferences by matching the other's presumed choice, following gender-based stereotypes about food [55]. Such social norms, including the influence of peers, have tremendous potential for understanding dietary patterns and designing public health interventions [75, 250, 293, 295].

A large fraction of the transactions recorded in our dataset were made by students, i.e., adolescents and young adults. Focusing on similar age groups, social influence in dietary habits has been examined in the context of school children [35, 119, 269, 318] and adolescents [91, 92, 346], who are theorized to be most susceptible to social pressures. In particular, effects of peer influence have been observed in children's and adolescents' diets and activity patterns [25, 317].

Systematic reviews of social network analyses of young people's eating behaviors and body weight reveal consistent evidence that school friends are significantly similar in terms of their body mass index. Friends with the highest body mass index appear to be most similar [124]. Prior work further reveals that the family context is essential when implementing healthy eating interventions, as parents, not friends, are the most prominent influencers of adolescents' healthy eating [197, 272].

Previous work has particularly been focused on unhealthy behaviors and their contagious effects, observing that obesity [66], overeating [237], fast food [357], high-fat [117, 168], and alcohol and snack consumption [267, 399] are contagious. In fact, the strongest evidence of social influence in food choices has been found for unhealthy behaviors (e.g., snack foods) [38, 80]. Beyond food consumption, peer influence and social norms have been observed to play a role in unhealthy weight-control behaviors among adolescent girls: self-induced vomiting, laxatives, diet pills, and fasting were all shown to be contagious among adolescent girls [109]. Rich literature tackles the problem of unhealthy behaviors through interventions to promote healthy dietary habits, and physical activity [121], losing weight [189], reducing the risk of chronic illnesses [143], and reducing food waste [290].

The issue of social influences in diets is a controversial one as there is a heated debate about whether unhealthy behaviors are indeed contagious, or whether the observed similarities should instead be attributed to homophily, i.e., people's tendency to form ties with others who are similar to oneself, to begin with. Disentangling social influence from homophily poses a fundamental challenge. Without strong assumptions about the structure of ties or the ability

to measure confounding factors, homophily and contagion are generically confounded (i.e., the effect of social influence cannot be identified) [21, 329, 330].

Knowledge gaps and relation to thesis work. In Chapters 4 and 5, we monitor social influences outside of experimental setups. Having access to a multi-year history of all transactions made on a large campus allows us to observe behavioral changes for longer time periods and in a more fine-grained way, by measuring a wide set of purchasing behaviors that occur in the real world. Our work attempts to minimize the effect of confounding variables in previously infeasible ways. Based on the rich transaction data, we measure a set of relevant confounding variables and carefully control for them in our quasi-experimental setups.

2.3.4 Further implications and societal challenges linked to on-campus food consumption

Lastly, beyond sustainability, health, and social influences, which are in the focus of this thesis, on-campus food consumption reflects other challenges and circumstances faced by individuals who are part of the campus community. We consider two other prominent difficulties in the existing literature about food on-campus: *economic factors* and *skills*.

Food consumption in university environments cannot be considered in isolation from the broader environmental factors of the society that the university is part of. The members of the university communities face circumstances that the broader society faces, and these factors, in turn, affect food consumption. Previous work has thus examined food insecurity among students on campus [54, 258, 270]. Work aiming to understand food insecurity among college students found that obtaining a degree, securing a better job, and improving their living standards were priorities that outweighed hunger concerns among students [167]. University students' risk of food insecurity is partly attributed to inadequate income support and the price of available options [179].

Similarly, food consumption is affected by social context and the level of cooking skills. In a related study, university students' food intake was characterized as particularly unhealthy among students who left their parents' homes and became responsible for their food [33]. The extent to which students successfully take on the role of self-catering depends on the student's competencies and skills acquired before independent living, living situation, and, most importantly, the student's ability to create dietary habits, including regular grocery shopping and cooking [36]. Therefore, adequate food skills may improve diet quality [398].

2.4 Worldwide context: Population-level information seeking behaviors

Finally, we provide a broad background for Chapter 6. We start by reviewing related work focusing on the online information-seeking behaviors during COVID-19 and work that measured changes in user behavior during the pandemic.

2.4.1 The COVID-19 infodemic

COVID-19 was the first event of its magnitude to take place in the era of social media and user-generated content. The important role that online platforms have in today's society has prompted researchers to study patterns of virality, information seeking, and information sharing in social media during the pandemic.

Kouzy et al. [205] studied the spread of COVID-19-related misinformation on Twitter by analyzing a sample of tweets, collected in February 2020, with trending hashtags and keywords related to the disease, finding a high prevalence of mis- or unverifiable information. Other studies that followed have found similar and complementary results. Yang et al. [404] also found a high prevalence of low-credibility information, which is disproportionately spread by bots. Ahmed et al. [9] identified and analyzed the drivers behind one of the main COVID-19-related conspiracies, which postulates the spread of the infection to be related to the 5G standard for cellular networks. Analyzing data containing conspiracy-related hashtags, they found that a handful of users were driving the conspiratory content and that many of those using the hashtag were denouncing the conspiracy theory.

Knowledge gaps and relation to thesis work. Whereas these studies detected and measured content that may impact behaviors and beliefs (e.g., conspiracy theories), in Chapter 6, we focus on the inverse question: *Can we uncover changes in behavior via digital traces?* We argue that both these directions are important to the overarching goal of understanding the social dynamics of the spread of the virus [97], since to understand the impact of misinformation, one must be capable of measuring people's behaviors, needs, interests, and concerns.

2.4.2 The Web in times of COVID-19

Recent work has also more broadly characterized how COVID-19 has altered people's online behavior, and how digital traces can be used to understand the pandemic better. A first and broader theme in this direction concerns how COVID-19 has increased Internet traffic. Feldmann et al. [115] showed that, due to COVID-19-induced lockdowns, Internet traffic of residential users increased by 15–20%. Traffic increases were observed in applications that people use when at home, such as Web conferencing, VPN, gaming, and messaging. Results in the same direction were also found in survey-based studies analyzing Internet time [74], and by a smaller-scale study measuring the increase in the stress on a campus network [112].

Second, and more related to the contributions of the thesis, are works leveraging digital traces to understand the impact of COVID-19 on mental health, economy, society, and human needs [1, 157, 363]. We highlight two recent papers based on search data. Lin et al. [219] used Google search data on COVID-19-specific keywords to predict the speed of the spread of the disease. They found, e.g., that searches for “*wash hands*” are correlated with a lower spreading speed of the disease. Suh et al. [348] measured changes in human needs using Bing search logs, finding that, for a variety of different “need categories”, there was an elevated increase in search activity, and that subcategories related to the most basic needs received the largest boost.

In further related studies, researchers used Wikipedia pageviews to monitor and forecast diseases at a global scale [138, 171] and to study anxiety and information seeking about infectious diseases, such as influenza [238], H1N1 [355], and Zika [359]. Our work extends rich literature studying behavioral changes of web users during unexpected events, crises, and catastrophes [132, 409].

Knowledge gaps and relation to thesis work. The outlined related work is complementary to the contributions of this thesis presented in Chapter 6. Signals from multiple distinct digital traces such as network usage, social media activity, and search logs may allow stakeholders to paint a more comprehensive picture of how the pandemic has impacted society. Resonating with existing work, our contributions investigate the integral role played by the Web in times of crisis and its usefulness in understanding how these events impacted the behavior of its users.

Studying campus-wide dietary behaviors with purchase traces

Part II

3 Food consumption on campus: Dataset and opportunities

3.1 Introduction

The concept of a modern campus evolved to be more than a collection of buildings and grounds that belong to an institution. A modern campus is a complete eco-system. In such a situated context, people spend significant parts of their time working, educating, or being educated, but also socializing, learning, doing sports, or entertaining themselves. On campuses, people also consume food, regularly and globally.

The food consumed on campuses has broad implications for the people on campus and the general population. First, the food we eat impacts our health. In the specific case of university campus environments, during on-campus time, young adults transition into adulthood [337], which often results in poor dietary behaviors and weight gain [398]. Therefore, students are at risk of establishing unhealthy habits during their first years in college [28, 127, 371] and developing negative health consequences including diabetes, heart disease, and hypertension [4, 286]. Furthermore, in universities, through the impact on health, food consumption is linked to academic performance [289]. A literature review of the food intake of university students [33] found that university students tend to exhibit unhealthy eating behaviors, such as a high intake of fast foods, snacks, and sweets. Prevalent and potentially harmful dietary habits include consuming energy drinks, alcohol, and vending machine foods.

Second, the food consumed on campus impacts the environment through its carbon footprint and generated waste [342, 361]. The food supply chain corresponds to over a quarter of all the anthropogenic greenhouse gas emissions [278]. Reducing this quantity and taking urgent action to combat climate change and its impacts to mitigate the effects of climate change is a UN sustainable development goal [73]. Given the need to actively address the challenges of climate change, university leaders have a growing interest in reducing their campuses' environmental impact. Key actions and initiatives relate to energy, buildings, water, waste, transportation, grounds, air and climate, and food [16].

Chapter 3. Food consumption on campus: Dataset and opportunities

To be able to take action towards the goals of ensuring on-campus health and sustainability, it is important to know, to begin with, what foods are consumed on campus and in what context. How sustainable and healthy are the dietary behaviors on campus? How do they vary across subpopulations? How are they affected by the regular campus operations?

However, answering these questions faces major challenges. Traditional population-level methods such as surveys face limitations, such as under-reporting [43, 137] and cannot capture temporal dynamics accurately. Hence, much about fundamental campus dietary behaviors with implications for health and the environment—such as meat vs. meat-free meal consumption and consumption of caffeinated or alcoholic drinks—remains unknown. New measurement and monitoring methods are needed to overcome these challenges and collect dynamic, fine-grained, and generalizable data about the diet of campus populations.

In this chapter, we introduce an anonymized dataset of passively collected logs of food purchases made on the EPFL university campus and perform descriptive statistical analyses of food-purchasing behaviors. In particular, we describe how passively sensed food purchase logs typically collected as part of regular business operations can be used to monitor and model on-campus food consumption.

The descriptive statistical analyses are motivated by existing knowledge about human dietary behaviors, analyzing dimensions along which behaviors vary, such as time, space, and demographics, and the known issues that campus food consumption faces, including ensuring health and sustainability. In particular, we describe statistical analyses of transaction logs that can help the stakeholders have a detailed understanding of food consumption on campus. The dimensions along which we perform descriptive data analyses are outlined below.

Monitoring and modeling regularities in campus food consumption. What spatio-temporal patterns and regularities characterize food purchasing behavior on the studied campus? We describe temporal (Section 3.2.4), spatial (Section 3.2.4), food choice patterns and regularities (Section 3.2.4), and how they vary across subpopulations.

Sustainability and health. How are the specific dietary behaviors with implications for health and sustainability reflected in the passively sensed purchase data? Given the sustainability issues and prevalence of potentially unhealthy behaviors, we describe behaviors important for sustainability (vegetarian meals, Section 3.2.5) and health (energy drinks, alcohol, vending machines Section 3.2.5). We describe the prevalence of specific dietary habits in subpopulations since understanding heterogeneity is an important step toward modifying unhealthy dietary habits.

Chapter outline. In this chapter, we introduce the purchase logs dataset and present descriptive statistical analyses of the dataset (Section 3.2). We then formulate a concrete set of research questions (Section 3.3) that illustrate the breadth of insights that can be sourced from the purchase logs. We outline a research agenda showcasing research questions that can be addressed through quantitative analyses, alongside a description of methodological

challenges. We aim to make a case for re-purposing such data, which are often available by default and can serve as a valuable source of information to harness in university campus environments to measure nutrition. The following two chapters, Chapters 4 and 5, address two specific research questions in detail by leveraging the introduced dataset.

The present chapter introduces a dataset collected on a university campus. However, the insights, opportunities, and challenges refer to corporate, medical, industrial, and other, more or less closed environments. The notion of campus used throughout refers to such different types of environments.

3.2 Dataset

3.2.1 Anonymized purchase logs

We introduce an anonymized dataset of food purchases made on the EPFL university campus. The data spans eight years, from 2010 to 2018, and contains about 38 million transactions, of which about 18 million were made with a badge that allows linking to an anonymized person's ID. The data includes 38.7k users, who, on median, are observed for a time period spanning 578 days and make 188 transactions. The transaction data is passively collected—it is available by default and not collected specifically for this research [26, 227, 260, 324].

Each transaction is attributed with the time it took place, information about the sale location (shop, restaurant, vending machine, or café), the cash register where the transaction took place, the purchased items, their quantity, and price. Items are associated with unstructured textual descriptions (e.g., “coffee”, “croissant”, “Coca-Cola can”). The unstructured textual descriptions were additionally mapped to categorical labels such as “sandwich” or “dessert” by a research assistant who labeled the 500 most frequently purchased items, which account for 95.4% of the total volume of item purchases observed in the dataset. In the data collection and analysis process, we worked directly with campus food-providing administration and transaction system managers who exported and anonymized the data, to understand the information about the food items and restaurants encoded in the dataset.

On campus, transactions can be executed using cash, credit card, or rechargeable identifying badge. Students, staff, and visitors receive an identifying badge which enables them to access campus facilities and services and settle transactions directly with this card without using cash or a credit card. In case the transaction was executed with the identifying badge, it is additionally attributed with the ID of the person who has made the transaction. In this chapter, the descriptive analyses are based on the period from 2012 to 2018 with menu data available (31M transactions, 16.6M with person ID).




Figure 3.1: Sustainability challenge participants: demographics. In (a), the histogram of sustainability challenge participants’ status (on the left) and gender (on the right). In (b), the histogram of sustainability challenge participants’ age by status.

3.2.2 Sustainability challenge participants and the demographic information

We also introduce a smaller-size enriched dataset gathered during a three-week campus-wide sustainability challenge in November 2018, during which 1,031 consenting and volunteering participants formed 278 teams to compete in taking sustainable actions (e.g., taking the stairs instead of the elevator, or consuming a vegetarian meal). For this subset of individuals, we leverage demographic information: gender (illustrated in Figure 3.1a: 584 females, 447 males), status at the campus (illustrated in Figure 3.1a: 724 students, 280 staff members, 27 “other”), and birth year (illustrated in Figure 3.1b: average 1991, median 1994, Q1 1988, Q3 1998).

The sustainability challenge participants are represented with a consistent ID in the transaction logs. Therefore, we enrich the transaction logs with sustainability challenge participant information. In total, 0.6M transactions are attributed with the gender, age, and status of the person who executed it. When describing this subset of transactions, we will refer to them as sustainability challenge participants’ purchases, to keep in mind that an analysis is based on a subset of the full set of transactions.

3.2.3 Ethical considerations

Nutrition is a potentially sensitive personal behavior. To protect user privacy, the log data used here was accessed exclusively by EPFL personnel involved in this project, and stored and processed exclusively on EPFL servers. The data was obtained with approval from EPFL’s Data Protection Officer and was anonymized before it was made available to the researchers for analysis. Finally, we note that our work leveraging purchase logs was conducted retroactively on data that had been collected passively in order to support campus operations. Thus, our analysis did not influence users in any way. In what follows, we present a set of descriptive statistical analyses of the introduced dataset.




Figure 3.2: The yearly heartbeat of the campus. Across the week of the year (on the x-axis), the number of executed transactions (on the y-axis), across the years. Background color marks the yearly academic schedule.

3.2.4 Descriptive analyses: Spatio-temporal and choice regularities in food purchases

Temporal regularities: The heartbeat of the campus. Having introduced the dataset, we start descriptive analyses by characterizing temporal regularities in on-campus food purchasing.

We first describe yearly regularities in food purchases. During a year, the academic calendar dictates life on campus. There are three important periods ¹ in the academic calendar of the studied campus:

1. The **fall** and **spring semesters**, when the lectures take place. We will refer to these periods as *semesters*.
2. The **winter** and **summer exam sessions**, when the exams take place.
3. The **winter** and **summer breaks**, the holidays for students.

In Figure 3.2, we monitor the number of transactions peaks during the spring and fall semesters when students are on campus and attending lectures. The number of transactions drops during the exam sessions and reaches the minimum during the winter and summer breaks when the students are away. Another drop occurs during the spring semester, for the week of Easter break, which does not occur on the same date across years. Note that in 2017 and 2018, the number of transactions during the semester decreased (dashed lines), likely due to opening of shops close to campus that do not support purchases with the identifying badge.

We next focus on the daily regularities (Figure 3.3). The rhythm of food purchases during the course of a day is the following: the transactions peak in the morning (breakfast 7 -11.30h),

¹Note that there is also one additional week of Easter holidays during the spring semester, which occurs at varying times over the years.




Figure 3.3: The daily heartbeat of the campus. The histogram of the transactions by the time of the day.

during the time of lunch (11.30 - 14.30), and the afternoon or evening snack time (14.30 - 20.30). Note that at each sharp hour, spikes in the purchasing volume occur.

Spatial regularities: Migration patterns between shops. Since human behaviors are known to exhibit spatial regularity and predictability [81], after temporal, we describe spatial regularities in the purchasing behaviors. We are interested in understanding the relationship and regularities in visits to shops. To that end, we perform an association rule analysis. Across all individuals, we monitor the transactions done by an individual during one week and the shops where the transactions occurred in. We consider the list of the shops where a given individual has executed transactions during a week. We then apply the Apriori algorithm [7], an algorithm for the discovery of association rules between shops. Association rules describe regularities between items in transaction data. For example, a rule $\{X\} \rightarrow \{Y\}$ found in the transaction logs would indicate that if a person visited shop X , they are likely to also visit shop Y during the same week. In Figure 3.4, the graph depicts the confidence of the association rules found using the Apriori algorithm, defined as the percentage of all transactions satisfying X that also satisfy Y .

This approach lets us monitor shop migration patterns (Figure 3.4). For instance, a thick arrow from node *Shop 10* to node *Shop 4* means that there is a high probability that an individual who went to *Shop 10* also went to *Shop 4* during the same week (shop names are anonymized). Furthermore, we find that the distribution of the edges is linked to the geographic locations of the shops. For instance, *Shop 4* seems to be the central place in the graph as there are many arrows with high confidence coming to it, and this cafeteria is indeed at the geographical center of the campus. Moreover, *Shop 8*, *Shop 11*, *Shop 1*, *Shop 9* (the campus bar) and *Shop 2* are near-by cafeterias on the campus, frequently visited by students, and we observe that they form a cluster in the shop co-occurrence graph. The remaining nodes are connected in the co-occurrence graph and have close locations on campus, except *Shop 10*, which is a disconnected cafeteria in this graph and not near other clusters of shops on the campus.




Figure 3.4: Shop co-occurrences. The directed weighted graph representing association rules between different shops (shop names are anonymized). Nodes represent shops, and the edges represent association rules. The edge thickness is proportional to the confidence of the rule.

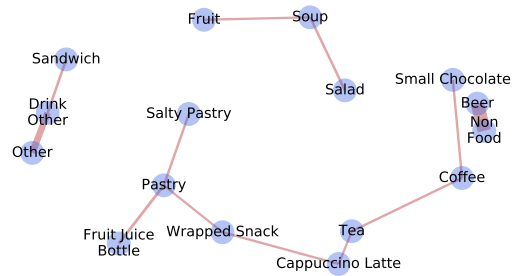


Figure 3.5: Food item co-occurrences. The weighted graph representing association rules between different food categories. Nodes represent food categories, and the edges represent rules. The edge thickness is proportional to the lift of the rule.

Food-choice regularities: Co-purchasing patterns. Next, after spatial, we describe regularities in food choice. We aim to describe the regularities in purchases of specific foods. To that end, we perform association analysis, as explained above. We consider purchases composed of food items belonging to different categories, i.e., for each purchase, a list of the food categories a person has purchased.

In Figure 3.5, we show the lifts of the association rules found using the Apriori algorithm, defined as the ratio of how frequently the rule appears in the dataset to that expected if X and Y were independent. Note that the lift is a symmetric measure; therefore, the graph is not directed. We obtain a graph where neighbor nodes are food categories that tend to be purchased together, e.g., *Salty Pastry* and *Pastry*. *Soup*, *fruit*, and *salad* are often co-purchased. Note that the *Non-Food* category contains only products linked to beer, such as glasses on deposit at the on-campus bar, explaining its high lift with the *beer* category.

To summarize, we demonstrate that food-choice regularities such as the described co-purchasing patterns can be discovered using association analysis on the anonymized purchase logs. In what follows, we focus on characterizing specific behaviors important for sustainability and health on campus.

3.2.5 Descriptive analyses: Specific dietary behaviors

In this subsection, we turn to describing specific dietary behaviors relevant for sustainability and health on campus.

Food sustainability on campus: Vegetarian meals. We first analyze vegetarian meal purchases since the consumption and production of meat, and red meat, in particular, has a negative




Figure 3.6: The supply and demand for vegetarian meals. In (a), across years (on the x-axis), the number of available vegetarian meals (on the y-axis). In (b), across years (on the x-axis), the volume of purchased vegetarian meals (on the y-axis), as the fraction of purchased vegetarian meals out of all the purchased meals, averaged across individuals. The evolution is shown separately for all individuals and individuals of different statuses. Error-bars mark bootstrapped 95% confidence intervals.




Figure 3.7: Vegetarian meals across subpopulations. By status, age, and gender strata (on the x-axis), vegetarian purchases (on the y-axis), as the fraction of purchased meals that are vegetarian, among all purchased meals, averaged across individuals. Error-bars mark bootstrapped 95% confidence intervals.




Figure 3.8: Beers, energy drinks, and vending machine purchases across subpopulations. By status, age, and gender strata (on the x-axis), the fraction of all transactions (on the y-axis) that contain beer, energy drink, or a vending machine product, averaged across individuals. Error-bars mark bootstrapped 95% confidence intervals. Note the varying y-axis.

impact on the environment [400]. We consider the frequency at which vegetarian meals are purchased out of all purchased meals. We use a vegetarian tag which indicates if the product is vegetarian or not.

We find that the supply of vegetarian meals was monotonically increasing from 5 in 2012 to 17 in 2018 (Figure 3.6a). Additionally, the demand for vegetarian meals increased too (Figure 3.6b), i.e., it overall doubled within the studied period (fraction up from 4% in 2012 to 8% in 2018), and within the subpopulations by status. The biggest increase is observed for Ph.D. students, reaching 18.7% in 2018. In other words, both the number of proposed vegetarian menus and the purchases of vegetarian meals increased, across all statuses.

On average, across individuals, we find that 5.3% of purchased meals are vegetarian (9.9% of meals are vegetarian among sustainability challenge subpopulation). In Figure 3.7, we examine the fraction of purchased meals that are vegetarian separately across individuals in the sustainability challenge dataset, by the status. Vegetarian meals are the most popular among Ph.D. students (11.2%), followed by students (8.3%), staff (5.9%), and other statuses (2.4%). Depending on age, vegetarian meals are the most popular among 21 to 30-year-olds, and more popular among females, than among males (8.5% vs. 7.2%).




Figure 3.9: Sustainability challenge subpopulation vs. entire population. Separately for sustainability challenge subpopulation and the entire population (on the x-axis), the fraction of all transactions (on the y-axis) that contain a vegetarian meal, beer, energy drink, or a vending machine product, averaged across individuals. Error-bars mark bootstrapped 95% confidence intervals. Note the varying y-axis.

Potentially harmful dietary behaviors on campus. We next analyze purchases of products with potentially harmful effects on health:

1. **Beers:** Beer is offered at the on-campus bar. The consumption of alcoholic beverages in excessive amounts is not recommended [339].
2. **Energy drinks:** The consumption of energy drinks has been reported in association with adverse health effects [325].
3. **Vending machine items:** Food products available in vending machines often have a high amount of sugar, and vending machines tend to be nutritionally poor [152].

In Figure 3.8, we observe the fraction of transactions including such products, out of all purchased products, across subpopulations of status, gender, and age. We focus separately on all purchases and sustainability challenge participants' purchases.

Regarding beers, we find that, on average, across individuals, 2.9% of transactions contain a beer (4.1% among sustainability challenge subpopulation). Beer purchases are the most prevalent among “other” statuses (e.g., interns and visitors), students, Ph.D. students, 21 to 30-year-olds, and males (3.7% males vs. 1.9% females).

On average, across individuals, 0.15% of transactions contain an energy drink (0.06% among sustainability challenge subpopulation). Energy drinks are the most prevalent among students and Ph.D. students, 26-30 year-olds, and males (0.1% males vs. 0.041% females).

Lastly, monitoring vending machine purchases, on average, across individuals, we find that 6.2% of transactions contain a vending machine item (4.2% among sustainability challenge subpopulation). Vending machine items are the most prevalent among students and 21 to 25-year-olds.

Overall, we find that purchases reflecting potentially harmful dietary behaviors are relatively prevalent, especially vending machine purchases (overall, 6.2% or 1 in 16 transactions a person makes contains a vending machine item). Students, Ph.D. students, younger subpopulations, and males are the most susceptible to potentially harmful dietary behaviors.

Figure 3.9 illustrates purchases among sustainability challenge subpopulation vs. among the entire population. Besides beer purchases which are more frequent among sustainability challenge participants (likely due to their student status), the sustainability challenge participants tend to execute fewer potentially harmful transactions (fewer energy drinks and products at vending machines and buy more vegetarian meals, both statistically significant with 95% CI). The discrepancy between the general population and the sustainability challenge subpopulation can be explained by both participants self-selecting to participate in the sustainability challenge and participants improving their behaviors due to the participation.

3.2.6 Descriptive analyses: Impact of the academic schedule

Exam session. We described how *the number of transactions* varies during the year and we observed that the purchasing volume variations coincide with the exam sessions (Figure 3.2). Now, we turn to investigate the impact of the exam sessions on *the type of purchased products*. Does purchasing behavior change during the exam sessions?

To answer this question, we compare the differences in purchases between semesters and exam sessions. Monitoring weeks of the entire studied period, in Figure 3.10, we find that, during the exam weeks, compared to semester weeks, there is a significant increase in the relative frequency of purchases of energy drinks (+22.4%), coffee (+9.9%), and a decrease in the relative purchasing frequency of beer (-20.0%), pizza (-15.3%), drinks (-6.0%), and fruit (-5.0%). For instance, the fraction of purchased energy drinks, which has the greatest change between exam weeks and semester weeks, peaks during the spring semester exams (peak occurs on the 24th week of the year when 0.5% or 1 in 200 purchased items is an energy drink) and during the fall semester exams (peak occurs on the last week of the year when 1% or 1 in 100 purchased items is an energy drink).

Overall, we conclude that the purchasing behavior shifts significantly during the exam session, with notable increases in purchasing frequency of energy drinks and coffee, and a corresponding decrease in the purchasing frequency of beers, pizzas, drinks, and fruit.

Hourly 15-minute break. In Figure 3.3, we observed the overall rhythm of food purchases during the course of the day, with transaction volume peaking in the morning, during the time of lunch, and in the afternoon. Additionally, we observed that spikes in purchasing volume occur at each sharp hour. During the fall and spring semesters, each class begins after the first quarter of each hour (e.g., at 8:15, 9:15, ...) and ends at the end of the hour (e.g., 9:00, 10:00, ...), with a break of 15 minutes between two classes. Next, we investigate the temporal patterns within an hour.

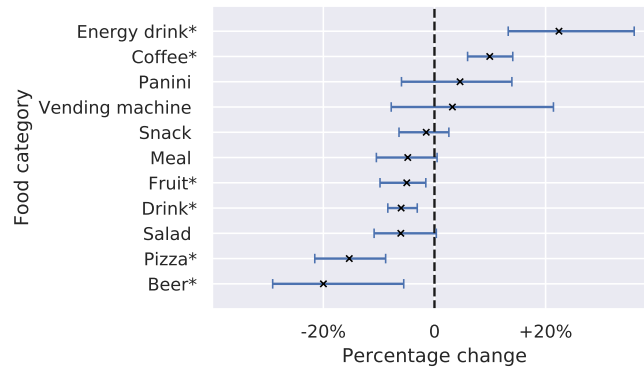


Figure 3.10: Exam period vs. lectures. For different food categories (on the y-axis), the percentage change (on the x-axis) in the fraction of all purchases that contain the food item during exam weeks, compared to lecture weeks. Error-bars mark bootstrapped 95% confidence intervals obtained by resampling respective weeks. Stars mark food categories with percentage change significantly different from zero.

The impact of the fifteen-minute break can be observed in Figure 3.11. We observe different behavior depending on the individual's status and whether the transactions are made during the semesters or not. During the spring and fall semesters, students' transactions peak at the 8th minute in the hour, during the 15-minute break, and consequently drop (Figure 3.11, top). Transactions executed by the staff members do not exhibit such a pattern, as the hourly break does not impact staff members, and staff members potentially target hours when students do not crowd the shops. Ph.D. students are in between staff and students, as one would expect. Furthermore, we find that 44% of the coffees purchased by students during the course semesters are purchased during the 15-minute break.

These differences between staff members and students disappear during the exam sessions and breaks when there are no fifteen-minute breaks (Figure 3.11, bottom), implying that hourly patterns are indeed linked with the academic calendar since students tend to take advantage of the 15-minute break to buy drink or food.

Based on these analyses, we conclude that, within an hour, the 15-minute break impacts food consumption on campus, particularly among students, and during the semesters.

3.3 Research opportunities leveraging purchase logs

Based on the insights from the descriptive statistical analyses, we now derive a set of five specific research questions that can be addressed solely with anonymized passively sensed purchase logs, to illustrate the breadth and depth of potential contributions. For each question, we outline the associated methodological challenges. The identified opportunities may serve as a call to action for researchers and stakeholders aiming to leverage passively sensed behavioral traces globally, across institutions and campuses. The following two chapters,

3.3 Research opportunities leveraging purchase logs



Figure 3.11: Lectures vs. breaks. By the minutes of the hour (on the x-axis), the fraction of executed transactions during the minute, separately during the semesters (top) vs. during the exams sessions and breaks (bottom), and separately by status. Error-bars mark bootstrapped 95% confidence intervals.

Chapters 4 and 5, address in detail a specific question outlined below, by leveraging the introduced dataset.

Impact of food availability

Motivation. Our descriptive analyses reveal that vending machine purchases are relatively frequent (on average, across individuals, 6.2% of transactions contain a vending machine item, cf. Section 3.2.5). One hypothesis is that vending machine usage relates to the (lack of) availability of other, healthier food sources. It is necessary to understand the factors driving vending machine purchases and further investigate how vending machines can be placed and stocked for an optimal trade-off between availability and healthy behavior. Similarly, we observed that both the number of proposed vegetarian menus and the purchases of vegetarian

Chapter 3. Food consumption on campus: Dataset and opportunities

meals increased (Section 3.2.5). However, it remains unclear whether the supply is driving the demand or the other way around. Specifically, the question emerges:

RQ: To what extent is food choice determined by the availability of options?

Challenges. Food availability is correlated with biasing factors that impact purchasing behaviors, such as trends and seasonalities, the academic schedule, and the period during the year. For instance, in summer, fewer foods are available since students are on a break, but food consumption patterns also change due to temperature changes. To isolate the effect of availability, it might be necessary to exploit haphazard situations such as vending machines breaking and being removed due to chance, or shops haphazardly running out of vegetarian meals (i.e., natural experiments). Furthermore, it is necessary to account for how food is offered (in a vending machine, over the counter, in separate queues per meal option, or in a single queue). Understanding the impacts of food availability and offer on purchases will allow designing food offer and shop layouts that encourage healthy diets and managing waste.

The price of laziness

Motivation. Related to the question of availability is the question of geographical distance to options. Studying the spatial patterns and analyzing co-occurrence patterns between shops, we found evidence that, within the same week, people tend to visit nearby shops (Figure 3.4). Choosing the closest shop can be a potentially dangerous habit if the close-by cafeteria happens to serve unhealthy food. How readily do users sacrifice healthy options for geographic distance? And vice versa, how far away do unhealthy options need to be moved before people shift to nearby healthy options? We ask:

RQ: To what extent is food choice determined by geographical proximity to options? How much work would people be willing to put in to reach healthier options?

Challenges. To answer these questions, an open challenge is inferring the default location of individuals. Since individuals are anonymized, it is challenging to link their physical location (e.g., office number or most frequently visited lecture hall locations) with the locations where transactions are executed. Inference of the location solely based on the transaction logs might be feasible but require careful evaluation.

Unaffiliated providers on campus or in its vicinity

Motivation. When options on campus do not meet their needs, people might turn to off-campus options. Studying the temporal patterns, we observed a period when the number of transactions during the semester decreased, likely due to opening shops close to campus that

do not support purchases with the identifying badge (Section 3.2.4), where the food purchases partially spilled over. This insight begs the question:

RQ: How does the opening of unaffiliated premises near campus affect on-campus meal consumption?

Challenges. Answering this question is challenging since data from unaffiliated providers on campus or in its vicinity is typically unavailable. The estimation, therefore, needs to be performed solely based on the incomplete on-campus purchase logs by carefully contrasting otherwise comparable periods with the modified availability of nearby shops.

Prediction of harmful dietary patterns

Motivation. We found that transactions containing alcoholic drinks and energy drinks are prevalent. On average, across individuals, 9% of transactions contain a beer, and 0.15% of transactions contain an energy drink (rising to 1% during the exam session, cf. Section 3.2.5). It remains unclear what mechanisms drive unhealthy alcohol drinking and energy drink consumption. Identifying the onset of such habits would enable prompt mitigation of adverse effects on well-being and health. Specifically, we ask:

RQ: Can the onset of on-campus alcohol drinking and energy drink consumption be predicted?

Challenges. Identifying the moments of increased risk of starting potentially harmful dietary habits requires careful crafting of contextualized indicators, such as academic timetables, social interactions, and analysis of purchasing routine forming and breaking, which may not generalize across subpopulations, and across campuses.

Social interactions

Motivation. Finally, the descriptive analyses revealed important differences between persons depending on status on campus, age, and gender. Overall, purchases reflecting potentially harmful dietary behaviors are most prevalent among students, Ph.D. students, younger subpopulations, and males (Section 3.2.5). However, beyond individual's features, social relations play an important role in dietary behavior as individuals on campus interact and share meals, beverages, and snacks. It remains unclear:

RQ: How does eating with others impact on-campus food choice?

Challenges. Measuring how dietary behaviors are affected by the behaviors of others is challenging due to numerous confounding factors and the interplay between selection biases

and social influences [330]. Identification of the inter-personal influences and determining whether social relations have heterogeneous effects on healthy vs. unhealthy foods requires careful causal analyses. The following two chapters, Chapters 4 and 5, address two specific questions related to social interactions by leveraging the introduced dataset.

3.4 Discussion

3.4.1 Summary of the descriptive analyses and their implications

The food offered and consumed on university campuses can significantly impact the environment and health, motivation, and academic performance of students and staff. In this chapter, we introduce a novel purchase dataset collected via identifying badges and present descriptive statistical analyses of food consumption on-campus. In what follows, we summarize the descriptive statistical analyses and discuss their implications for policies and decision-making aiming to improve the food offer and consumption on-campus.

Regularities and co-occurrences. Descriptive statistical analyses reveal regularities in how persons on-campus transition between cafeterias during a week (Figure 3.4) and how food items belonging to different food categories are combined (Figure 3.5). Since the transitions between cafeterias are governed by geographical proximity (Section 3.2.4), our analyses imply that the decisions about offering at cafeterias should not be taken in isolation, but should account for such spatial migration patterns. Similarly (cf. Section 3.2.4), the offer should take into account the frequent item pairings because it is not enough to consider foods in isolation (e.g., a sandwich can have good nutritional properties, but if often purchased together with a beverage with high sugar content, promoting it might not be optimal).

Sustainability. Characterizing specific dietary behaviors related to sustainability and focusing on meat consumption on campus (Section 3.2.5), we observe an increase in both the supply and the demand of vegetarian meals between 2012 and 2018 across all statuses (students, Ph.D. students, employees in Figure 3.6b). However, there are important differences between subpopulations, with vegetarian meals being, on average, most popular among Ph.D. students, people between 21 and 30 years of age, and females (Figure 3.7). Our findings highlight the key importance of sustainability campaigns geared toward the right subpopulations and individuals who do not already purchase vegetarian meals frequently. For sustainability interventions, it could be most effective to target staff, males, and the oldest and the youngest subpopulations on the studied campus.

Health. In Section 3.2.5, characterizing sales of products with potentially harmful effects on health, we observe that students tend to consume more unhealthy products compared to Ph.D. students and employees (Figure 3.8). Gender also plays a role, with females purchasing fewer unhealthy products than men (Figure 3.8). Similarly, these insights could help stakeholders design interventions targeted toward the right subsets of individuals. Health interventions on

the studied campus could be most effective if targeting males, younger subpopulations, and students.

The heartbeat of the campus. We find that academic schedules drive food consumption on campus, both at the yearly level (lecture season vs. exam season) and the daily level (lectures vs. breaks). On the yearly level, exam sessions are associated with surges in the consumption of coffee and energy drinks (Figure 3.10). On a daily level, the 15-minute break between lectures drives food consumption, particularly among students and during lecture weeks. For example, 44% of the coffees purchased by students during lecture weeks are purchased during the 15-minute break. These regularities in both yearly and daily academic schedules should be taken into account in efforts to promote sustainable and healthy habits by modifying offerings of items such as energy drinks or coffee. Furthermore, the insights regarding the daily patterns and variations are potentially useful for optimizing logistics and staffing decisions. The fact that staff visit less during the break (Figure 3.11) indicates that there might be overcrowding. Analyses of purchase logs can detect these issues and inform measures taken to improve the experiences around food on campus for everyone.

Exam session. Exams are associated with increases in purchases of potentially unhealthy products, likely due to stress and performance desires. A possible approach towards reducing energy drink consumption could be, for instance, to explore offering fewer energy drinks and, instead, consider proposing potentially healthier alternative energizing products to students, such as tea. By capturing nearly all on-campus food consumption, the purchase log analysis approach complements survey-based methodologies, which likely under-report [43, 137] stigmatized consumption of unhealthy items. For policymakers, it should be a priority to mitigate the adverse effects of the exams. Making exam sessions a better experience for students by encouraging socialization and organizing social events might be a promising direction, as there is evidence that socialization is reduced during exams since beer and pizza purchases decrease, cf. Figure 3.10. These insights highlight the need to do more to promote student well-being as stress and anxiety levels are elevated among university and college students [158, 383].

To summarize, the descriptive statistical analyses of transaction logs imply that such approaches leveraging passively sensed anonymized data should be an essential component in efforts to monitor the evolution of food consumption on campus while being aware of complex spatio-temporal variations and differences between subpopulations. The analyses that rely solely on passively sensed purchase logs have an untapped potential to support stakeholders in their efforts to understand the evolution of food consumption on campus. We end the chapter by considering the limitations of the introduced dataset that should be kept in mind when interpreting the findings.

3.4.2 Dataset limitations

Passively sensed digital traces analyzed here make new kinds of analyses feasible and have advantages compared to the traditional methods. However, the presented analyses are not without their limitations. Digital traces are not collected for scientific research purposes, creating methodological challenges [214, 315]. Therefore, there is a trade-off between the advantages of digital traces and the associated challenges. In what follows, we highlight some of the most prominent limitations of analyses of food consumption relying on passively sensed purchase logs.

Incompleteness. The unstructured food item labels and information derived from them are incomplete. Additionally, cash transactions cannot be mapped to individuals. Similarly, the vegetarian tag was deduced from the name and type of the product and is, therefore, not necessarily always correct.

Construct validity. The log data does not directly capture food consumption but provides indirect proxies via purchasing. It is not guaranteed that the purchased items were consumed. Conversely, other items that were not purchased might have been consumed on campus (e.g., a soda brought from home). Individuals can borrow badges from others to make a transaction for someone else or pay with cash. In other words, we study purchasing as visible in the transaction logs, which might not perfectly mirror the actual complete food consumption on campus.

External validity. Future work should determine to what extent behaviors measured on campus reflect other settings off-campus, e.g., food consumption. People who primarily eat home-cooked food might present a skewed representation of their diet in the studied cafeteria purchase logs, and the fact that individuals might consume food prepared at home introduces unobserved variables into our analyses.

Representativeness. The participants of the sustainability challenge are a biased subset of all the individuals from the overall dataset (due to self-selection). Therefore, differences between persons based on demographics might not generalize to all individuals.

Causality. Beyond the descriptive statistical analyses, more careful controls are necessary to identify the true role of demographics since gender, age, and status are correlated. A potential approach may involve matching on the demographic covariates while keeping other characteristics fixed, which is not possible here, given the small number of individuals. While we describe the purchasing changes coinciding with the exam session at the population level, more careful controls are necessary to identify the causal impact of the exam on purchasing habits of individuals.

4 Social tie formation and food consumption on campus

4.1 Introduction

In the previous chapter introducing the purchase logs dataset, one of the identified knowledge gaps (cf. Section 3.3) relates to the question of social interactions. How does eating with others impact on-campus food choice? In designing interventions and policies that promote healthier diets, such situational food norms, including social influences, play a prominent role, as they are theorized to have a powerful effect on food intake [75, 170, 172, 250, 334].

However, despite the postulated importance of social factors, measuring how dietary behaviors are affected by the behaviors of others remains challenging. On the one hand, experimental studies to date have been limited to observing people in small-scale scenarios with a short duration [294, 295]. On the other hand, observational studies have relied on survey-based methods [66], employing questionnaires [399] and personal food journals [77, 78], which are costly to organize and prone to biases [46].

Furthermore, making causal inferences regarding the influence of social ties on food intake faces the challenge of numerous confounding factors. Although similarities in diet and eating behaviors among persons connected via social links (e.g., friends, family, and peers) have been observed in a number of experimental and survey-based studies [119, 162, 227, 269, 318, 346], it is not clear whether the similarity in the consumption patterns arises from social influence, or if confounding factors, such as self-selection in tie formation (homophily) and environmental influences, can explain the similarity [21, 204, 329, 330]. In real-world settings, it remains challenging to measure and disentangle properties that are relevant in the context of food consumption, such as attributes of the individuals and of the environment (e.g., food options available in different locations and settings). Researchers have only recently been addressing this gap by studying social media and other digital traces of human behavior in the context of food consumption [3, 56, 268].

In order to shed new light on the influence of social factors on food choice, we harness the novel data source: logs of 38 million on-campus food purchases introduced in the previous

chapter. The large scale and long duration of the data enable studies with greater statistical power, compared to prior setups, and allows for reducing the influence of confounding factors, and thus for identifying the causal effect of social influence, by carefully selecting a suitable subset of individuals whose food choice behaviors are monitored throughout the individuals' long-term transaction histories.

Based on this dataset, we design a longitudinal observational study to address the question of how a person's food choice is affected by eating with someone else whose own food choice is healthy vs. unhealthy. To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design, where we identify focal users who at first do not have any regular eating partners but then start eating with a fixed partner regularly, and we match focal users into comparison pairs such that paired focal users are nearly identical with respect to covariates measured before acquiring the eating partner, but the two focal users' new eating partners diverge in the healthiness of their respective food choice.

4.1.1 Research questions

Specifically, in this chapter we seek to answer the following questions:

- RQ1** How is the *overall healthiness* of a focal person's food choice affected by the healthiness of an eating partner's food choice? Does the focal person's food choice change, and if so, in what direction?
- RQ2** How is a focal person's choice of *specific food categories* affected by the healthiness of an eating partner's food choice? Does the distribution over food categories change, and if so, what items are purchased more, and what items less?

4.1.2 Summary of main findings

Regarding question 1, we observe that, when a focal person acquires a new eating partner, the healthiness of the focal user's food choice shifts significantly in the direction of their new eating partner's dietary patterns. In a difference-in-differences analysis of 415 comparison pairs of focal persons (identified among a total of around 39,000 persons in eight years' worth of log data), which carefully controls for a number of confounding covariates, we find clear evidence of social influence: focal persons acquiring a healthy-eating partner change their habits significantly more in the direction of healthy foods than focal persons acquiring an unhealthy-eating partner. We quantify the robustness of this finding in a sensitivity analysis, and we provide further evidence by observing a dose-response relationship between the difference in exposures and the difference in effects.

Regarding question 2, we observe that focal persons who start eating with healthy-eating partners show an increase in the purchase of coffee and lunch meals, items generally purchased in




Figure 4.1: Distribution of purchases across item categories, across a campus. Most purchased items are coffee and lunch meals.

large numbers, with the strongest effect. On the other hand, items purchased at higher rates by the matched counterparts, who start eating with unhealthy-eating partners, loosely form a cluster of potentially unhealthy items that should not be eaten in large quantities (soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes).

4.1.3 Implications

Students and staff consume large amounts of food on campuses, daily and globally. The present chapter shows the value of employing novel methods relying on population-scale digital traces to measure social influence on food choice behaviors in this context. The derived insights have the potential to support interventions aimed at encouraging more healthy and sustainable dietary habits in university environments and beyond.

4.2 Materials and methods

4.2.1 Transaction log data

In this chapter, we leverage the anonymized dataset of food purchases introduced in the previous chapter (cf. Section 3.2.1). To summarize, the data spans eight years, from 2010 to 2018, and contains about 38 million transactions, of which about 18 million were made with a badge that allows linking to an anonymized person's ID. Each transaction is labeled with the time it took place, information about the sales location, the cash register where the transaction




Figure 4.2: **Left:** Annual distribution of food purchases. The trends mirror the university schedule: the number of transactions drops at the end of the spring semester (around week 25), and increases again at the start of the fall semester (around week 40). A similar pattern is observed before the beginning of the spring semester (around week 10). **Right:** Annual distribution of detected onsets of social ties. The ties emerge disproportionately often when classes start at the beginning of the fall semester (by a factor of 3.5 times, compared to a baseline sampled at random from the distribution of purchases).

took place, and the purchase items. The distribution of purchases across categories is shown in Figure 4.1. Purchases are not evenly spread over the course of the year, but, as expected, are higher during semesters, and lower during the breaks between semesters (Figure 4.2, left). In this chapter, we also leverage the smaller-size enriched transactional dataset gathered during a campus-wide sustainability challenge (introduced in Section 3.2.2). The enriched dataset was not used for the analyses, but only for assessing the accuracy of our heuristic method for inferring frequent eating peers (described next).

4.2.2 Inference of co-eating onset from proximity in transaction logs

To measure the effect of the emergence of new social ties, we first infer frequently co-eating persons based on the proximity in the transaction logs. Frequently co-eating persons are likely to share a social tie, i.e., they are persons likely to be friends, colleagues, or classmates who often eat together. Previous work has shown that such social ties can be reliably inferred from geospatial proximity [79].

To infer frequent eating peers, we monitor a sequence of transactions made on the same day with the badge in the queue of a fixed cash registry, in a given shop. We identify situations when two individuals are adjacent in the queue and make a transaction within one minute between each other, with no one in between them. We use a lower threshold of 10 such high-confidence proximity indicators to infer a likely social tie. The first appearance of proximity in the logs is then considered to be the onset of co-eating. We observe a spike in tie formation coinciding with the start of classes in the fall (Figure 4.2, right).

Furthermore, we evaluate the precision of our heuristic by comparing the inferred co-eaters with ground-truth team membership information from the sustainability challenge. We observe that team membership in the sustainability challenge, a ground-truth indicator of a

social tie, is correlated with sharing an inferred tie based on the transaction logs: out of all the pairs of individuals from the sub-population taking part in the sustainability challenge who are detected as frequent eating partners, 72% are also members of the same team.

4.2.3 Inference of nutritional properties from transaction logs

We infer a set of summary nutritional properties from raw transaction logs by relying on a set of pre-established criteria. We derive healthiness labels based on food-pyramid recommendations [381]. Products that should be consumed in the least amounts possible, i.e., items at the top of the Swiss food pyramid (with high amounts of saturated fats, salt, added sugars, refined grains, and highly processed foods) were considered as “unhealthy” (e.g., sodas, chips, candies, and chocolate bars). Other products that are not at the top of the Swiss food pyramid are considered to be “healthy” (including non-sweetened beverages, fruits, vegetables, whole grains, meat, fish, and nuts). When insufficient information was available from the name of the product, “unclassifiable” was selected.

Two professional epidemiologists specialized in nutrition independently assessed each food item and categorized them into healthy vs. unhealthy vs. unclassifiable. The reviewers had access to the unstructured textual description of the item (e.g., “coffee”, “croissant”, “Coca-Cola can”). The reviewers did not have access to any other meta-information about the items. Disagreements were resolved by a third reviewer. Labels are used to create a healthiness score of a set of purchases by averaging individual product scores, coded numerically as 1 for healthy (25% of items), -1 for unhealthy (46% of items), and 0 for unclassifiable (29% of items).

4.2.4 Matched incident user design with active comparators

Recall that we are interested in determining whether and how eating with others impacts the nature of food consumption. As depicted in Figure 4.3(a), a naïve approach to answering those questions would be a *cross-sectional design*: at any given absolute point in time, some people are regularly eating with their peers (indicated with green) while others do not (indicated with gray). Starting from a certain absolute point in time t_0 , by identifying persons with different habits, one could compare what is consumed by the persons who do not have a regular eating partner with what is consumed by the persons eating with a regular eating partner. One could also compare the food consumed by persons who are eating regularly with partners who have different habits.

The problem with this setup is that those persons who do not eat with others might have done so in the past (e.g., Person 1 in Figure 4.3(a)). Those who do eat with others might have been doing it for a long time or might have just initiated. Also, some people stop eating with others, whereas other people continue. It could be that those who stop do so because they prefer the diet they seek when eating alone (i.e., selection bias). Additionally, people who eat with others might differ in fundamental ways from those who do not.

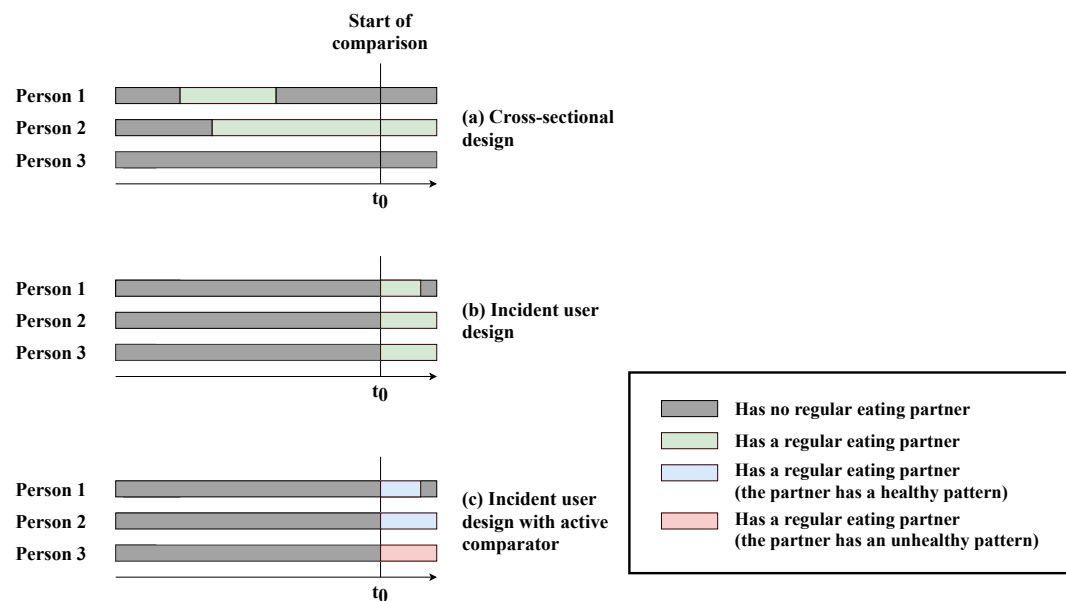


Figure 4.3: Study design diagrams. We illustrate three potential observational study designs to estimate the effect of eating with other persons on food choices, (a) *cross-sectional design*, (b) *incident user design*, and (c) *incident user design with active comparators*. At different points in time, a person either does not have a regular eating partner (marked in gray), or she does (marked in green, red, or blue). A cross-sectional design observes food consumption at the start of the monitored period, a fixed time t_0 , which is the same across all participants. Incident user design isolates the effect of the onset of co-eating with another person on subsequent food consumption. In incident user design, time is tracked relative to the moment of onset t_0 , which may be different across participants. The active comparator design additionally allows for comparisons of the effect of onset among persons who all start to eat with someone, but their partners have different characteristics (marked in red and blue). The present chapter is based on an incident user design with active comparators (presented in more detail in Figure 4.4).

For these two reasons, looking at everyone at the same moment in a cross-sectional way can be problematic. To overcome these challenges, we can turn to an *incident user design* (Figure 4.3(b)), which restricts the population to those people who newly initiate the treatment—starting to eat together with another person. We are interested in the causal effect on food consumption of initiating eating with a peer. Among people who had no regular eating partners in the past, what is the causal effect of starting to eat with a peer? In this way we isolate the causal effect of initiation. We restrict the observed population so that none of the persons have a history of eating with someone. Note how Person 1 in Figure 4.3(b) starts eating together with a regular partner, but then after a while no longer has a regular eating partner. This is not an issue because we are interested in the effect of the onset.

As opposed to the cross-sectional design, where time is absolute, the incident user design offers the flexibility of tracking time relative to an onset t_0 that may be different for different

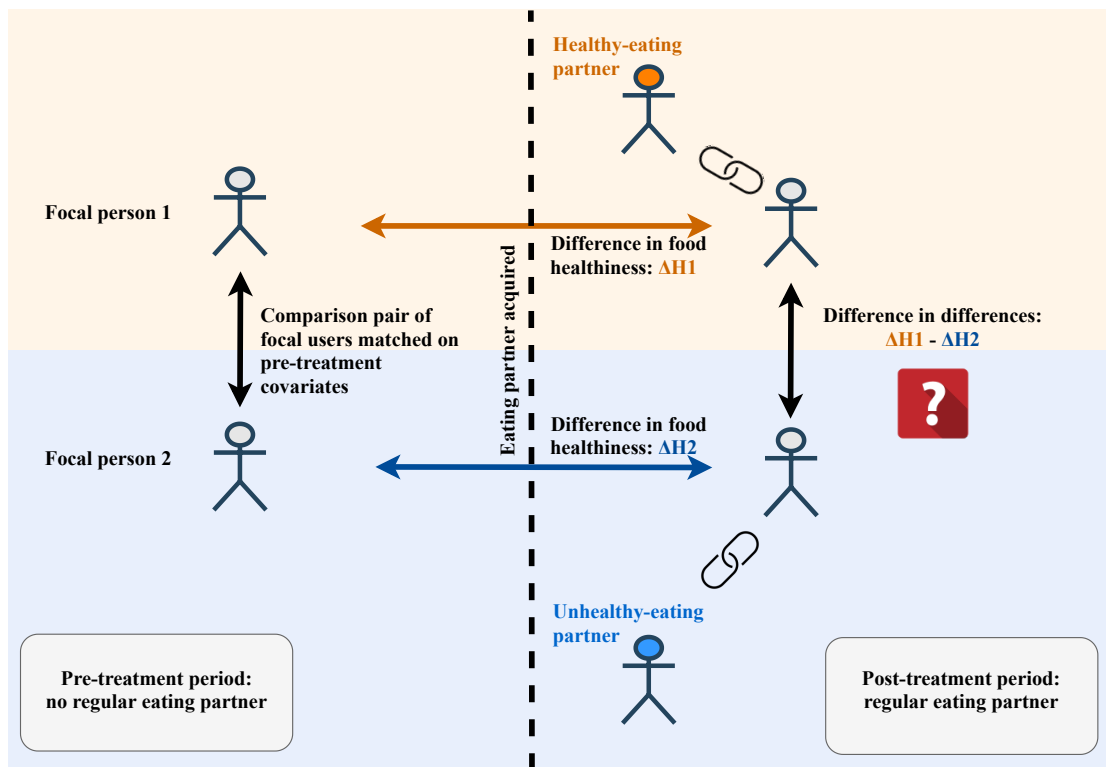


Figure 4.4: The matched incident user design with active comparators on which the present study is based. We identify comparison pairs of focal persons 1 and 2, who are indistinguishable in the pre-treatment period and have no regular eating partners, until the moment of co-eating onset, when they each acquire a regular eating partner. Focal person 1 starts regularly eating with a healthy-eating partner, while focal person 2 starts regularly eating with an unhealthy-eating partner. The comparison pair of focal users is then observed in the pre-treatment period (no regular eating partner) and post-treatment period (regular eating partner). The effect of the co-eating onset is estimated using a difference-in-differences analysis.

participants. Although this design allows us to compare different treatments, the problem with this setup, which persists from the above-described cross-sectional design, is that, if the comparison group is “no treatment” (i.e., no initiation of co-eating), it is not apparent when the follow-up should start for the “no treatment” group. Additionally, selection bias remains and is not accounted for, as people who do not initiate might in other fundamental ways differ from those who do initiate.

Our study design addresses these challenges by implementing a variant of incident user design, *incident user design with active comparators* (Figure 4.3(c)). Here, before initiation, no user included in the study had a regular eating partner (i.e., was treatment-free). We compare the effect of initiating to eat with partners who have different habits among persons who all initiate to eat with someone (illustrated with blue and red in Figure 4.3(c)). Active comparator designs tend to involve significantly less confounding [191, 224, 406], as people who eat with




Figure 4.5: Daily fraction of purchases annotated as potentially unhealthy, tracked over five years. A seasonal pattern emerges. Drops in the daily fraction of unhealthy purchases coincide with between-semester breaks.

different kinds of others are more alike among themselves than when compared to people who do not have regular eating partners.

Our study design is illustrated in more detail in Figure 4.4. We identify persons (referred to as *focal persons*) who had no regular eating partners and, at a moment t_0 specific to that focal person, initiate eating with someone (referred to as *eating partner*). Here, as defined in Section 4.2.2, a person qualifies as a focal person's potential eating partner if the two were observed making subsequent purchases in the same queue within one minute of one another on at least ten occasions in the entire dataset, and the onset of co-eating is defined as the first one of these occasions. We then isolate pre-treatment and post-treatment periods of the focal person's food purchases comprising all transactions made six months before the first purchase together (moment t_0) and six months after, respectively. We ensure that the focal person does not initiate eating with anyone else in the pre- and post-treatment six months. The length of the pre-treatment period is chosen so that it is feasible to expect that an individual will be present on campus given the typical stay in the logs (the total observed 12 months of pre- and post-treatment correspond to one school year).

Some persons initiate co-eating with a person who has a positive healthiness score in the aligned pre-treatment period. In contrast, some initiate co-eating with a partner who has a negative score. These are the two groups that we seek to compare (we refer to the two types of partners as *healthy-eating partner* and *unhealthy-eating partner*).

For a focal person who starts to eat together with a partner who has a healthy dietary pattern, an *active comparator* (or *counterpart*), will be another focal person who starts to eat together

with a partner who has an unhealthy dietary pattern. The potential counterparts start to eat with their partner in the same month as the other counterpart. This is done in order to control for temporal confounds that might arise from a seasonal variation of food popularity: as seen in Figure 4.5, unhealthy foods are especially popular at certain times of the year. The healthiness of the partner's dietary pattern is determined according to its numeric value (greater or less than zero), and not relative to the focal person.

Comparing incident users with active comparators is an important step toward reducing the impact of biases. However, in the assignment of the type of treatment, there can still be confounding. For example, it might be the case that only people who already have healthy habits start eating together with a partner who has healthy habits, due to a preference for similar others. The influence of the partners would then be indistinguishable from the impact of selection biases caused by homophily.

Hence, we turn to a *matched* incident user design with active comparators. We introduce an improvement over the previously discussed setup, where the incident users are matched to the potential active comparators while additionally controlling for pre-treatment covariates. Our goal here is to balance potential confounding variables within pairs, to be able to observe how the onset of co-eating with partners with different dieting patterns is associated with subsequent changes in the focal person's dieting pattern. We achieve this by performing a propensity-score-based causal analysis. We approximate randomized treatment assignment by modeling the propensity to experience the assigned intervention, relying on a number of pre-treatment covariates describing the focal persons' eating profiles. Due to the balancing property of propensity scores [300], matching on propensities results in similar covariate distributions between groups that differ in their assigned interventions.

The covariates capture important dimensions of the pre-treatment dietary pattern of the focal person: *where* the food is purchased (what is the shop where the person most frequently buys food), *when* the food is purchased (what is the fraction of items occurring during lunchtime), *what* types of items are purchased (what fraction of purchased items are meals, and what is their estimated healthiness), and *how often* the person purchases food on campus (number of transactions). We measure these confounding covariates up to time t_0 .

We use a random forest model that predicts the type of treatment based on pre-treatment covariates of the focal person (area under the ROC curve: 0.87). This implies that past purchases allow us to accurately predict whether the tie will be formed with a healthy- or an unhealthy-eating partner, and that confounding is a real problem that needs to be addressed. The distribution of the propensity to start eating with a partner who has a high healthiness score is presented in Figure 4.6a. We also examine the feature importances in predicting the treatment assigned, i.e., the initiation of eating with a partner who has a healthy or unhealthy eating pattern (Figure 4.6b). We observe that the focal person's pre-treatment healthiness score is in fact the most important predictor of the type of partner the focal person will start to eat with, pointing at homophily.




Figure 4.6: **(a)** Distribution (before matching) of propensity to start eating with a healthy-eating partner. **(b)** Importance of most indicative features for predicting treatment assignment, i.e., initiation of eating with a healthy- vs. unhealthy-eating partner (shop names anonymized). Most important feature: pre-treatment healthiness score of focal person's purchases, which indicates homophily.

Focal persons in the two sets are then matched while ensuring that two potential matches have propensity scores (likelihoods of receiving the treatment) within a caliper of 0.1. The size of the caliper was chosen so that balance in covariates is achieved. Moreover, an exact match on the sign of the mean pre-treatment healthiness score and the most frequented shop is required to achieve tight control. We then create matched pairs based on possible candidates by performing maximum weight matching on the weighted bipartite graph, where nodes are focal persons, and the weights use similarity based on the Mahalanobis distance in covariates. We maximize the total similarity to find a maximal matching.

The result is a set of matched pairs of focal persons, indistinguishable up to the moment of initiation, who initiated co-eating with partners with different dietary patterns in the same month. This approach yielded 415 matched pairs of 830 focal persons who started to eat with different partners. We require at least ten high confidence indicators of eating together with the partner (Figure 4.7a). Partners' distribution of pre-treatment healthiness scores is shown in Figure 4.7b.

Our matched analysis then moves on to comparing focal people who initiate co-eating with a person with a healthy dieting pattern, to their counterparts who have the same dieting patterns up to the moment of initiation, but initiate co-eating with a partner who has an unhealthy dieting pattern. The post-treatment patterns are then compared across treatments within the matched population.

Before moving on to the analysis of the outcomes, we ensure that the matched persons are comparable by measuring the balance of their pre-treatment covariates (Table 4.1). We use the standardized mean difference (SMD) across covariates in the two groups to measure the balance. We observe that matching greatly reduces the SMD, as the largest SMD across covariates (the one of the pre-treatment healthiness score of the focal person) changes from 0.301

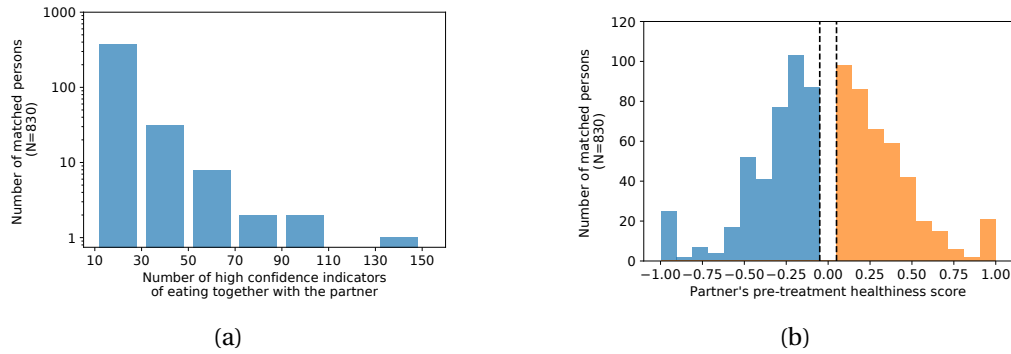


Figure 4.7: **(a)** Histogram of the number of high confidence indicators of eating together with their respective partners, for matched focal persons. **(b)** Histogram of the partner's pre-treatment healthiness score, across matched persons. Orange bars correspond to healthy-eating partners, blue bars to unhealthy-eating partners. A margin of 0.1 is ensured to differentiate the treatments.

before matching, to 0.023 after matching. Groups are considered balanced if all covariates have SMD lower than 0.2 [203], a criterion that is satisfied here.

4.3 Results

Recall our research question: we want to understand how a person's food choices are affected by the healthiness of a co-eating partner's food choices. Do people's choices change and, if so, in what direction (i.e., towards more or less healthy)?

4.3.1 Regression analysis of pooled data

First, we aim to determine if there are any significant differences between the outcomes of the matched focal persons. Is the pre-treatment healthiness of the partner predictive of post-treatment healthiness of the focal person? We start by performing a regression estimation of the effect of the partner's pre-treatment healthiness score on the focal person's post-treatment score.

Table 4.1: Pre-treatment covariate balance. To ensure that matched persons are comparable, we evaluate the balance of their pre-treatment covariates, via the standardized mean difference (SMD) across covariates in the two matched groups.

Pre-treatment covariate	SMD before matching	SMD after matching
Preferred shop (i.e., where the largest fraction of pre-treatment transactions is made)	exact match required	
Pre-treatment percentage of lunchtime transactions	0.109	0.045
Pre-treatment percentage of meal transactions	0.207	0.075
Pre-treatment mean healthiness score	0.301	0.023
Pre-treatment mean weekly number of transactions	0.071	0.023

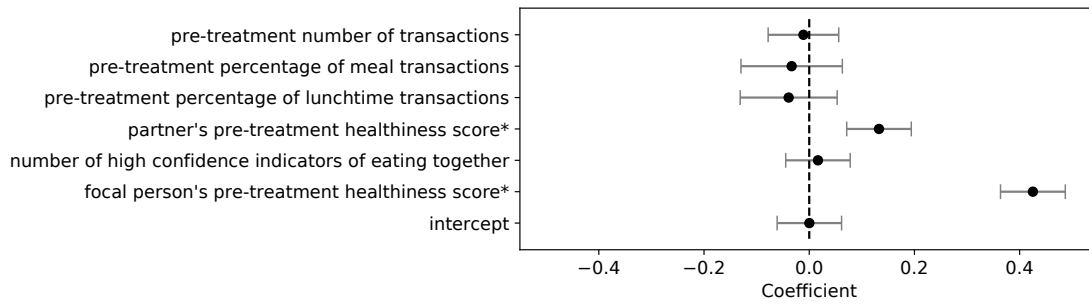


Figure 4.8: Linear regression effect estimates. The effect of the focal person's pre-treatment covariates, the partner's pre-treatment healthiness score, and the number of detected high-confidence indicators of eating together on the focal person's post-treatment healthiness of purchased items. The effects are estimated with linear regression ($R^2 = 0.194$); 95% confidence intervals are approximated as two standard errors. Significant coefficients ($p < 0.05$) are marked with an asterisk (*). The focal person's own healthiness score and their eating partner's healthiness score are the only two statistically significant factors associated with the focal person's post-treatment healthiness score.

We fit a model where the focal person's post-treatment healthiness score is the dependent variable. We include the focal person's pre-treatment healthiness score, the partner's pre-treatment healthiness score, the number of high-confidence indicators of eating together, and the focal person's pre-treatment covariates as the independent variables. The focal person's pre-treatment covariates (number of transactions, percentage of transactions that are meals, percentage of lunchtime transactions, and the pre-treatment healthiness score) are already controlled for by matching, but they are included in the model to account for possible residual confounding. The predictors and the outcome are standardized, so the coefficients are interpreted as increases in the healthiness score per standard deviation of the predictor.

In Figure 4.8, fitting the linear regression, we measure a significant positive effect of 0.13 (95% CI [0.07, 0.19]) of the partner's pre-treatment healthiness score. The focal person's own pre-treatment healthiness is the strongest predictor of post-treatment healthiness (coefficient 0.43, 95% CI [0.36, 0.49]). This is the first indicator that the pre-treatment score of the partner is associated with the focal person's patterns.

4.3.2 Contingency-table analysis

Next, to obtain fine-grained insights about patterns taking place at the pair level, we analyze the outcomes with a contingency table. We binarize the outcome to look for either an increase or no increase post-treatment, compared to pre-treatment. Four possibilities exist for any matched pair: both increased, only one or the other increased, and none increased. The contingency table is presented as Table 4.2. The table counts the frequency of the four possible results. Using a chi-squared test, we reject the null hypothesis of no treatment effect ($p = 0.00017$).

Table 4.2: Contingency table counting number of pairs of matched focal persons in each condition. Post-treatment healthiness score is compared to pre-treatment score to determine if there was an increase; in columns, for focal persons who start to eat with healthy-eating partners, and in rows, for their matched counterparts, i.e., focal persons who start to eat with unhealthy-eating partners.

		<i>Focal person with a healthy-eating partner</i>		Total pairs
		Increase	No increase	
<i>Focal person with an unhealthy-eating partner</i>	Increase	126	67	193
	No increase	103	119	222
Total pairs		229	186	415

It is particularly informative to observe the discordant pairs (off-diagonal entries in the contingency table) among the matched pairs. Such pairs correspond to situations when the outcome (increase or no increase) differs in the matched pair. The intuition is the following: if there is no effect, the two types of discordant entries should be balanced. However, we observe that in 103 pairs, the focal person with a positive intervention increased, and the focal person with a negative intervention did not. The reversed situation, in comparison, occurs in 67 pairs. We test the null hypothesis of no effect in a paired randomized experiment using McNemar’s test [210], which relies directly on the evidence that comes from the discordant pairs (their number and the ratio between them). Here, too, we reject the null hypothesis of no treatment effect ($p = 0.007$).

4.3.3 Difference-in-differences analysis

We move on to further exploit the matched setup in order to estimate the difference-in-differences [217] effect for pairs of matched focal persons. The idea is to first calculate the difference between post-treatment and pre-treatment healthiness scores for each focal person separately. Then, we can calculate the difference in treatment effects between two matched focal persons in each pair. Averaging the differences in differences across all pairs yields the overall treatment effect.

Regression model. In practice, following the standard approach, we estimate the difference-in-differences effect via a regression model. Here, each focal user adds two data points (one pre-treatment, one post-treatment), each of which specifies, as predictors, the type of partner with whom the focal user started to eat as a treatment (healthy-eating or unhealthy-eating) and the time period (pre- or post-treatment); and, as the outcome, the healthiness score of the focal user’s food choice during the respective period. Each matched pair thus contributes four data points, and the modeled dataset consists of $4 \cdot 415 = 1,660$ data points. Formally, the model takes the following form:

$$y_{it} = \alpha + \beta \cdot \text{healthy_treatment}_i + \gamma \cdot \text{treated}_t + \delta \cdot (\text{healthy_treatment}_i \cdot \text{treated}_t) + \text{error}_{it}, \quad (4.1)$$

where the dependent variable y_{it} is the focal user i 's healthiness score in period t , and the independent variables indicate whether i 's partner has a positive or negative pre-treatment healthiness score ($\text{healthy_treatment}_i = 1$ or 0 , respectively) and whether the respective data point captures the pre- or post-treatment period ($\text{treated}_t = 1$ or 0 , respectively). The coefficient δ of the interaction term, then, is the difference-in-differences effect of starting to eat with a healthy- vs. unhealthy-eating partner.

Results. Calculating the average difference-in-differences effect with a linear regression across all matched focal persons, we observe a larger post-treatment increase in focal persons with healthy-eating partners compared to the post-treatment increase in matched counterparts, $\delta = 0.051$ (95% CI [0.021, 0.076], $R^2 = 0.07$). This means that, accounting for possible temporal drifts between post-treatment and pre-treatment that are not associated with the initiation, focal persons starting to eat with a healthy-eating partner significantly diverge from their matched counterparts starting to eat with an unhealthy-eating partner. Quantitatively, the effect size of $\delta = 0.051$ means that, compared to matched counterparts who start eating with an unhealthy-eating partner, focal persons who start eating with a healthy-eating partner increase their healthiness score by an additional 5.1% of the full range spanning from a neutral healthiness score (i.e., 0) to the maximum healthiness score (i.e., 1).

Similarly, to estimate the effect of social tie formation on the absolute number of healthy and of unhealthy purchased items, we repeat the regression analysis described in Equation 4.1, but now with different dependent variables that capture the *total number of healthy and the total number of unhealthy items purchased* by focal user i in period t . We observe that the focal persons who start to eat with a partner with a healthy pattern purchase an additional 5.71 (95% CI [3.21, 8.21], $R^2 = 0.17$) healthy items, and an additional -1.13 (95% CI $[-3.04, 0.78]$, $R^2 = 0.12$) unhealthy items in the six months following the tie formation, compared to their matched counterparts.

Sensitivity analysis. The above finding relies on the assumption that there are no unobserved variables that create the differences between the matched focal persons that could explain away the measured effect. Sensitivity analysis is a way of quantifying how the results of our calculations would change if the assumptions were violated to a limited extent. If the conclusions of the study would change little, then the study is insensitive to a violation of the assumptions, up to the specified limited extent. In contrast, if the conclusions would change substantially, then the study is sensitive to a violation of the assumption.

The key assumption made in our analysis is that the treatment assignment is not biased, or in other words, that after balancing the pre-treatment covariates, the co-eating initiation with a healthy-eating vs. an unhealthy-eating partner is randomized (i.e., it is effectively decided by a coin flip). We measure by how much that assumption needs to be violated in order to alter our conclusion that there is a significant difference-in-differences effect on the healthiness of purchased items among the matched focal persons. Specifically, sensitivity analysis lets us answer the following question: if there is a violation of randomized treatment assignment




Figure 4.9: Sensitivity analysis. For the sensitivity $\Gamma = 1.17$, the amplification (Λ, Δ) is plotted (see text for explanation). Horizontal and vertical dashed lines indicate Γ , i.e., the asymptotic value of Λ for $\Delta \rightarrow \infty$, and vice versa.

(i.e., a deviation from a fair 50/50 coin flip), how large would it need to be in order to alter the conclusion that the null hypothesis of no difference between focal persons can be rejected? This notion is quantified by the *sensitivity* Γ , which specifies the ratio by which the treatment odds of two matched persons would need to differ in order to result in a p -value above the significance threshold. We always have $\Gamma \geq 1$, with larger values of Γ corresponding to more robust conclusions.

For the chosen $p = 0.05$, we measure a sensitivity of $\Gamma = 1.17$, which implies that, within matched pairs, an individual's probability of being the treated one could take on any value between $1/(1 + \Gamma) = 0.46$ and $\Gamma/(1 + \Gamma) = 0.54$ without changing our decision of rejecting the null hypothesis of no effect. In other words, if the assignment of the treatment after matching were not approximating the ideal 0.5, but a third variable made some people more likely to initiate eating with a healthy-eating or an unhealthy-eating partner, the randomized treatment assignment would have to be violated by deviating from the fair 0.5 by at least four percentage points.

Additionally, we conduct an amplification of the sensitivity analysis [301]. Amplification is particularly relevant when the concern is not about a violation of the randomized treatment assignment, but rather about the potential existence of a specific unobserved covariate. It then becomes useful to consider possible combinations of Λ and Δ , two parameters describing the unobserved covariate, that would result in the measured Γ . The strength of the relationship between the unobserved covariate and the difference in outcomes within the matched pair is defined by Δ , whereas Λ defines the strength of the relationship between the unobserved covariate and the difference in probability of being assigned a treatment. With these definitions, the sensitivity Γ can be expressed in terms of Λ and Δ , as $\Gamma = (\Lambda\Delta + 1)/(\Lambda + \Delta)$.




Figure 4.10: Dose-response relationship. **(a)** Histogram of pre-treatment differences in healthiness scores between partners of paired focal users. **(b)** Difference-in-differences effect between focal persons in matched pairs, stratified by the pre-treatment differences in healthiness scores between the partners they were exposed to.

The result of sensitivity analysis amplification is presented in Figure 4.9. For combinations of Λ and Δ in the orange area, significant effects would be detected (leading to $p < 0.05$), whereas for the combinations in the blue area, no significant effects would be detected (leading to $p > 0.05$). An infinite number of (Λ, Δ) combinations fall on the border; e.g., $(\Lambda, \Delta) = (2.0, 1.6)$ corresponds to an unobserved covariate that doubles the odds of treatment and multiplies the odds of a positive pair difference in the outcomes by 1.6.

Overall, we conclude that the study design is insensitive to small biases [299].

4.3.4 Dose-response relationship

Next, we analyze the dose-response relationship in our matched setup. Similar focal persons initiate eating with differing partners. We observe systematic changes in the dieting patterns of the focal persons after the tie formation. But do more drastic difference-in-differences effects occur when the differences between partners are more drastic? In the case of a true causal effect, one would expect a dose-response effect where focal persons diverge more post-treatment if their partners diverged more pre-treatment.

Although large differences in the pre-treatment scores between matched focal persons' partners are rare (Figure 4.10a), Figure 4.10b shows evidence of a dose-response relationship: the difference-in-differences effect is stronger when the partners are more different (i.e., the more extreme difference in partners leads to more extreme effect estimates). If there were other confounding factors that could explain the observed difference-in-differences effects, and those factors had nothing to do with the onset of eating together, we would not expect to find a dose-response relationship. The observed dose-response relationship thus further supports the conclusion of a causal effect.




Figure 4.11: Post-treatment increase in healthiness score, stratified by pre-treatment healthiness score of focal person (with 95% bootstrapped confidence intervals). The difference is shown separately for focal persons who start eating with a person with a positive (orange) vs. negative (blue) healthiness score. The difference is monitored in the first 3, 6, and 12 post-treatment months (left, center, and right panel, respectively).

4.3.5 Stratification by pre-treatment healthiness

Additionally, we would like to understand for whom the treatment is effective. Are there changes across the board with respect to the initial healthiness, or only for specific sub-populations? For whom is the intervention most efficient? We again monitor the differences between post- and pre-treatment healthiness scores, but now stratified into quartiles by pre-treatment healthiness score of the focal person (Figure 4.11). Moreover, we repeat this analysis for post-treatment observation periods of varying length (3, 6, and 12 post-treatment months). In the aligned, post-intervention period, persons who start eating with partners with healthy dieting patterns are characterized with consistently higher healthiness scores compared to the matched counterparts, across strata of the focal person's pre-treatment healthiness score. Note that the fact that the slopes are decreasing may be a simple regression to the mean. The key observation is that, within each stratum, when comparing the outcomes in orange and blue, people who initiate eating with a healthy-eating partner (orange) see a greater post- vs. pre-treatment difference compared to people who initiate eating with an unhealthy-eating partner (blue).

4.3.6 Analysis of affected food-item categories

Finally, we set out to understand the influence of new co-eating partners on the rates at which categories of food items are subsequently purchased. Since we observed that the behaviors are modified, we now ask: what items are eaten more, and which less? What foods being purchased and eaten in group settings on campus have the largest influence on others?

To estimate category-specific difference-in-differences effects, we repeat the regression analysis described in Equation 4.1, but now with a different dependent variable y_{cit} , which captures the number of items from food category c purchased by focal user i in period t . By fitting a separate regression for each food category c , we obtain category-specific effects δ_c .




Figure 4.12: Estimated difference-in-differences effects of co-eating onset with healthy- vs. unhealthy-eating partner on the frequency of purchased food categories (with 95% confidence intervals approximated as plus/minus two standard errors). Categories with a significant effect are marked in orange (positive) and blue (negative), whereas categories with no significant effect are marked in gray.

The estimated effects δ_c , together with 95% confidence intervals, are presented in Figure 4.12. We observe that the focal persons initiating to eat with a healthy-eating partner purchase more coffee, lunch meals, coffee from vending machines, soup, fruit, dessert, tea, salad, and wraps, compared to their matched counterparts, who purchase more soft drinks, drinks from vending machines, water, condiments, pizza, kebabs, and crêpes. The values on the x -axis can be interpreted as the number of purchased items by which the matched focal persons diverge in the post-treatment period. For example, in the six months following tie formation, people who start eating with healthy-eating partners purchase, on average, around two additional meals and around four additional coffees, compared to the matched counterparts. The matched counterparts who start eating with unhealthy-eating partners, by contrast, on average purchase around one additional soft drink in the six months following tie formation.

Coffees and lunch meals are the items that see the largest increase after tie formation with a healthy-eating partner. These items are in general purchased in large numbers (Figure 4.2).

Conversely, items with the strongest effect among the matched counterparts, with the exception of water, loosely form a cluster of potentially unhealthy items that should not be eaten in large quantities. The remaining items with a significant positive effect, soups, fruits, desserts, tea, salad buffet, and wraps are overall less indicative of an unhealthy dietary pattern.

4.4 Discussion

We report on a longitudinal observational study of the effect of the formation of social ties on food choice, leveraging a novel source of data: logs of millions of food purchases made over an eight-year period on a major university campus. To estimate causal effects from the passively observed log data, we control confounds in a matched quasi-experimental design: we identify focal persons who start regularly eating with a fixed partner and match focal persons into pairs such that paired focal persons are nearly identical with respect to covariates measured before acquiring the partner, but the new eating partners diverge with respect to the healthiness of their respective food choice behaviors (before tie formation with the focal person).

4.4.1 Summary of main findings

We observe that the people who acquire a healthy-eating partner change their habits significantly more toward healthy foods than those acquiring an unhealthy-eating partner. We further identify foods whose purchase frequency is impacted significantly by the partner's healthiness of food choice: coffees and lunch meals are the items that see the largest increase on behalf of those who initiate eating with a healthy-eating partner, whereas the matched counterparts, with the exception of water, increase purchases of items that loosely form a cluster of potentially unhealthy items that should not be consumed in large quantities: soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes.

4.4.2 Implications

Our findings show that digital traces can be used as a valuable tool for monitoring dietary habits, and can provide valuable insights into the effects of social ties on dietary choices. This chapter establishes the feasibility of relying on transactional logs in order to monitor food consumption and derive meaningful insights about behavioral patterns taking place at a population scale. Such digital traces can complement small-scale field experiments, making it possible to observe large populations over long time periods.

By relying on a novel transactional dataset and using carefully designed quasi-experimental methodology, we confirm theories of social influence on food choice postulating that social influence plays a prominent role and has a powerful effect on food intake [75, 170, 172, 250, 334]. Conforming to the behavior of others is adaptive, and individuals find it rewarding [172]. Hence, dietary choices are expected to converge with those of our close social connections.

Additionally, eating norms are known to reduce the intake of unhealthy foods [293] in specific contexts. For example, exposure to social eating norms is known to result in a reduction in the weight of consumed high-energy snack foods [293]. The fact that persons exposed to a partner eating healthy foods eat fewer items that should not be eaten in large quantities (soft drinks, drinks from vending machines, condiments, pizza, kebabs, and crêpes) compared to the counterparts indicates the presence of such positive influence of others.

Finally, we contribute to the rich literature about social influences in eating by demonstrating that, beyond studying social modeling in specific situations [370], the naturally or experimentally occurring event of tie formation is an important dimension to consider in order to understand the mechanisms of social influence and their potential for influence and interventions.

The most imminent utility of measuring the impact of social tie formation on food choice relying on large-scale passively sensed transaction signals lies in its potential to inform public health interventions on campuses. Designing large-scale nutritional interventions is challenging and logistically complicated. Additionally, it is difficult to predict their impact through experimentation due to ethical concerns stemming from the danger of eliciting undesirable effects. This work shows how observational insights based on passively sensed data can be used to evaluate the impact of potential interventions by estimating the impact of similar interventions that occur naturally, without external experimentation.

For instance, in order to incentivize healthy eating habits, university or corporate stakeholders might consider launching programs with disclosure and consent that help students or staff connect onsite with “lunch buddies” and incentivize consumption of meals in a company. Such programs leveraging peer-led nutritional interventions [405] would need to take into consideration self-selective disclosure and consent as confounding factors. First, by relying on passively sensed observational data and evaluating the impact of similar interventions, it would be possible to anticipate the impact on the involved individuals ahead of implementing any interventions in the real world. Second, it would be possible to optimize the pairing of people, by estimating on whom the effect of tie formation with others could be strongest. Finally, being informed about what products are most likely to be purchased in modified quantities after tie formation, the stakeholders could optimize the offering of products.

The question of whether the estimated impact of such naturally occurring interventions is identical to the true causal effect, and whether it is expected to mirror the impact of intentional externally induced interventions, remains. In what follows, we discuss the assumptions that would need to hold, and the limitations that should be considered.

4.4.3 Limitations and causal assumptions

Limitations. This study is subject to certain limitations, some of which suggest promising directions for future work. Inherent to observational studies, we recognize the inability to

infer true causality. Controlling for all the possible confounds is fundamentally infeasible. Still, we make a step towards understanding the effects of a certain “treatment” on food consumption by developing a quasi-experimental design based on propensity-score matching and difference-in-differences methods, whereby we seek to minimize biases due to *observed* confounding variables, enriched with a quantification of the danger of *unobserved* confounding variables via sensitivity analysis. Our work, therefore, provides insights based on passively sensed behavioral signals that go far beyond simpler correlational analyses.

We note that the inference of social ties might be imperfect. However, the fact that a large fraction of ties forms precisely at the beginning of the academic year (with the fall semester), when students are exposed to new fellow students (Figure 4.2), and the fact that there is a correlation with ground-truth team membership points towards reliability. In addition, we note that we might potentially be detecting the onset of co-eating with a delay (i.e., it actually occurs earlier than detected) if peers eat together as part of a larger group and are not directly adjacent in the queue, or if they use cash for the transaction. That said, we note that any potential delay in estimating the onset would lead to more conservative estimates of the effect of partnering up, as potential changes in the patterns would be counted as purchases before the tie formed.

Causal assumptions. We discuss what assumptions are necessary for our study design to let us isolate the causal effect of social-tie formation on food consumption. In particular, we consider the assumptions of the potential-outcomes framework [305] and the extent to which they can be assumed to hold in the present work.

Stable unit treatment value. Units are assumed not to interfere with each other. In other words, the treatment assignment of one unit does not affect the outcome of another unit, i.e., there is no “spillover” or “contagion.” Recall that we require initiation with no more than one peer during the monitored pre- and post-treatment periods. While this restriction ensures that there are no spillovers, our study is unable to capture complex network interactions occurring in on-campus settings. Future work should therefore generalize beyond the studied setup by identifying time-varying treatments and dynamic treatment regimes using g-formula methods [354, 407].

Consistency. The potential outcome under a treatment is assumed to be equal to the observed outcome when the actual treatment is received. In other words, the counterfactual outcome for treated units is the observed outcome for controls. While our study design attempts to compare people who are similar up to the moment of receiving the treatment by modeling the propensity to be treated, the assumption that the outcome of the matched person would be exactly the outcome of the treated person is untestable.

Positivity. The probability of every treatment for every set of covariates is assumed to be non-zero. In our study, it is reasonable to assume there are no people who could never possibly receive a given treatment.

Ignorability. Finally, the ignorability assumption refers to the absence of unmeasured confounding. By performing a sensitivity analysis, we have attempted to assess the possible impact of unobserved biases if the ignorable treatment assignment assumption is violated. This analysis leads us to conclude that our findings are insensitive to small biases.

Construct validity. We also consider the issue of the construct validity of our study design, i.e., the degree to which the obtained indirect measurements (transaction logs of food consumption) are reflecting the true phenomenon that is intended to be measured (actual food consumption). It is reasonable to assume that students and staff indeed consume the food that they purchase. However, one cannot eliminate the possibility of persons borrowing the card, or paying for items consumed by other people. Conversely, people on campus may consume food not recorded in the purchase logs (e.g., food purchased using cash, or food prepared at home or in off-campus restaurants). Future work should determine the extent to which the assumption that purchasing implies consumption, and vice versa, holds.

External validity. Future work should also determine external validity, that is, to what extent behaviors measured on campus reflect other settings off campus, e.g., food consumption—and the complex interplay with food consumption of their social ties at home—in families or in restaurants. People who primarily eat home-cooked food might present a skewed representation of their diet in the studied cafeteria purchase logs, and the fact that individuals might consume food prepared at home introduces unobserved variables into our analyses. For example, people who do not have regular eating partners in extended periods might eat alone because they prefer to eat a peculiar type of food that they cook at home and that is not served on campus, or measured via purchase logs.

4.4.4 Future work

This study opens the door for future research directions and potential follow-up studies of the mechanisms of social influence beyond those described, such as the role of the campus-wide sustainability challenge (cf. Section 4.2.1), and if eating with a healthy-eating partner is correlated with other behavioral changes (e.g., switching cafeterias). Future work should further understand the measured behavioral changes by understanding the underlying mechanisms and the social contexts in which the foods are consumed in modified quantities (e.g., socializing by meeting for a cup of coffee versus consuming a meal together). Future work should demonstrate whether modified habits are retained in those individuals who return later to eating alone for long stretches, or, if habits are not retained, how the individual regresses to the state they were in before the tie formation, and if so, what is the rate of the decay back.

5 Purchasing mimicry in food consumption on campus

5.1 Introduction

As described in the previous chapter, on-campus, formation of social ties is associated with changes in dietary behaviors. Outside of the campus environment, existing body of work has consistently observed similarities between connected persons in social networks, e.g., friends [124] and family [197, 272], in a number of experimental and survey-based studies [119, 162, 227, 269, 318, 346]. Food consumption has been found to be influenced by eating with others [170, 172, 218, 334], and the food choices of others, including people one does not know, have been observed to influence food choices [67, 294].

However, although such similarities in food consumption driven by social factors have been consistently observed, much less is known about the precise governing mechanisms. Numerous postulated mechanisms about how others influence our food consumption exist, including the processes of information gathering, self-presentation, minimizing regret, social desirability through appropriateness, and integration concerns [293].

Such mechanisms can result in both dish variety seeking and dish uniformity seeking [22, 89, 256]. Uniformity seeking was tested across a range of studies, for instance, in the form of uniformity by matching the other's food choice [55]. Conversely, variety seeking [285, 331] has also been observed since individuals make choices that diverge from those of others to communicate the desired identity effectively [32].

In particular, a potentially prominent postulated mechanism of interpersonal influence on eating behavior might be *behavioral mimicry*, which occurs when a person copies the behavior of another. For instance, individuals automatically mimic the gestures and hand movements of others, as an unconscious attempt to make the other individual like them since mimicry eases social interactions [58, 60].

It has been shown that viewing another individual performing an action activates an immediate reaction in an individual's motor system [169, 182]. Behavioral mimicry is also consistent

with the least effort principle when making choices—sometimes, simply doing what others do is the easiest choice [412]. Since eating is often habitual, i.e., automatically driven by external cues, unconscious behavioral mimicry may be a key interpersonal influence mechanism when eating with others. Previous studies have shown evidence of mimicry in behaviors linked to food consumption—people tend to adjust their intake directly to their eating companions by eating more when others eat more and less when others eat less [59, 353].

5.1.1 Research question

In order to understand dietary behaviors and design behavioral interventions in university environments, it is necessary to understand the answers to these questions. Identifying the role of purchasing mimicry in social norms is the first necessary step toward determining whether and to which extent purchasing mimicry can be exploited for behavioral interventions.

Specifically, in this chapter we ask:

RQ How prominent is food purchasing mimicry in campus environments? What foods are the most associated with purchasing mimicry, and what subpopulations are the most affected?

Despite the postulated importance of social factors, identifying and measuring mimicry in food consumption remains challenging. On the one hand, experimental studies monitor behaviors in artificial settings where people are aware they are being observed [349], which involves participation effect challenges, referred to as the Hawthorne effect [236, 358]. Furthermore, experimental studies to date have been limited to observing people in small-scale scenarios with a short duration, often in a laboratory setting [294, 295]. Most notably, such studies rely on confederate design, testing whether pairing a participant with an actor (i.e., a confederate) influences the amount and type [30, 294] of food eaten by the participant and their biting pattern [333], i.e., whether individuals take a bite of their meal in congruence with their eating companion rather than eating at their own pace [169]. More naturalistic experimental settings attempt to increase the validity of the findings by instructing participants to perform an unrelated activity while food is provided and consumption patterns recorded [168].

On the other hand, observational studies face limitations due to the presence of confounding factors and biases. In real-world settings, it remains challenging to measure and disentangle properties that are relevant in the context of food consumption, such as attributes of the individuals and the environment (e.g., food options available in different locations and settings). Another challenge is homophily, people's tendency to form ties with others similar to themselves to begin with [21, 90, 204, 218, 329, 330].

At present, characterizations of food purchasing mimicry originate from different experimental conditions. “Higher resolution” data and design paradigms are needed to identify behavioral mimicry and how it varies across foods and subpopulations. Researchers have only recently

been addressing gaps in the knowledge about human dietary behaviors by studying digital traces in the context of food consumption [315]. Such large-scale passively sensed signals have the potential to be harnessed in university campus environments to measure factors of well-being related to nutrition [151] and beyond [26, 227, 260, 324, 350].

The present chapter addresses the challenges of understanding the role of mimicry by leveraging the dataset of cafeteria records made on the EPFL university campus that captures the order of decision-making, allowing us to measure whether early decision-makers influence late decision-makers.

We test for evidence of purchasing mimicry, aim to identify the effect and understand the factors guiding it. In particular, we aim to determine the extent to which individuals on campus are influenced by others when deciding what to buy. What type of foods and what types of individuals (age, gender, status on campus) are most affected? What are the dyad characteristics associated with purchasing mimicry (e.g., is the effect stronger among the individuals with the same or differing attributes)?

Based on the transactional data, we design an observational study to identify the mimicry in food purchases. We leverage the sequential queue nature of cafeterias and the fact that with passively sensed data, we can monitor many persons in many situations. We consider a large number of situations where a dyad of partner (early decision-maker, i.e., the person who goes first in the purchasing queue) and focal user (late decision-maker, i.e., the person who goes second in the purchasing queue) are adjacent, and both make a purchase. We identify 0.5M such adjacent purchases, which we refer to as *situations* (Figure 5.1).

The large number of situations, rich data about the environmental context, and information about historical patterns of the individuals let us measure with high granularity and scale. We identify comparable situations and perform causal analyses. In particular, we devise a matching-based methodology to identify the effect of early decision-makers on late decision-makers, while minimizing the impact of biasing factors.

5.1.2 Summary of the main findings

Analyzing purchasing behaviors, we find significant mimicry of partners' purchases affecting all food types. We find that the partner's influence on the focal person diminishes once the ordering of the queue is randomized (cf. Section 5.3.1). The observed effect is robust across subpopulations and affects all genders and statuses, while it is the strongest for students and younger persons (cf. Section 5.3.2).

Our analyses of purchase logs provide novel insights into purchasing mimicry. First, the novel dataset, its scale, and a large number of studied situations make it possible to study mimicry with greater statistical power compared to previous research.

Chapter 5. Purchasing mimicry in food consumption on campus

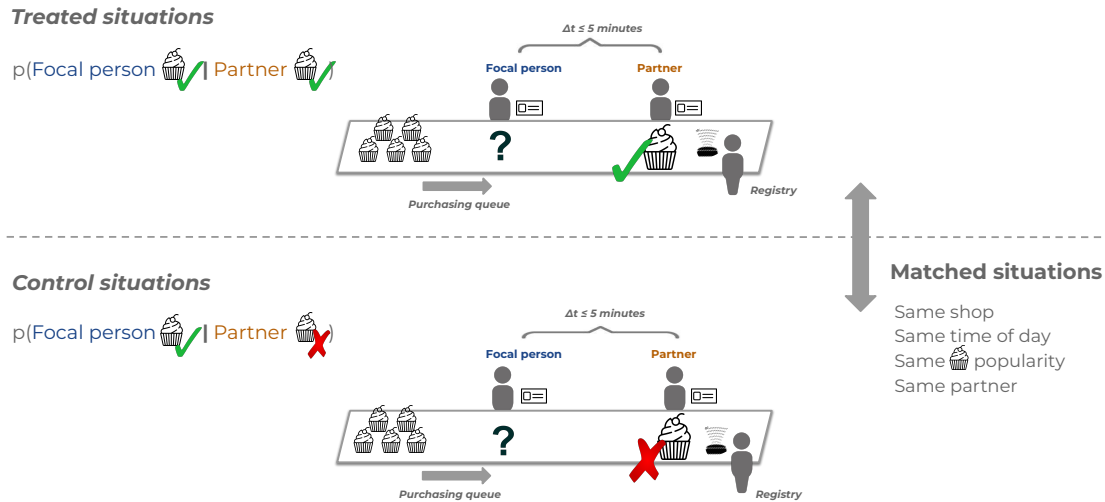


Figure 5.1: Study design. We identify situations where two individuals make purchases within five minutes of each other, with no one in between, adjacent in the purchasing queue. The first person to make the transaction in the queue is referred to as the partner, and the second person as the focal person. We are interested in identifying the impact that the purchasing behavior of the partner has on the focal person, i.e., in particular, purchasing a dessert in the illustration. To that end, the situations are matched, such that the situations are comparable (i.e., they occur in the same shop, time of day, same partner identity, same availability, and popularity of the dessert), but in treated situations, the partner purchases a dessert. In contrast, in the control situations, the partner does not purchase it. Our study then contrasts the focal person's probability of purchasing the dessert, given that the partner purchased (treated situations) or not (control situations).

Second, we perform careful causal analyses. Given the available information about individuals, who can be consistently tracked across many situations, and about the environment, such as the popularity and availability of different foods at shops in time, we can minimize the impact of numerous important confounding factors and isolate the mimicry in purchases. We carefully select suitable situations where the purchase decisions are monitored, aiming to disentangle homophily and influence. Additionally, we consider unobservable biases through sensitivity analysis.

Third, having access to a multi-year history of all transactions made on a large campus allows us to measure a wide set of purchasing behaviors that occur in the real world, as opposed to the artificial setting of lab-based studies.

In the next section, the study design and the estimation framework are described in detail.

5.2 Materials and methods

5.2.1 Sequential choices

In this chapter, we leverage the anonymized dataset of food purchases introduced in Chapter 3, Section 3.2.1. In order to identify purchasing mimicry, we monitor a sequence of transactions made with the badge in the queue of a cash registry, in a given shop. We identify situations when two individuals are adjacent in the queue and make a transaction within five minutes of each other, with no one between them (we investigate the impact of the delay in the purchasing queue in Section 5.3.3). We monitor two individuals making a purchase sequentially in a purchasing queue with the badge, as illustrated in Figure 5.1. A unit of analysis is an instance of two persons having a meal together, operationalized as two individuals executing transactions consecutively in the same shop. Henceforth, for brevity, we refer to such an instance of two adjacent purchases as *a situation*.

Co-eating matrices (Appendix A, Tables A.1, A.3, and A.2) outline the dyad frequency among the subset of the studied situations with demographic data available. The tables illustrate a preference for eating with others of the same gender, age, and status. We also note that the order female-male is more common than the order male-female. Similarly, the order staff-student is more common than the order student-staff, likely reflecting social norms of politeness and giving way to others depending on their gender and seniority.

There are three daily three peaks of transactions. The studied situations occur during the time of breakfast (before 11:00am), lunch (11:00am–14.30pm), or afternoon (after 14:30pm). During the three periods, persons purchase an anchor—a meal during lunch or a beverage (coffee or tea) during breakfast or afternoon (Figure 5.3a). In addition to the anchor food item, individuals might purchase an additional item (such as a dessert or a condiment), referred to as *an addition*.

In our main analyses, we study the effect of purchasing mimicry of the frequent food item additions (e.g., a dessert or a fruit). The additions were selected to include all food items where among the situations with the anchor, in at least 1% of situations, the partner buys the addition (Table 5.1) (i.e., at least 1% of the situations is treated). In total, there are three types of additions frequently purchased together with a beverage during breakfast and afternoon hours (fruit, dessert, and pastry), and seven types of additions frequently purchased together with a meal during lunch hours (condiment, salad, pastry, dessert, soup, soft drink, and fruit).

Overall, we analyze 509,220 identified adjacent purchase instances that took place in the twelve major shops. The studied adjacent purchase instances are executed by 5,504 unique unordered pairs of individuals. The instances are selected such that the two individuals make at least ten transactions together adjacent in the purchasing queues, in order to be able to monitor the same pairs repeatedly.

Table 5.1: Food addition items. For the three meals, the studied food addition items, the frequency with which the addition is purchased by the partner within the studied situations (i.e., treatment frequency), and the number of matched treated-control situations where the addition is purchased vs. not.

Time of day	Food addition item	Treatment freq.	# matched situations
Breakfast/morning snack time	Dessert	8.77%	1004
	Fruit	3.62%	1226
	Pastry	7.85%	16898
Lunch time	Condiment	1.49%	5590
	Dessert	1.21%	3954
	Fruit	8.49%	22424
	Pastry	1.93%	7400
	Salad	1.51%	5286
	Soft drink	2.8%	8970
	Soup	7.79%	18956
Afternoon/evening snack time	Dessert	6.39%	1288
	Fruit	2.7%	466
	Pastry	16.54%	3524

5.2.2 Causal assumptions and Directed Acyclic Graph (DAG)

Given a partner a and a focal person b , let $Y_a(t)$ be the partner's choice (set of purchased items within the transaction) and $Y_b(t)$ be the focal person's choice (set of purchased items within the transaction). We observe a person b (focal person), choosing items to purchase in situation t , $Y_b(t)$. The focal person's choice $Y_b(t)$ is governed by the focal person's eating profile X_b . Additionally, we consider common environmental factors in the specific situation t , $P(t)$, that can influence the choices of both observed individuals. Common environmental factors are operationalized as the location, the time of day, popularity, and availability of the item at the shop on the given day.

Furthermore, positioned in front in the queue, before person b , there is a frequent peer, person a (partner), choosing items to buy. Similarly, the partner's eating profile X_a impacts their choice $Y_a(t)$. Focal person b can be influenced by person a (partner) in their food choice $Y_a(t)$, corresponding to the causal path of food purchasing mimicry between $Y_a(t)$ and $Y_b(t)$. Due to the theoretical importance of "matching" the social norm and uniformity seeking through behavioral mimicry, we are interested in the causal path of purchasing mimicry, i.e., estimating the causal effect of the treatment ($Y_a(t)$) on the outcome ($Y_b(t)$).

However, the peer's choice can influence the observed person's choice through other biasing paths. In the presence of homophily, the social tie between persons a and b , $S_{a,b}$ is influenced by the traits of each individual X_a and X_b since more similar people tend to be closer friends given homophily, and in turn, influences the observed behavior $Y_b(t)$ through homophilic biasing paths, for closer friendship might make mimicry stronger. Eating profiles composed of habits and preferences are unchanging and independent of individual choices t . Social tie strength is a property of the network and is independent of the timing of individual choices t .

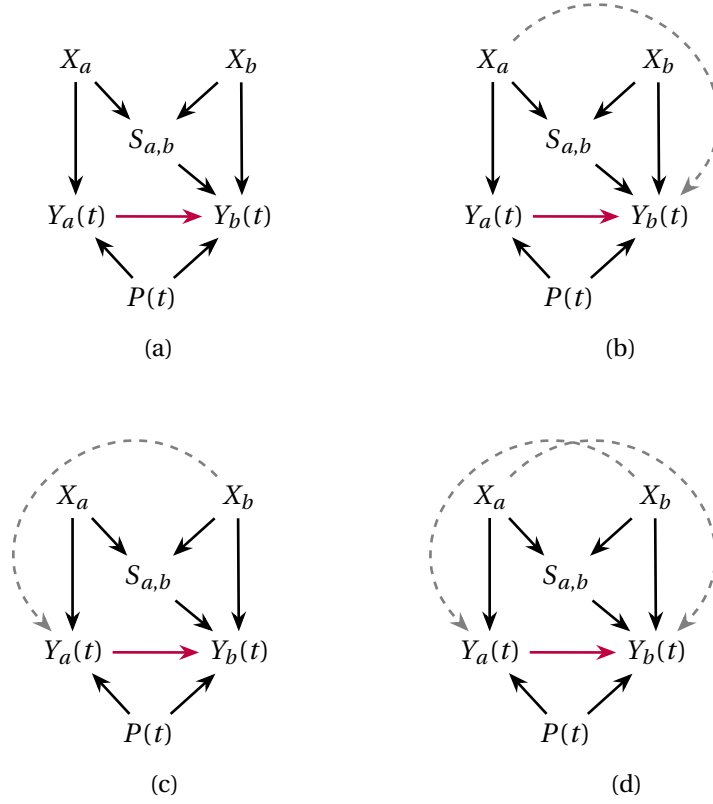


Figure 5.2: Directed acyclic graphs (DAGs) encoding the assumptions about the causal relationship between variables. X_a and X_b are partner's and focal person's eating profile respectively; $S_{a,b}$ is the social tie strength; $Y_a(t)$ and $Y_b(t)$ are partner's and focal person's sets of purchased items at time t respectively, and $P(t)$ are common environmental factors at time t . The purple arrow marks the causal path of purchasing mimicry. In (a), the assumed DAG. In (b), (c), and (d), the variations of the assumed causal relationships where the Assumption 1 is violated such that the traits of the individuals can influence the observed purchasing behavior through factors not related to friendship strength $S_{a,b}$.

In other words, we make the following assumptions:

Assumption 1. *The traits of partner X_a can influence the observed behavior of the focal person $Y_b(t)$ only through $S_{a,b}$ (we investigate this assumption further in Appendix A by considering alternative DAGs).*

Assumption 2. *We assume that $Y_a(t)$ influences $Y_b(t)$ through the ordering in the queue, while $Y_b(t)$ is not influenced by $Y_a(t)$, i.e., no coordination before making the decision (we investigate this assumption further in Appendix A).*

Assumption 3. *There are no other unobservable biases (we investigate this assumption further in Section 5.3.4 via sensitivity analysis).*

The causal graph reflecting these assumptions is presented in Figure 5.2a. The illustrated graph is the standard DAG assumed in the literature in order to identify the causal effect of social influence under the presence of homophily in a pairwise setup when examining the causes behind why a person manifested a behavior at a given time [221, 330]. The DAG is equivalent to the causal graph allowing for latent variables to influence both manifest network ties and manifest behaviors when the behaviors are time-independent, e.g., the choices are independent of each other, and there are no other unobservable biases.

According to backdoor criterion [153], the minimal sufficient adjustment set of variables for estimating the total effect of $Y_a(t)$ on $Y_b(t)$ is $\{X_a, P(t)\}$. Therefore, in our main analyses, we match on partner's identity to control for X_a and common environmental factors to control for $P(t)$. In Section A.2.1 we consider how our estimation framework and the subsequent estimates vary as Assumption 1 is violated and additional controls are necessary.

5.2.3 Matched estimation framework

The setup. To estimate the total effect of the partner's purchase on the focal person's purchase ($Y_a(t)$ on $Y_b(t)$), we perform matched estimation. In Section 5.2.2, given the assumed relationship between variables, the sufficient adjustment set of variables is the identity of the partner and the common environmental factors. Common environmental factors are operationalized by measuring the important dimensions of the dietary context: where the food is purchased (shop), when the food is purchased (time), the availability, and the popularity of the food, that day, in that shop as the fraction of all transactions that contained the food item.

Matching. We do matching of situations in order to find the pairs of comparable situations where in one situation the partner buys the addition i ($i \in Y_a(t)$), whereas in the other the partner does not buy the addition i ($i \notin Y_a(t)$). Within the pair of comparable situations, we ensure that the partner is the same person and that the situations took place at the same shop and during the same time of the day (breakfast time vs. lunch time vs. afternoon snack time). Additionally, we require that within the pair of comparable situations, the item was available in both situations and equally popular (up to 10% caliper) and that both the focal person and the partner purchase the anchor item (meal or a beverage). The size of the popularity caliper was chosen to achieve the balance in covariates, before analyzing the outcomes.

Covariate balance. For all the covariates except food item popularity, an exact match is required. For popularity, we ensured that after matching $SMD < 0.2$ (before matching $SMD = 1.23$, after matching $SMD = 0.08$). Groups are considered balanced if all covariates have $SMD < 0.2$, a criterion satisfied here [203].

Outcome analysis. The matched study design is illustrated in Figure 5.1. After matching, we analyze 96986 comparable situations, matched into 48493 pairs of situations. The distributions of matched situations across additions are outlined in Table 5.1. The result is a set of matched pairs of comparable situations, indistinguishable in the observed attributes, except that in

one, the partner buys the additional food item, whereas, in the other, the partner does not buy it. By monitoring different items, we apply our framework to measure the effect of different interventions in different subpopulations. To quantify the effect of an intervention partner “buying item”, our main analysis compares the purchases of the focal person in the matched situations.

Given a food item i , $Y_a(t)$ partner's choice (set of purchased items within the transaction) and $Y_b(t)$ focal person's choice (set of purchased items within the transaction), we measure risk difference (RD_i) and risk ratio (RR_i), calculated based on the 2x2 contingency matrix illustrated in Table 5.2. The two outcome statistics are defined as:

$$RD_i = p(i \in Y_b(t) | i \in Y_a(t)) - p(i \in Y_b(t) | i \notin Y_a(t)), \quad (5.1)$$

and

$$RR_i = \frac{p(i \in Y_b(t) | i \in Y_a(t))}{p(i \in Y_b(t) | i \notin Y_a(t))}. \quad (5.2)$$

The risk difference and risk ratio describe the absolute and the relative difference in the observed risk of events between treated and control situations. For a focal individual, they describe the absolute and the relative increase in the probability of purchasing the item when the partner buys the item vs. when the partner does not buy the item. Within the comparable situations, we resample situations to obtain the 95% bootstrapped confidence intervals.

Moreover, we consider a randomized baseline. In each situation, instead of the partner, we choose a random person from the purchasing queue (during the same day, at the same shop and same line). The objective of the randomized baseline is to understand similarities stemming from the contextual factors and not directly caused by the actual ordering of the queue and the partner's choice. The estimation, as previously described, is then performed on such randomized situations.

Table 5.2: Contingency table counting the number of situations in each condition depending on whether or not the partner purchased the item (rows), and whether or not the focal person purchased the item (columns).

		<i>Focal purchased</i>		Total situations
		No	Yes	
<i>Partner purchased</i>	No	40230 (41.48%)	8263 (8.52%)	48493 (50%)
	Yes	33332 (34.37%)	15161 (15.63%)	48493 (50%)
Total situations		73562 (75.85%)	23424 (24.15%)	96986 (100%)

5.2.4 Amplified asking

To make heterogeneous estimates depending on status at the campus (beyond the subpopulation of participants in the sustainability challenge), we rely on the paradigm of amplified asking to, first, build a model that can predict status in the sup-population where the status information is available, and then, second, amplify the entire dataset with the estimated class belonging, by making out-of-sample predictions over the whole population [315]. We train the classifier based on the features that capture temporal patterns typical of personnel and staff. For instance, students make summer and winter breaks, while staff might still be on campus. Similarly, students might make transactions in the later hours.

We use the total number of transactions, the number of years at the campus, and the distribution of transactions across months, weekdays, and hours in the day. The classifier uses a random forest model and achieves on a 20% held-out test set precision with respect to students 88.33%, with respect to personnel 78.26%, and recall with respect to students of 90.60%, and with respect to personnel 76.60%. Note that status estimation does not rely on the variables linked with the studied phenomena (purchased items) but merely on the temporal distribution reflecting when the individuals are present on the campus.

5.3 Results

The study design is illustrated in Figure 5.1. To summarize, individuals make their purchases at the registry with identifying badges. We identify situations ($N_{\text{before matching}} = 509220$) when individuals make purchases within five minutes of each other, with no one in between. As a reminder, we study scenarios where two individuals make a transaction while adjacent in the queue. We are interested in identifying the impact that the purchasing behavior of the partner (early decision-maker) has on the focal person (late decision-maker). We identify the change in the probability that the focal person will buy a food item when the partner buys the item vs. when the partner does not buy the item. We study dyads where the partner and the focal person are observed together at least ten times.

We perform matching of situations such that the situations are comparable (cf. Methods), but, in treated situations, the partner purchases a food item of interest. In contrast, in the control situations, the partner does not purchase the food item of interest. Matching is informed by the assumed causal graph.¹ Situations occur during breakfast (until 11h, 17.4% situations), lunch (11h - 14.30h, 72% situations), afternoon/evening snack (after 14.30h, 10.6%), as illustrated in Figure 5.3a. We ensure that the matched situations are comparable by requiring that the partner and focal person both purchase an anchor item (a meal during lunchtime or a hot beverage during the morning or afternoon/evening), and monitor the purchasing of one of the 13 food items frequently purchased with the anchor, selected based on frequency among the situations (Table 5.1, Methods).

¹In the Appendix A, we examine the robustness of our estimates as the causal assumptions are violated.

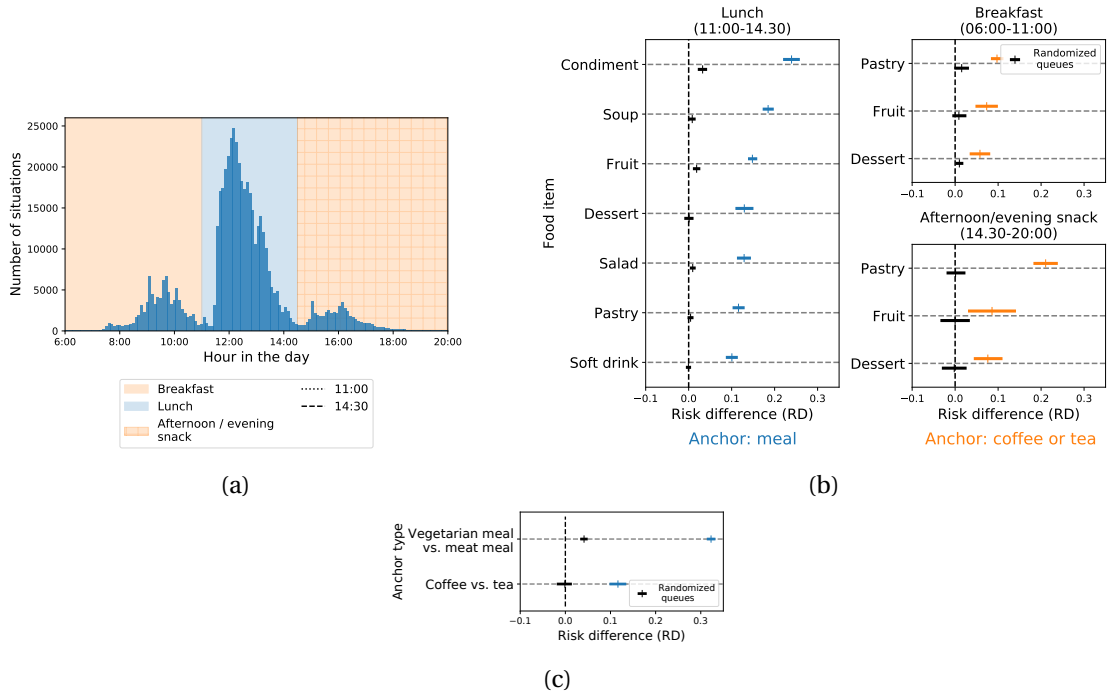


Figure 5.3: In (a), the histogram of situations during a day, on the x-axis, the hours in the day, and on the y-axis, the number of situations. The three peaks correspond to breakfast, lunchtime, and afternoon or evening snack time (shaded regions). In (b), separately for lunch, breakfast, and afternoon or evening snack, the estimated risk difference (on the x-axis), for the different food item additions (on the y-axis). In (c), the estimated risk difference (on the x-axis), for the anchor type (on the y-axis), type of meal, vegetarian vs. not, and type of beverage, coffee vs. tea. The error bars mark 95% bootstrapped CI. Risk difference estimates are colored (blue for lunch where the anchor is the meal, orange for breakfast and afternoon or evening snack where the anchor is a beverage). The randomized baseline is presented in black.

After matching, within the matched sets of comparable situations ($N_{after\ matching} = 96986$ into $N_{after\ matching}/2$ pairs of situations, where one of the 13 food items is bought or not, Table 5.1), we contrast the focal person's probability of purchasing the food item of interest, given that the partner purchased the item (treated condition) or not (control condition). The discrepancy between the two probabilities is expressed in absolute terms (risk difference) and relative terms (risk ratio). In what follows, we analyze the matched situations.

Table 5.3: Contingency table counting the number of pairs of matched situations in each condition (treated and control). In columns, situations where the partner purchased the item, and in rows, matched situations where the partner did not purchase the item.

Partner purchased (treated)		\neg Partner purchased (control)		Total pairs of situations
		\neg Focal purchased	Focal purchased	
	\neg Focal purchased	28111 (57.97%)	5221 (10.77%)	33332 (68.74%)
	Focal purchased	12119 (24.99%)	3042 (6.27%)	15161 (31.26%)
	Total pairs of situations	40230 (82.96%)	8263 (17.04%)	48493 (100%)

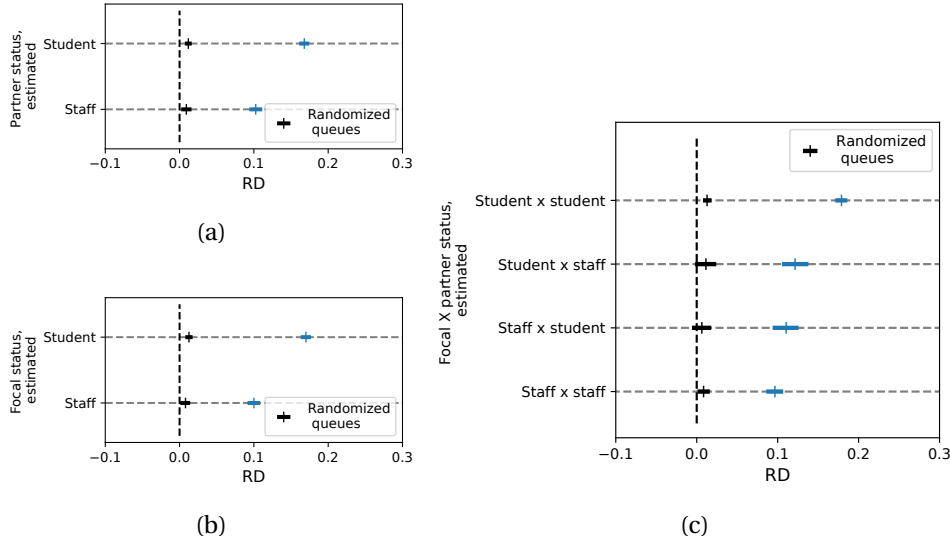


Figure 5.4: The estimated risk difference across the matched situations (on the x-axis), depending on the individuals' estimated status (on the y-axis). The error bars mark 95% bootstrapped CI. Risk difference estimates are presented in blue, and the randomized baseline is presented in black. In (a), for partner's status, in (b), for focal person's status, in (c) for the four combinations of the focal-partner status.

5.3.1 Significant mimicry within dyads affects all food types

Paired analyses. As a first look into the matched situations, we test for evidence of purchasing mimicry and aim to identify the effect pooled across food items. The contingency table (Table 5.3) counts the frequency of the four possible outcomes, comparing matched pairs of situations where in one case partner buys, and in the other case, the partner does not buy the additional item. Note that the most frequent outcome is that in both matched situations, regardless of the partner, both focal users do not buy the addition. The least likely is that in both situations, regardless of the partner, both focal users buy (since purchasing probabilities are in general low, cf. Table 5.1).

In particular, the discordant instances among the matched situations are informative, i.e., the off-diagonal entries in the contingency table, which correspond to matched situations where the two focal persons' purchases differ. If there were no partner effects, the two types of discordant entries would be balanced. However, we observe that more frequently, focal persons mirror their partners vs. do the opposite (2.3 times more likely). In 25% of pairs of situations, focal persons purchase when partners do, and focal persons do not purchase when partners do not. In contrast, the opposite scenario (focal persons doing the opposite of their partners) is rarer as it occurs in 11% of pairs of situations. The disbalance between the discordant instances is the first evidence of mimicry. Based on the contingency table, we reject the null hypothesis of no treatment effect (chi-squared test of no treatment effect $p < 10^{-12}$).

Risk analyses. Next, pooling the matched situations across the different items, we quantify risk difference (RD) and risk ratio (RR), cf. Methods. Henceforth, the risk-based effect estimate is the main method of analysis. Overall, across all matched situations (13 additions), we measure the risk difference of 14.22% [13.73%, 14.74%] and the risk ratio of 1.83 [1.79, 1.88], implying a change in the focal person's probability of purchasing the food item, depending on the partner's choice. In comparison, in the case of the randomized baseline where the purchasing order in the queue is randomized, we measure a risk difference of 1.07% [0.69%, 1.45%] and a risk ratio of 1.07 [1.05, 1.1]. In other words, the partner's influence on the focal person diminishes once the ordering of the queue is randomized.

Risk analyses across food items. Since effect modification is expected given varying times of day, separately for the different times of day and across the thirteen additions (seven lunch additions and three breakfast and afternoon/evening snack additions), we quantify the risk differences. We find that all the risk differences are greater than zero with 95% CI (Figure 5.3b). The random baseline is much smaller for all additions, consistent within the time of day and among additions.

Risk difference for lunch additions varies between 10.06% [8.65%, 11.42%] for soft drink and 23.94% [22.11%, 25.76%] for condiment. Risk difference for breakfast additions ranges between 5.78% [3.39%, 8.37%] for dessert, 7.34% for fruit [4.73%, 9.95%], and 9.74% for pastry [8.25%, 11.18%]. Note that pastry is a separate category from dessert since it can be savory. For afternoon or evening snack, risk differences are 7.61% for dessert [4.35%, 11.02%], 8.58% fruit [3.85%, 14.16%], and 21.06% for pastry [18.22%, 23.89%].

Risk analyses across anchors. Although the main analyses focus on food addition items, we also analyze the mimicry of the anchor itself (Figure 5.3c). The meal anchor can be vegetarian or meat-based, whereas the beverage anchor can be coffee or tea. While the main analyses focus on the purchasing mimicry of the additions, matching comparable situations meal (vegetarian meal vs. meat meal) and comparable situations beverage (coffee vs. tea), we observe significant risk difference for meal type (32% [31.39%, 33.21%]) and beverage type (11.65% [9.72%, 13.52%]). Estimates corresponding to the randomized baseline in case of meal type (4.14% [3.30%, 4.95%]) and beverage type (0.10% [-0.02%, 0.01%]) are again much lower. We note that mimicry is stronger for the meal-type anchor since purchasing vegetarian food is a behavior related to health and sustainability and, therefore, potentially more likely to be impacted by social norms.

To summarize, among the matched pairs of situations, we find significant mimicry of partners' purchases affecting all food types. The partner's influence on the focal person diminishes once the ordering of the queue is randomized.

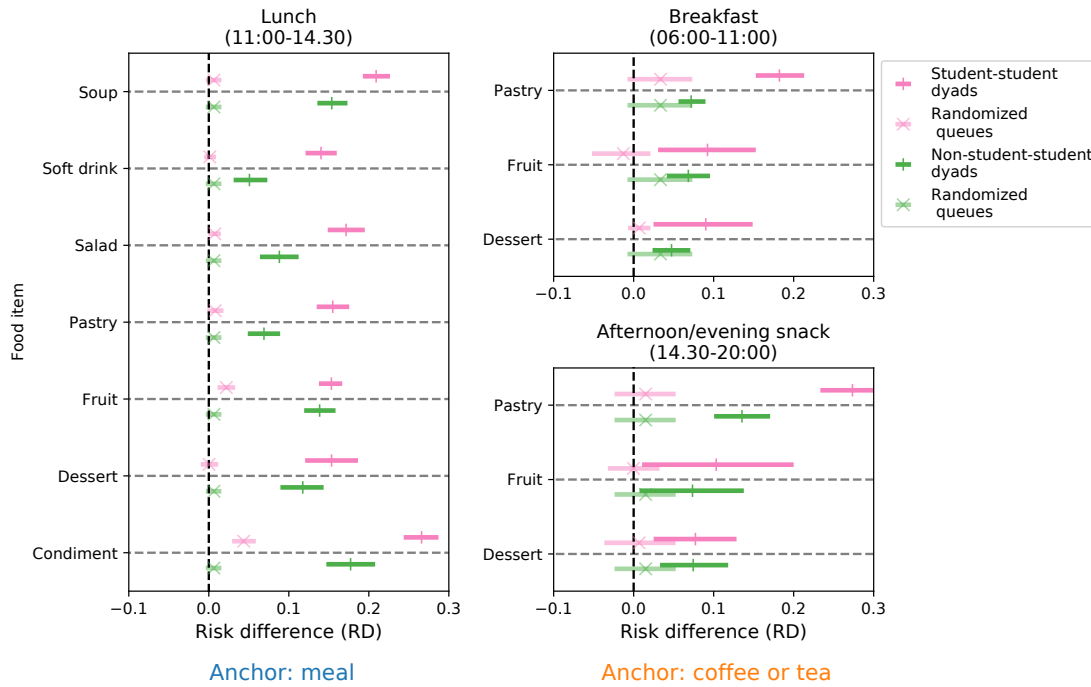


Figure 5.5: Separately for lunch, breakfast, and afternoon or evening snack, the estimated risk difference (on the x-axis), for the different food item additions (on the y-axis). The error bars mark 95% bootstrapped CI. Risk difference estimates are colored in pink for student-student dyads and in green for non-student-student dyads. Randomized baselines are presented in a lighter color.

5.3.2 Mimicry is strongest among students and the youngest subpopulations

Estimated status. We next measure the effect among subsets of matched situations based on the estimated status of the partner and the estimated status of the focal person in Figure 5.4. We find that the effect is stronger when *partner* is a student (risk difference 16.78% [16.10%, 17.46%]) vs. staff (10.25% [9.39%, 11.12%]; Figure 5.4a). Similarly, the effect is stronger when the *focal person* is student (17.0% [16.30%, 17.73%]) vs. staff member (10.01% [9.16%, 10.91%]; Figure 5.4b). Examining the four configurations of status within the partner-focal dyad (Figure 5.4c), we find that student-student is the condition with the largest risk difference (17.89% [17.11%, 18.60%]). In contrast, the staff-staff condition is the one with the smallest risk difference (9.66% [8.60%, 10.68%]).

The observation regarding students vs. staff differences holds across the different foods. In Figure 5.5, we measure the risk difference separately among estimated student-student dyads vs. all non-student-student dyads, where students can be focal or partner users, but not both. We find that across the three times of day and the different food items, the effect is consistently greater among the student-student dyads, implying that the difference depending on the status cannot be explained by any discrepancies in the preferred food items between students and staff.

Demographics: Status, age, and gender. We next investigate the effect across all the matched situations within the subpopulation with demographic data (Appendix A, Figure A.1). First, among the subpopulations with ground status (as opposed to inferred), in Figure A.1a and Figure A.1b), we consistently find that the effect is stronger both then partner is student (10.73% [5.67% , 15.59%]) vs. staff member (5.68% [1.85% , 9.38%]), and when the focal person focal is a student (14.15% [9.56% , 19.11%]) vs. a staff member (7.22% [2.44% , 11.38%]). Note that the differences are not statistically significant, likely due to the smaller sample size.

Second, we investigate the role of age, in Figure A.1c and Figure A.1d. Given the birthdate and the time of the transaction, we calculate the age at the time of the transaction, and we bin the age into terciles. We find that the effect is the strongest when both partner and the focal person are in the youngest group (≤ 22 years old at the transaction time). Monitoring the partner's age, we find that the effect monotonically decreases along the age bins (≤ 22 years old: 12.04% vs. 23-32 years old: 8.11% vs. > 32 years old: 4.98%). Similarly, monitoring the focal person's age, we find that the effect monotonically decreases along the age bins (≤ 22 years old: 17.72% vs. 23-32 years old: 13.18% vs. > 32 years old: 4.04%).

Third, regarding gender (Figure A.1e and Figure A.1f), we find a significant effect among all subpopulations, with a risk difference greater than zero when the partner is either male or female and when the focal person is either male or female. However, we observe no differences depending on the gender of the partner and the gender of the focal person.

To summarize, the purchasing mimicry effect is not restricted to particular subpopulations, as it is robust across subpopulations and affects all genders and statuses. However, it is the strongest for students. One of the driving mechanisms is age, as copying is the strongest among younger persons.

5.3.3 Mimicry diminishes as the proximity in the purchasing queue decreases

We next investigate the presence of a dose-response relationship. We measure the risk difference and risk ratio among subsets of matched situations with the different delays in the purchasing queue between the focal person and partner. In case of a true causal effect, one would expect a dose-response effect where the focal person's purchasing probabilities in the matched situations diverge more when the proximity is stronger, as they would correspond to instances where the focal person is more likely to have seen the choice of the partner.

As the distance (measured in seconds) between the adjacent persons in the purchasing queue increases (distribution illustrated in Figure 5.6a), the effect estimate decreases too (Figure 5.6b and Figure 5.6c). We measure a significant anti-correlation between the delay in the purchasing queue and risk difference ($\beta = -0.002$, two-sided $p = 8.71 \times 10^{-5}$) and between the delay in the purchasing queue and risk ratio ($\beta = -0.03$, two-sided $p = 2.24 \times 10^{-6}$). Overall, a larger effect is observed for smaller distances in the queue.

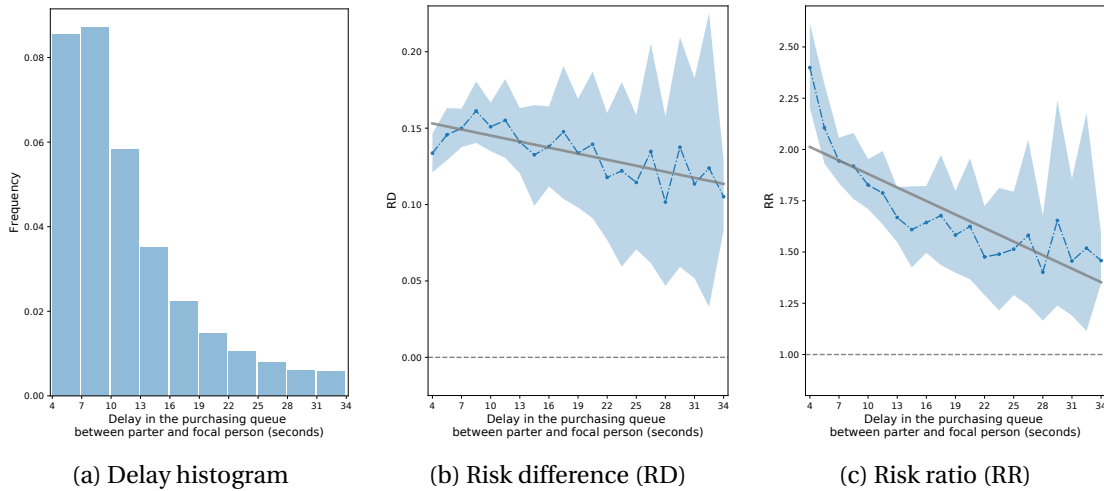


Figure 5.6: Dose-response. In (a), the histogram of the temporal delay between the partner's and focal person's transactions among the studied situations. On the x-axis, the delay, and on the y-axis, the frequency. In (b), the risk difference estimate within the subset of matched situations (on the y-axis), with the given delay in the purchasing queue (on the x-axis). In (c), the risk ratio estimate within the subset of matched situations (on the y-axis), with the given delay in the purchasing queue (on the x-axis). The shaded areas mark 95% bootstrapped CI. The gray dashed line represents the least square linear fit. Note the truncated x-axis. Situations with a delay of up to five minutes are considered. However, they are rare, as visible in (a).

If other factors were causing the purchasing similarity, for instance, a third-party present in the shop and convincing individuals to purchase a food item or not, and such factors had nothing to do with the ordering and the distance in the purchasing queue, we would not expect to see a dose-response relationship. The observed dose-response relationship supports the evidence of a causal effect.

Furthermore, we investigated an alternative hypothesis (cf. Appendix A, Section A.2.3) where the observed similarities between adjacent persons in the purchasing queue are driven by the fact that the two persons coordinated to go for a meal together and agreed on the food choice before lining up in the purchasing queue. Since we identify pairs where the null hypothesis of no effect of the order can is rejected, we argue that it does not appear plausible that pre-purchase coordination can entirely explain the measured effect in all pairs.

5.3.4 Sensitivity analysis

Finally, our findings rely on the Assumption 3 that there are no unobserved variables creating differences between the matched situations that could explain the measured purchasing similarity between partners and focal persons. We perform sensitivity analysis to quantify how the estimates made here would change if this assumption were violated to a limited extent. How strong would the unobserved biases need to be to explain the difference in outcomes

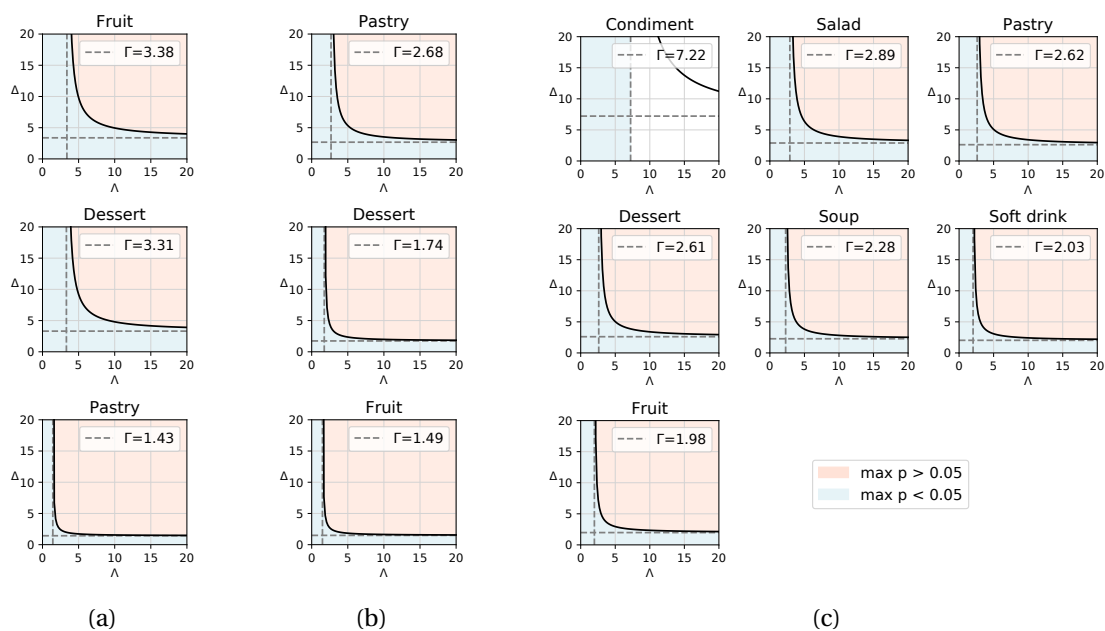


Figure 5.7: Sensitivity analysis. In (a) breakfast, in (b), afternoon snack, and in (c) lunch additions. For the measured sensitivity Γ , the amplification (Λ, Δ) is plotted. Horizontal and vertical dashed lines indicate Γ , i.e., the asymptotic value of Λ for $\Delta \rightarrow \infty$, and vice versa.

between the two sets of matched situations? Specifically, we measure, if there is a violation of the randomized treatment assignment among the matched situations (the choice of the partner), how large would it need to be in order to alter the conclusion that the null hypothesis of no differences depending on focal person's choice can be rejected? As in the previous chapter, this quantity is quantified with Γ , specifying the ratio by which the treatment odds in two matched situations would need to differ to result in a p-value above the significance threshold (larger values of Γ correspond to more robust conclusions).

Figure 5.7 summarizes the sensitivity analyses. For $p = 0.05$, we measure the sensitivities ranging between 1.43 (purchasing a pastry with a breakfast beverage) and 7.22 (purchasing a condiment with lunch). Additionally, we perform and illustrate (Figure 5.7) an amplification of the sensitivity analysis [301] where Γ is expressed in terms of Λ and Δ , as $\Gamma = (\Lambda\Delta + 1)/(\Lambda + \Delta)$. Δ is defined as the strength of the relationship between the unobserved covariate and the difference in outcomes within the matched pair, whereas Λ is defined as the strength of the relationship between the unobserved covariate and the difference in the probability of being assigned a treatment.

For combinations of Λ and Δ in the orange area, significant effects would be detected (leading to $p < 0.05$). In contrast, no significant effects would be detected for the combinations in the blue area (leading to $p > 0.05$). An infinite number of (Λ, Δ) combinations fall on the border. For instance, in the case of purchasing fruit during breakfast, $(\Lambda, \Delta) = (5.0, 9.8)$ corresponds to an unobserved covariate that increases the odds of treatment five-fold and multiplies the

odds of a positive pair difference in the outcomes by 9.8. Such amplification is relevant when the concern is not about the violation of randomized treatment assignment but about the presence of specific unobserved covariates with assumed Λ or Δ . Overall, we conclude that the study design is insensitive to moderate biases [299].

5.4 Discussion

To promote healthy and sustainable dietary choices in campus environments, stakeholders and commercial food providers need insights regarding food purchasing behaviors on campus. To that end, population-scale transaction data can be leveraged to measure, understand, and enhance purchasing behaviors.

Since social norms have long been suspected to play a crucial role in the design of such interventions, in this chapter, our goal is to identify and characterize purchasing mimicry, theorized as the driving mechanism behind the social influences between individuals on campus. The results presented in this chapter document the prominent role of purchasing mimicry and highlight the need for taking it into account when designing how the foods are offered on university campuses and beyond.

5.4.1 Summary of main findings

First, we find significant mimicry of partners' purchases affecting all food types. The partner's influence on the focal person diminishes once the ordering of the queue is randomized (cf. Section 5.3.1). Second, we find that the effect is not restricted to particular subgroups, as it is robust across subpopulations as it affects all genders and statuses. However, it is the strongest for students. One of the driving mechanisms is age, as copying the preceding person's food choice is the strongest among younger persons (cf. Section 5.3.2). Finally, we find that the mimicry diminishes as the proximity in the purchasing queue decreases, thus exhibiting a dose-response relationship (cf. Section 5.3.3).

5.4.2 Policy implications

We find evidence in favor of a specific behavioral mechanism for how dietary similarities between individuals occur. The evidence is based on a large-scale observational study monitoring many individuals for a long time in a natural setting. The behavioral mechanism of purchasing mimicry has implications for policies and interventions. Our findings imply that supplementary food items can be leveraged to increase or reduce the intake of specific foods and nutrients via food additions, not only meals [134]. The fact that we observe discrepancies between sub-population implies that policymakers should take these differences into account when designing the offering and offering interventions (students vs. staff). Efficiently targeting

subpopulations to achieve targeted effects resonates with previous work finding that policies aiming to encourage healthy eating need to be optimized for specific subpopulations [15].

Furthermore, our results indicate that behavioral interventions aiming to change diets should consider leveraging mimicry-based strategies when designing the dietary interventions. On the one hand, mimicry can be desirable when the offering is nutritious. On the other hand, mimicry can be undesirable when consuming the affected food items is not recommended. We find that mimicry is strong across the board for likely healthy items (such as fruit and salad) and potentially less healthy items (such as condiments, desserts, and soft drinks).

Therefore, concretely, to *promote* purchasing highly nutritious foods in on-campus environments, the additions to meals and beverages are a good place to intervene and ensure the availability of fresh, nutritious, and sustainable options, since their purchasing frequency is further boosted by mimicry. Previous work has found that such food availability interventions are effective in university environments [134]. Such desirable purchases can be further incentivized in sequential isles to boost purchases through the mechanisms of social norms.

Similarly, to *reduce* purchasing of calorie-dense, low nutrient foods, the additions are a good opportunity for intervening via point-of-purchase intervention strategies [96], since they are particularly affected by mimicry (strongest effect are measured for condiments, cf. Figure 5.3). Given the goal of reducing individuals influencing each other's purchasing behavior, future work should determine the effectiveness of interventions that aim to reduce the interactions, e.g., by enabling pre-ordering a meal through an application vs. deciding on the spot, since it is known that impulse buying is mediated by temporal proximity and making decisions in the proximity of others [331].

Such interventions to the offer should be explored in conjunction with manipulating food position (proximity or order) [50] since changing the order in which food items are presented at cafeteria counters has been proposed as a potentially effective way of altering food consumption [135].

Finally, our findings demonstrate that digital traces can provide valuable insights into the determinants of dietary choices. Digital traces can complement small-scale field experiments, making it possible to observe large populations over long periods. By studying behaviors as they occur naturally in a large population, our findings confirm and refine knowledge mostly derived from small-scale experimental studies. The approach leveraging passively-sensed purchase logs makes it possible to anticipate the impact of interventions before implementing them and identify the right subpopulations to target.

For instance, social influence in dietary habits has previously been examined in the context of school children [35, 119, 269, 318] and adolescents [91, 92, 346], who are theorized to be most susceptible to social pressures to diets and activity patterns [25, 317]. Although previous experimental studies found relationship type, gender, and age group not to be significant predictors of eating mimicry [30], a recurrent issue previous studies face is the small sample

size. Our findings relying on observations with a greater statistical power confirm the role of age since we find the effect to be the strongest in the youngest subpopulations and students.

5.4.3 Limitations

We now highlight potential limitations that put the bounds on the validity of our study.

Measurement errors. We note we study purchasing situations as they are visible in the transaction logs, which might not perfectly mirror the actual complete food consumption [264, 327]. A further source of measurement error is that the status estimation is imperfect. Individuals with demographic information are not a representative subpopulation of the complete campus population because they self-selected to participate in the sustainability challenge.

Further biases stem from the fact that purchasing behavior and choice mimicry might be driven by other unobserved factors, e.g., purchasing power, personal relationships, overall health and wellbeing, or calorie need. We perform sensitivity analyses to contextualize such unobservable factors by establishing how strong they would need to be to fully explain the observed effects. Lastly, our analyses monitor dyads, and future work should study more complex group dynamics beyond dyads that might occur in purchasing queues.

Population errors. Our study examines the behavior of those who eat in close temporal proximity to others and might therefore be different from those who only visit shops on less busy occasions and might not exhibit the mimicry patterns described here. Consequently, we only make claims regarding the studied purchase instances and the monitored individuals. Future work should determine the extent to which these results generalize beyond the university campus environment to the general population and standard settings where people make food choices while exposed to the choice of others (such as supermarkets, coffee shops, bakeries, food courts, and food trucks).

5.4.4 Future work

Future work should design and deploy on-site interventions to test the potential of behavioral nudges exploiting mimicry to promote healthy and sustainable eating on campuses. Lastly, it is important to note that behavioral mimicry is pervasive in human interactions beyond dietary behaviors. Moderators of mimicry include various motivational, social, emotional, and personality factors that lead to more or less mimicry in a given situation [59, 353]. Future work should focus on further understanding what factors drive the differences between the individuals depending on their age and status on campus.

Studying worldwide dietary behaviors with information seeking traces

Part III

6 COVID-19-induced shifts in dietary interests

6.1 Introduction

We now shift from studying diets in a campus-wide context to worldwide analyses focusing on the effect of COVID-19 on dietary behaviors. The coronavirus disease 2019 (COVID-19) pandemic has led to the implementation of unprecedented non-pharmaceutical interventions, including case isolation, social and physical distancing measures, business and school closures, travel restrictions, and full-scale national lockdowns [123]. For instance, in mid-May 2020, more than one third of the global population was under lockdown [222]. These interventions have caused important shifts in people's lives, which in turn created challenges that did not originate directly in the virus itself, but in the social, economic, and psychological implications of the population-scale measures taken to prevent the spread of the virus [42, 307], transforming education [112], exercise habits [74], mental health [352], online behaviors [115, 150], labor markets [157], transport, and mobility [42, 102], to name a few. Identifying how the pandemic has broadly impacted human needs and interests [311, 348] is therefore critical.

A thorough understanding of changes in food-related interests is particularly pressing, as changes in diet can have important ramifications for health, and dietary monitoring can help improve the well-being of populations. Diets are suspected to have become less balanced during the COVID-19 pandemic [188], and changes in diet and physical activity during the pandemic are known to increase the risk of cardiovascular disease [234] and are suspected to be associated with negative mood during lockdowns [185]. The implemented interventions disparately impact population segments within a country, depending on people's demographics, health, and habits [18, 186, 249]. Therefore, the pandemic can negatively impact the diet especially of those populations and individuals who are already most vulnerable [194], such as those affected by malnutrition [165, 266], eating disorders [165, 296], addictions [72], or obesity [29, 276]. Furthermore, in general, diet and nutrition are prominent factors in maintaining overall health and are important for developing a healthy immune response, which affects the speed of recovery and the probability of developing severe symptoms [183].

Public health and nutrition researchers and stakeholders have therefore issued a number of warnings about the potential nutritional public health issues that might emerge as a consequence, such as alcohol misuse [72] or weight-gain and obesity [186, 199]. There are concerns about the long-term implications of stress and boredom associated with lockdowns, as well as emotional eating [49, 255, 410], potentially linked with alcohol misuse and weight gain [72, 199]. However, it is not clear which aspects of the many potential adverse impacts of confinement on diets are most pressing, and on which of the many potential public health issues to focus first.

Beyond health, the question of COVID-19-induced shifts in dietary interests is also of economic importance [335]. It is necessary to understand emerging consumer needs, subsequent market readjustments [83, 166], and supply chain issues [155] that impact global access to food and food security [99, 208]. Many emerging customer behaviors are of interest to retailers and business owners during lockdowns, such as stockpiling, more frequent cooking, online purchasing, and changes in shopping locations [31, 62, 249, 298, 321, 382].

Early on in the course of the pandemic, anecdotal reports about changes in dietary habits during lockdowns emerged, e.g., about increased interest in baking [341, 394]. Existing research has studied the impact of COVID-19 stay-at-home orders on health behaviors and physical and mental health [231], finding initial evidence of increased sedentary behaviors and reduced physical activity [18, 190, 347, 390], less eating out, increased cooking and baking from scratch [122, 139, 283], and generally increased consumption of [27, 141], and interest in [211], food. Overall, food consumption and meal patterns were mostly found to be more unhealthy during confinement [139, 413], with the exception of a decrease in alcohol consumption [18, 413].

Current evidence, however, relies primarily on surveys and does not leverage passively collected large-scale observational data [62, 98, 100, 320]. It remains challenging to quantify shifts in food interests globally and holistically, across different types of food, and fundamental questions about food interests during the pandemic remain unanswered.

6.1.1 Research questions

The present chapter aims to bridge this gap by asking the following guiding question:

RQ How did dietary interests shift during COVID-19-induced mobility restrictions in 2020?

To address this question, we quantify people's change of interest in foods when they spend more time at home and how long these shifts in interests persist as mobility reverts to normal.

The fact that—unlike most previous events that directly impacted so many lives worldwide—the COVID-19 pandemic unfolded in a time of widespread Internet access allows us to conduct a population-wide infodemiology [111] study by relying on passively sensed digital trace data.

Specifically, we use time series capturing the popularity of Google search queries related to 1,432 foods (e.g., “bread”, “pizza”), as well as ways of accessing food (e.g., “recipe”, “restaurant”, illustrated in Figure 6.1), obtained in aggregated form via the publicly available Google Trends tool, to analyze changes in food-related interests across 18 countries. Google is the world’s largest Web search engine, and Google Trends search volumes have been shown to be a powerful population-scale sensor for numerous human behaviors, including unemployment [64], trading decisions [281], and voting [345]. We thus add to a rich literature that, well before COVID-19, has begun to analyze health and nutrition behaviors using digital trace data [15, 154], such as search engine logs [376, 392], purchase logs [12, 51], online recipes [297], reviewing platforms [65], social media such as Twitter [3, 88, 241, 395] or Instagram [261, 332], and geo-location signals [309].

6.1.2 Summary of main findings

Methodologically, drawing meaningful conclusions from the longitudinal Google search volume time series is challenging due to the presence of trends and seasonalities. We overcome these hurdles via quasi-experimental time series analyses (outlined in Figure 6.2), isolating the effect of the 2020 discontinuity in mobility patterns on food interests and going beyond simple correlations by accounting for 2019 baseline trends. This study design lets us identify the immediate, short-term increases in interest in all food types, which is found to be stronger and longer-lasting than those that coincide with end-of-year holidays (Figure 6.3a). The increased food interest is not uniform across types of food. The most prominent increases, in absolute and relative terms, occur for calorie-dense carbohydrate-based foods such as pastries and bread.

6.1.3 Implications

The identified shifts in interests, many of which persisted for months and some of which continued past our observation period (Figure B.12), represent a potential danger for public health and should be taken into account to inform decisions made by stakeholders in efforts to mitigate the effects of the COVID-19 pandemic on diets worldwide.

6.2 Materials and methods

6.2.1 Search interest time series

Our analyses rely on a curated and calibrated set of interest time series collected from Google Trends, an important tool for researchers [64, 142] that makes aggregate statistics about the popularity of search queries in the Google search engine publicly available. We collect time series of search interest in entities related to foods or ways of accessing foods. Search queries may be specified as plain text (e.g., “Cookie”) or as entity identifiers (e.g., “/m/021mn”) from

the Freebase knowledge base [41]. We use Freebase identifiers to conduct a multilingual study of interest since they allow for grouping various surface forms relating to the same topic. For instance, the entity “Cookie” (“/m/021mn”) captures “cookies”, “cookie”, “Cookie”, or “cookie jar”, etc., while the entity “Recipe” (“/m/0p57p”) captures all recipe queries across languages.

Google Trends provides time series of search interest for the specified input queries. Since search interest is not returned in terms of absolute search volume, but normalized by time and location and rounded to integer precision, we use Google Trends Anchor Bank (G-TAB) [391] to calibrate the time series. The benefit of calibration is that the interest is expressed on the same scale, and the combined interest in a set of entities can be estimated by adding up the interest in individual entities.

We collect interest data for two types of Freebase food entities: (1) entities related to the ways how people access food (such as “recipe”, “restaurant”), and (2) specific food entities (such as “cookie”, “pizza”).

1. Food-access mode entities: we curate entities that reflect ways of accessing food, starting from seed entities (recipe, take-out, restaurant, picnic), and inspecting related entities. Food-access mode entities are aggregated into four groups. Entities can be related to consuming food at home or outside of the home; orthogonally, entities can be related to consuming food prepared by persons within the household or food prepared by a third party (see Table B.3 for details about individual entities). We refer to the four groups of entities related to food:
 - (a) prepared within the household, consumed at home: recipe, cooking, baking, grocery store, supermarket
 - (b) prepared by third party, consumed at home: food delivery, take-out, drive-in
 - (c) prepared within the household, consumed outside: picnic, barbecue, lunchbox
 - (d) prepared by third party, consumed outside: restaurant, cafeteria, cafe, diner, food festival
2. Foods entities: we start from ids of food entities from Freebase. These are entities of type “food”, “dish”, “beverage”, or “ingredient”. Food entities are aggregated into categories. Category creation: we enrich Freebase entities with Wikidata knowledgebase [377] properties using the Wikidata query API. For each Freebase entity id, we query Wikidata with the Freebase id to get its “instance of” or “subclass of” properties. We derive a taxonomy of 28 categories based on “subclass of” and “instance of” relations. To ensure that the food classes are general and representative, we keep all classes with at least ten entities. Note that not all entities have a “subclass of” or “instance of” field available in Wikidata and therefore cannot be automatically categorized. To achieve higher coverage, we manually annotate a set of popular entities. We monitor global time series of all food entities in 2019–2020. We select the top entities that covered 95.7% of global food search volume and annotate all such entities that do not already have a category derived based

on Wikidata. This process resulted in a set of $N=1432$ entities, categorized either based on Wikidata or manually. Categories are presented in Table B.1. An author who is a professional epidemiologist specialized in nutrition assessed and refined the entities and the corresponding categorization.

6.2.2 Search interest data collection

Overall, we collected time series for 1,432 food entities and 16 food access mode entities in 18 countries, spanning from January 1, 2019, to December 31, 2020, at weekly granularity. The time series were collected and calibrated with the Google Trends Anchor Bank library. The full list of 1,432 food entities and 16 food-access mode entities is available in our data repository. The 1,432 food entities are categorized into 28 food categories, and the 16 food-access mode entities are categorized into four groups. After data collection, we obtain country-specific time series for 28 food categories, and four aggregate food-access mode by adding up time series of respective individual entities.

The countries were pre-selected with the goal of achieving global coverage across continents, studying countries with a large number of Internet users [184], and including countries with varying severity of mobility restrictions. Additionally, we collected global time series of interest in the “recipe” and “restaurant” entities, spanning from the beginning of 2019 until the end of 2020, at weekly granularity, in all countries and territories (Figure 6.1).

6.2.3 Mobility time series and COVID-19-induced mobility decreases

To capture variation in the mobility of the populations in the 18 studied countries, we use mobility reports [13] published by Google, which capture population-wide movement patterns based on cellphone location signals. We use country-wise mobility data from February to the end of December 2020. The mobility reports specify, for each day, by what percentage the time spent in residential areas differed from a pre-pandemic baseline period in early 2020.

We chose to rely on mobility data and not the official start of lockdown dates. The problem with employing the official start of lockdown date in statistical analyses is that it is not guaranteed that they would impact movement patterns across different countries homogeneously (e.g., it could be that for some of the countries people stayed more at home even before the lockdown was enacted). Similarly, the official lockdown date might vary within a country.

We automatically detect changes in the mobility time series caused by both government-mandated lockdowns as well as self-motivated social distancing measures [291]. We refer to these points as mobility changepoints. We use mobility changepoints as heuristic dates for when people started or stopped spending substantially more time in their homes. Unlike choosing one of the official dates of lockdown implementation or relaxation, this leads to a meaningful onset of decreased mobility across different countries.

Figure B.1 depicts three important mobility changepoints dates that occur at different moments throughout 2020 in the studied countries:

1. The first sharp mobility decrease occurring in March and April 2020 when people started to spend substantially more time at home.
2. The mobility increase occurred as people stopped spending substantially more time at home.
3. The second mobility decrease occurred between October and December 2020 (occurred in some of the studied countries), when people started spending substantially more time at home during the second wave of the pandemic.

We detect the three changepoints for each country independently by smoothing and thresholding: we consider the weekly rolling average mobility. We monitor the percentage of time spent at home. The first date when time spent at home increased by 10% is the start of reduced mobility in the first wave. We repeat the same to detect the onset of the second wave. In this way, the period when percentage of time spent at home consistently stays above 10% compared to pre-pandemic baseline (defined as pre-pandemic mobility levels by Google) is a period of decreased mobility, in the first, or in the second wave. Note that the period of decreased mobility is very short in Sweden, the country with no government-mandated mobility restrictions. Sweden is still included in our analyses as a contrasting case.

The three changepoint dates are marked in Figure B.1 in the 18 studied countries. The first mobility decrease, and the second mobility decrease (in case it occurs) serve as cutoff dates in our modeling approach. The date of the mobility increase serves to limit the possible duration of the studied period with decreased mobility.

6.2.4 Modeling approach

To estimate the potential effects of the sudden mobility changes on food interest time series, we devise a regression discontinuity design (RDD) with a local regression in time. Additionally, we incorporate a fake discontinuity separating before vs. after the cutoff date in 2019, the year before the pandemic, to account for seasonal trends. The model of a given interest time series in a given country is given by the following regression discontinuity design (RDD) in quadratic form:

$$\begin{aligned}
 \log y_{iT} = & \alpha' + \beta' \cdot t + \gamma' \cdot t^2 \\
 & + \alpha'' \cdot i_t + \beta'' \cdot i_t t + \gamma'' \cdot i_t t^2 \\
 & + \alpha''' \cdot j_T + \beta''' \cdot j_T t + \gamma''' \cdot j_T t^2 \\
 & + \alpha \cdot i_t j_T + \beta \cdot i_t j_T t + \gamma \cdot i_t j_T t^2,
 \end{aligned} \tag{6.1}$$

where T is the year (2019 or 2020); t is the week in the year relative to the week in which the discontinuity occurred in 2020 (but not in 2019), for $t \in [-t_{\min}, t_{\max}]$; $t_{\min} = 10$; y_{tT} is the calibrated (see above) search interest volume in week t of year T of an entity (or set of entities) in the respective country; i_t is a binary variable equal to 1 if $t > 0$ and 0 otherwise; and j_T is 1 in 2020 and 0 in 2019. This way, for all weeks where $i_t = j_T = 1$, a unit is “treated”, otherwise it is not. Logarithmic outcomes are used in order to make the model multiplicative. The outcome is modeled as a separate quadratic function of time before and after the discontinuity in order to capture nonlinear temporal patterns. By comparing observations lying closely on either side of the temporal threshold, we estimate the treatment effect while minimizing potential bias from unobservable confounders.

The interaction coefficients α, β, γ model the effect of the discontinuity, controlling for baseline trends in 2019. The short-term increase in interest is captured by the fitted coefficient α , which estimates the short-term effect of the mobility decrease on search interest. The approach is outlined in Figure 6.2.

With this model, we measure the time-dependent trends because the model is expressive enough (i.e., quadratic terms capture the temporal evolution, see illustrations in Figure B.5). We also provide the main results with constant and linear models in Appendix B. Bandwidth choices are made in the following way: $t_{\min} = 10$, since it is the maximum number of weeks in 2020 before the cutoff, $t_{\max} = 30$, since it is the maximum number of weeks we can have so that across all the studied countries, the second mobility decrease shock is not included. We also investigated the impact of the choice of the bandwidth (see Appendix B).

In our analyses, we fit a model of this general form (Equation 6.1) to interest time series, separately for each studied entity or groups of entities, in each of the studied countries. We use the modeling approach to investigate three key quantities illustrated with the example of interest in pastries and bakery products in Brazil, and Australia in Figure 6.2:

1. Short-term increase in interest, captured by the fitted coefficient α . The model is multiplicative due to the logarithm. After fitting the model (Equation 6.1) with OLS, the relative increase over the baseline is then calculated by converting α back to the linear scale, via $e^\alpha - 1$; the 95% CIs (approximated with two standard errors) are also converted back to linear scale.
2. Time it takes for the interest to revert to normal. We measure how many weeks after the mobility decrease (within the $t_{\max} = 30$ weeks) the modeled interest in 2020 is no longer significantly different from the counterfactual prediction based on 2019 (based on non-overlapping 95% CIs).
3. Long-term increase in interest. In case the interest did not go back to normal within the 30 weeks after the mobility decrease, we measure how elevated the interest remains at the end of the modeled period, 30 weeks after mobility decrease, compared to the interest in the same week in 2019.

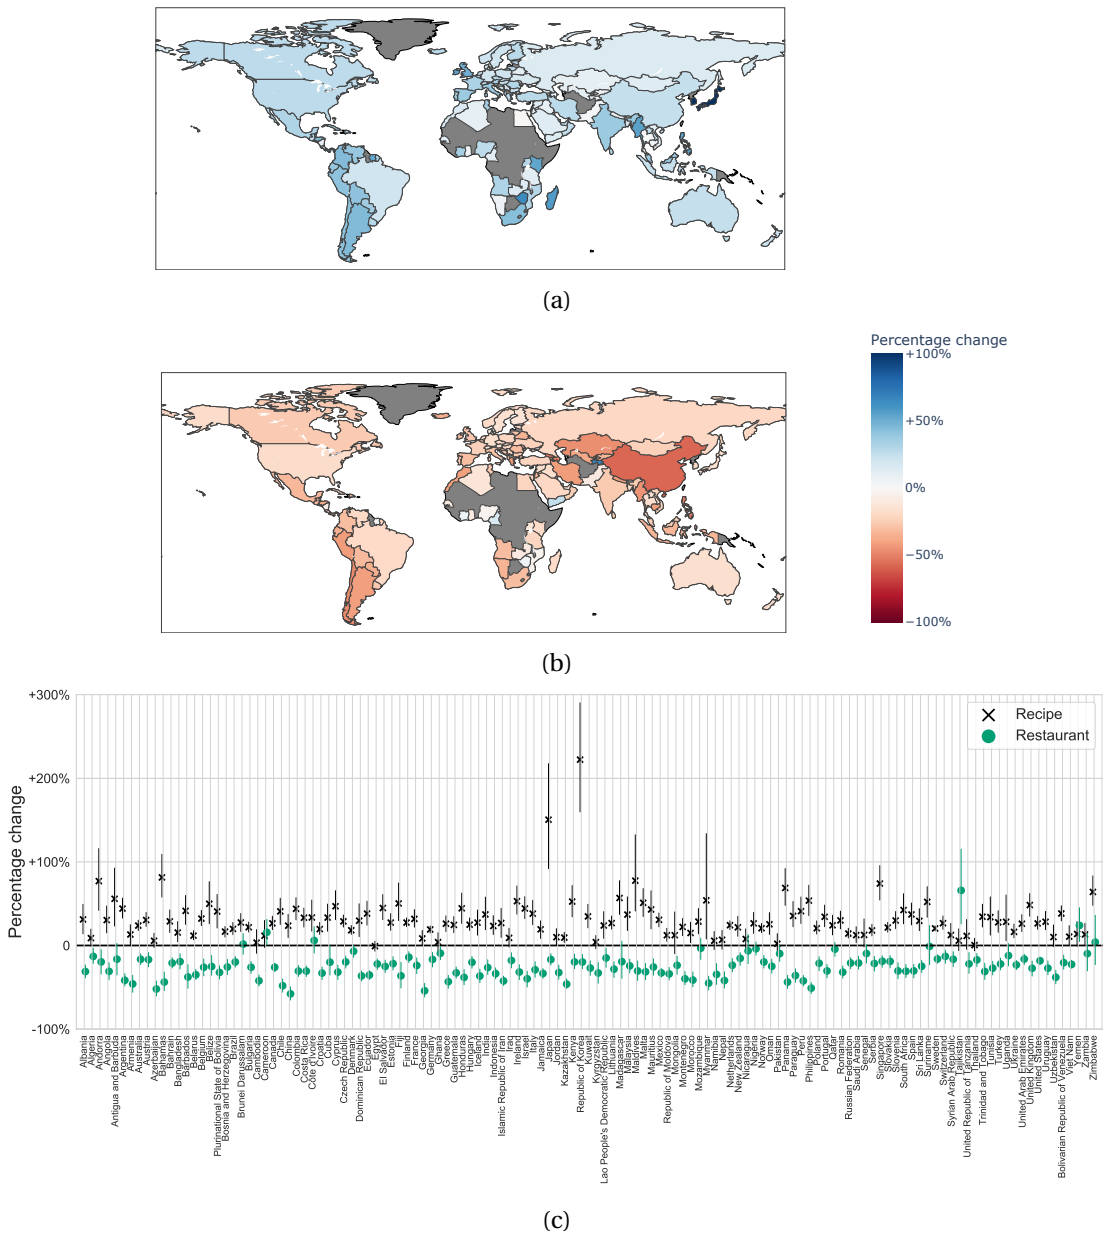


Figure 6.1: Global shifts in dietary interests. In (a), search interest in the concept of “recipe” across countries, in 2020 vs. 2019. In (b), search interest in the concept of “restaurant” across countries, in 2020 vs. 2019. In (a) and (b), color marks average relative change in interest in 52 weeks of 2020 in a country, compared to corresponding weeks of 2019. Countries with not enough search data are marked in gray. In (c), average relative change in interest in 52 weeks of 2020 in a country, compared to corresponding weeks of 2019. Points mark average relative change (across $n = 52$ weeks), illustrated on maps in (a) and (b). Error bars mark 95% confidence intervals. During 2020, compared to 2019, there was a global increase in interest in recipes, and a global decrease in interest in restaurants.

6.3 Results

Before describing the results, as a reminder, we briefly summarize the data to be analyzed. In total, we curated a set of 1,432 entities related to specific foods (e.g., “bread”, “pizza”) grouped in 28 food categories (details in Section 6.2, Figure 6.2), which covered 95.7% of the global food search volume in 2019 and 2020. Table B.1 summarizes the descriptions of food categories and contains examples of popular foods in each category. We also curated a set of 16 different entities related to ways of accessing food (e.g., “recipe”, “restaurant”), grouped in four categories: entities can be related to consuming food at home or outside of the home, and orthogonally, entities can be related to consuming food prepared by persons from within the household or food prepared by a third party (Table B.3).

As a first glimpse into the food entities and fluctuating interests, we estimate interest in recipes and restaurants, globally across 129 countries with enough search data available to estimate

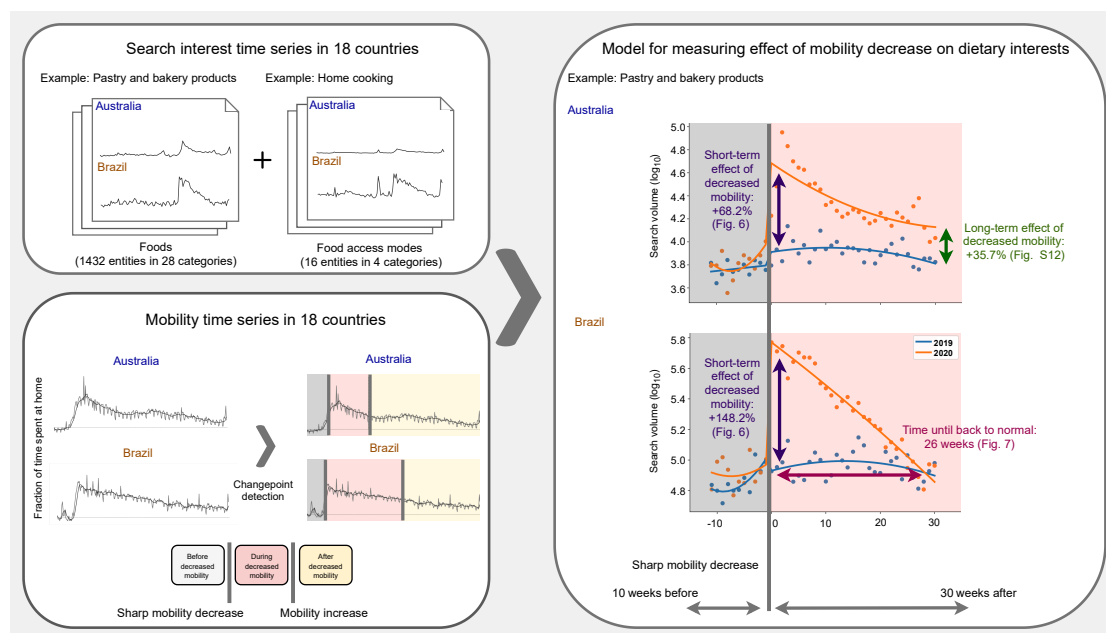


Figure 6.2: Study design. We start from interest time series in 18 countries, capturing search interest in food entities and in entities about ways of accessing food. In order to measure the effect of the changes in mobility time series on interest, we first detect mobility changepoints (the abrupt mobility decrease and the eventual mobility increase) via changepoint detection. On the right, we illustrate the modeling approach on an the example of interest in pastry and bakery products in Australia and Brazil, where on the x-axis is the week relative to the week of the mobility decrease, and on the y-axis is search interest. The modeling approach measures the effect of the shock of mobility decrease on dietary interests, controlling for pre-pandemic trends. With this model, we measure three key quantities: the short-term effect of decreased mobility, the time until interest reverts back to normal, and the long-term effect of increased mobility.

weekly interest volumes. Figure 6.1 illustrates the emerging shifts: a global increase in interest in recipes (Figure 6.1a), and a global decrease in interest in restaurants (Figure 6.1b) during the weeks of 2020, compared to the corresponding weeks of 2019.


We collected search interest time series in 18 countries. The studied countries were selected to achieve geographic diversity and comprise countries with a large number of Internet users across continents: Brazil, Canada, Mexico, United States, France, Germany, Italy, Spain, United Kingdom, India, Indonesia, Japan, Egypt, Kenya, Nigeria, and Australia. Additionally, to achieve a varying severity of lockdowns, Sweden and Denmark were added as contrasting cases, due to particularly lenient COVID-19-induced restrictions [159]. The interest time series were collected from the Google Trends platform [64, 142] and calibrated with Google Trends Anchor Bank [391] (so time series for different search queries can be aggregated via summation and compared with one another). Although absolute search volume—the number of issued queries—cannot be inferred, calibration can infer absolute search volume up to a constant multiplicative factor. This way, ratios of absolute search volumes can be validly estimated when working with calibrated Google Trends time series. The interest time series in the same regions in 2019 serve as baselines.

Note that, although different languages are spoken in the 18 studied countries, search queries did not need to be translated, as Google Trends allows language-independent entity descriptors from the Freebase knowledge base [41] as input. For instance, given as input the Freebase entity descriptor for “bread” (/m/09728), Google Trends will return the search interest for all queries related to the concept “bread” across languages.

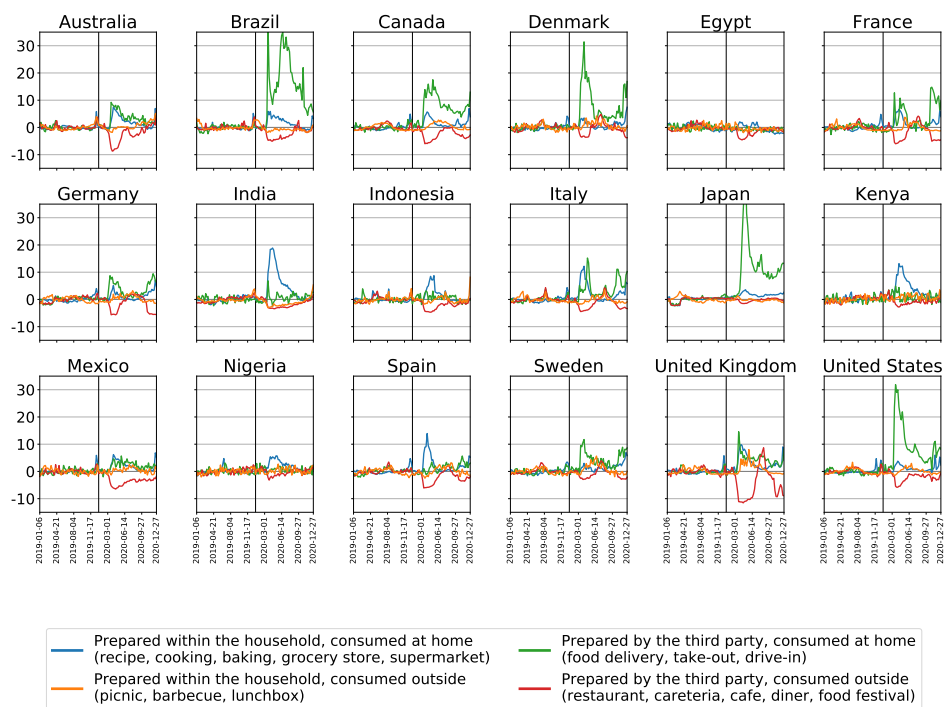
6.3.1 Overall surge in food interest larger than during end-of-year holidays

We examine how the total interest in food entities evolved in 2019 and 2020 (Figure 6.3a). We monitor interest in all food entities, normalized by the 2019 mean and standard deviation (z -scores). We refer to this quantity as the surplus of interest. Normalizing food interest allows us to quantify the surplus of food interest in a week relative to the Christmas week of 2019. In a given week, the surplus relative to the Christmas week is measured as the ratio between the z -score in the observed week and the z -score in the Christmas week.

First, note the peaks of food interest during the end-of-year holiday season in both 2019 and 2020. Second, note the increase in overall interest in food entities coinciding with the reduced mobility due to COVID-19 occurring in March 2020. These rises of food interest are larger in amplitude compared to the rises of interest during end-of-year holidays, and they last longer. For example, in the US, the surplus (compared to the 2019 mean) of food interest at its peak during the first wave of the COVID-19 pandemic equals the surplus of interest during the Christmas week of 2019, as well as that of the surplus of interest during the Thanksgiving week of 2019. In total, the surplus of food interest in the first six months of 2020 in the US is 9.0 times higher than the surplus of interest during the Christmas week of 2019, and 8.9 times higher than the surplus of interest during the Thanksgiving week of 2019.



(a)



(b)

Figure 6.3: Temporal evolution of dietary interests. In (a), total interest in food entities (z-scores), 2019-2020. Dashed line marks 10 standard deviations above the 2019 mean. In (b), interest in ways of accessing food (z-scores), 2019-2020.

We next compare the surplus of food interest at its peak during the first wave of the COVID-19 pandemic with the surplus of interest during the Christmas week of 2019 across countries. We exclude countries with a non-Christian majority (India, Indonesia, Japan, Egypt, and Nigeria), where there are no prominent increases in food interest during Christmas week (Figure 6.3a). When comparing to Christmas holidays, the surplus of food interest at its peak during the first wave of the COVID-19 pandemic is on average 1.9 times higher than the surplus of interest during Christmas week, while the total surplus of interest in the first six months of 2020 is on average 18.8 higher than the surplus of food interest in the Christmas week of 2019.

The increases in food interest are drastic in India, Indonesia, and Nigeria, too, with food interest at the peak of mobility restrictions surpassing 10 pre-pandemic standard deviations. Note that Sweden, Denmark, and Japan, the countries with the mildest government-mandated mobility restrictions [44, 159, 263, 402], contrary to all other studied countries, had no notable overall increase in food interest in 2020 (Figure 6.3a).

Next, similarly, in Figure 6.3b, we examine the temporal evolution of the interest in the four modes of accessing food reflecting whether they relate to consuming food at home or outside of home, and orthogonally, whether they are related to consuming food prepared by persons within the household or food prepared by a third party. In all countries, in 2020, there was a decrease in interest in food prepared by third parties and consumed outside (in red) and an increase in interest in food prepared within the household and consumed at home (in blue) coinciding with the onset of the first wave of the COVID-19 pandemic in the first half of 2020.

Compared to the end-of-year holidays, the surplus of interest in food prepared within the household and consumed at home (recipes, cooking, baking, grocery stores, and supermarkets) was at the peak 1.7 times higher than the surplus during the Christmas week of 2019, on average across countries (excluding countries with a non-Christian majority; cf. above). In the first six months of 2020, it was in total 13.7 times higher than the surplus of interest during the Christmas week of 2019. The increases in interest in recipes, cooking, baking, grocery stores, and supermarkets relative to the 2019 mean were large in India as well, surpassing 10 pre-pandemic standard deviations at the peak. Additionally, we note large increases in interest in food prepared by third parties and consumed at home (in green). In the US, Brazil, Japan, and Denmark, this interest increased by more than 30 pre-pandemic standard deviations at the peak.

6.3.2 Changes in food interests are strongly associated with mobility

Next, we combine search interest time series with mobility data published by Google (described in Section 6.2) which captures the relative increase in time people spend indoors compared to a pre-pandemic baseline. We find that interest in different ways of accessing foods and interest in specific foods are strongly correlated with mobility during the COVID-19 crisis (Figures 6.4a and 6.4b).

Across weeks in 2020, from February to the end of December 2020 (the period for which mobility data is available), we calculate the country-specific Spearman rank correlation between mobility time series and food interest time series. Here, in order to adjust for seasonal trends, the food interest for a given week of 2020 is expressed as the relative increase compared to the corresponding week of 2019.

We observe strong and significant associations between food interests and mobility. Interest in recipes (Figure 6.4b) is positively correlated with spending more time at home ($p < 0.05$ in all countries except Japan; Spearman's rank correlation coefficient ranging between 0.36 in Egypt and 0.95 in Mexico), and takeout is significantly and positively correlated in 13 out of the 18 studied countries (Spearman's rank correlation coefficient ranging between 0.34 in Indonesia and 0.82 in Australia). Interest in restaurants, on the other hand, is negatively correlated with spending more time at home, significantly in all studied countries (Spearman's rank correlation coefficient ranging between -0.40 in Kenya and -0.97 in Italy).

Regarding food categories (Figure 6.4a), despite some variation between countries, there are notable food categories that have a significant positive correlation with spending more time at home in most of the studied countries, such as desserts (ranging between 0.40 in Sweden and 0.84 in Brazil) and pastries and bakery products (ranging between 0.46 in Denmark and 0.88 in the UK).

The correlation between mobility and food interest normalized by the 2019 baseline is measured for individual entities (Table B.3). All entities related to consuming food at home are correlated positively on average over countries, whereas all entities related to consuming food outside of home are correlated negatively on average (except barbecue, likely due to the fact that barbecue food can also be consumed at home). Among specific foods, the strongest positive correlation is found for pancake, baking powder, bread, baker's yeast, cookie, chocolate brownie, chicken meat, chocolate cake, biscuit, and pasta. The strongest negative (although much smaller) correlation is found for foods such as tapas, Korean barbecue, sushi, and gelato, typically eaten in social contexts taking place outside of home.

In the analyses so far, we have examined the response of the interest as the mobility changed by measuring correlation. Next, given the abrupt nature of the change in mobility, we isolate the effect of the shock of mobility decrease on food interest via a modeling approach.

6.3.3 More interest in home food, less interest in out-of-home food

As depicted in Figure 6.2, to isolate the shock of the mobility decrease occurring in all studied countries in March 2020, we first automatically detect changes in the mobility time series caused by government-mandated lockdowns or self-motivated social distancing measures (Section 6.2). We refer to these points as mobility changepoints (Figure 1 B.1).

To measure the effect of decreased mobility on food interest time series, we employ a quasi-experimental design that isolates the impact of the mobility decrease shock (the discontinuity),

Chapter 6. COVID-19-induced shifts in dietary interests




Figure 6.4: Spearman rank correlation coefficient between mobility decrease and interest volume. In (a), correlation for categories of food entities, and in (b), for ways of accessing food. For each group, $n = 18$ values represent correlation coefficients for the 18 studied countries (calculated based on $n = 46$ samples corresponding to weeks of 2020). The boxplot summarizes the values across 18 countries. Significant correlations ($p < 0.05$), according to a two-sided t-test (with no correlation between interest and mobility as the null hypothesis) are marked in blue, and non-significant correlations in orange. No adjustments for multiple comparisons are made. Boxplots represent the 50th (center line), 25th and 75th percentile (box limits). The whiskers extend to the minimum and maximum values but no further than 1.5 times IQR.

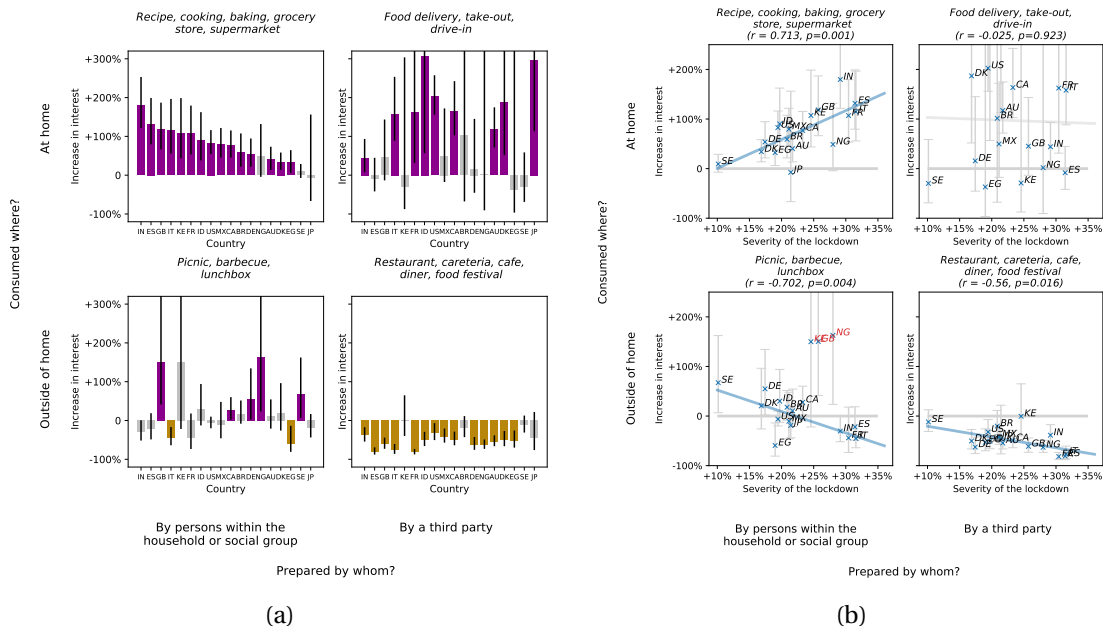


Figure 6.5: The short-term effect of mobility decrease on interest in accessing food. In (a), the short-term effect of the shock of mobility decrease on interest in accessing food, estimated with our RDD-based model. For each group of entities ($n = 4$), and each country ($n = 18$), the model (Equation 6.1) is fitted on $n = 82$ samples representing weekly interests. Bars represent effect estimates (fitted coefficient α). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects according to a two-sided t-test. Fitted coefficients and statistics are presented in Table B.4. In (b), the relationship between the severity of the lockdown, measured as the increase in the percentage of time spent at home at the peak of reduced mobility (x-axis) and the estimated short-term effect on interest (y-axis), across four groups of entities about ways of accessing food. For each way of accessing food ($n = 4$), Pearson correlation coefficient is reported (calculated over $n = 18$ countries) and p-value according to a two-sided t-test with no correlation between severity of lockdown and short-term effect as the null hypothesis. The short-term effect of the shock of mobility decrease on interest is estimated with our RDD-based model and presented with 95% confidence intervals. The straight blue line is a least square fit. Red country codes mark countries that are considered outliers. No adjustments for multiple comparisons are made.

controlling for patterns occurring in the same weeks of 2019 when COVID-19-induced mobility restrictions did not occur (Figure 6.2).

We find that in 15 out of the 18 studied countries (Figure 6.5a), there was a significant short-term increase in interest in food prepared within the household and consumed at home (e.g., recipes), with short-term boost in interest (α) ranging between +32.2% in Egypt and +179.8% in India. In 14 countries, there was a significantly decreased interest in food prepared by third parties and consumed outside of the home (e.g., restaurants), ranging between -32.1% in USA and -81.7% in France. There were major increases in interest in food prepared by third parties

and consumed at home (e.g., food delivery), with more than a +100% significant increase in eight of the 18 studied countries.

We next analyze the relationship between the amplitude of the short-term changes in dietary interest and the severity of lockdowns (Figure 6.5b), where the severity of a lockdown is defined as the percentage change of the fraction of time spent at home (with respect to the pre-pandemic baseline level) at the peak of reduced mobility [291]. The severity of a lockdown varies between +10.1% in Sweden, the country with no government-mandated mobility restrictions [44], and +31.6% in Italy, a country with severe lockdown measures [40]. All peaks of mobility decrease occurred between March and May 2020.

We find that the more drastic the lockdown severity, the more drastic the change in dietary interests. Changes in interest in food-access modes have an association with the severity of the lockdown: positive for food prepared within the household and consumed at home ($R = 0.71$, $p = 0.001$), and negative for consumption outside of the home, i.e., food prepared by third parties and consumed outside ($R = -0.56$, $p = 0.016$) and food prepared by third parties and consumed at home ($R = -0.70$, $p = 0.004$). Here, the UK, Kenya, and Nigeria were excluded because they are clear outliers. When not excluding these three countries, we still observe a negative, but non-significant correlation ($R = -0.11$, $p = 0.676$). The discrepancy between the UK, Kenya, Nigeria and the other countries might be linked to COVID-19 policies allowing congregation in open green spaces, including parks and beaches [192]. Note that Sweden, the country with no government-mandated mobility restrictions [44], had no significant shift of interests in ways of accessing food.

The fact that the effect of decreased mobility on the interest in recipes and similar entities across countries rises linearly with the severity of lockdown adds to the evidence that interests changed because mobility decreased. If there were other confounding factors that could explain the changes in dietary interests, and those factors had nothing to do with the shock of the mobility decrease, we would not expect to find such a clear dose-response relationship. Instead, we would need to envisage a more complex effect of an unobserved factor that could impact both the strength of the lockdown in a country and cause changes in the population's dietary interests, in ways that have nothing to do with spending more time at home.

Although significant increases in interest in food prepared by third parties and consumed outside of the home also exist, they are not correlated with lockdown strength. Presumably other factors are at play, such as the response of the market, availability of delivery companies, or how quickly restaurants adapted to do deliveries.

6.3.4 Drastic increases in interest for calorie-dense, carbohydrate-based foods

Having established the link between the sudden decrease in mobility and the shifting interests in ways of accessing food, we next examine how exactly the interest in specific types of food varied (Figure 6.6). Is the observed increase in food interest uniform across all food types, with



Figure 6.6: The short-term effect of the shock of mobility decrease on interest in categories of food entities. The effect is estimated with our RDD-based model. For each category of entities ($n = 28$), and each country ($n = 18$), the model (Equation 6.1) is fitted on $n = 82$ samples representing weekly interests. Bars represent effect estimates (fitted coefficient α). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects according to a two-sided t-test. Fitted coefficients and statistics are presented in Table B.4. Food categories are sorted by median effect across countries. No adjustments for multiple comparisons are made.

interest in all foods increasing proportionally, or does interest in certain foods increase more? We apply the modeling approach (Figure 6.2) on time series capturing interest in the 28 food categories in the 18 countries and measure the short-term effect of decreased mobility.

Overall, we find that there was a momentary increase of total food interest (gray bands in Figure B.10), significant in all countries except Sweden and Japan, ranging between +24.6% in Denmark and +99.4% in Spain. Similarly, there was an increase in interest in most of the individual food categories (Figure 6.6). The biggest increases, however, occurred for calorie-dense, processed, carbohydrate-based foods: pastry and bakery products, bread and flatbread, and pie. These effects are significant in most of the countries. Especially strong cases (with

Chapter 6. COVID-19-induced shifts in dietary interests

increases of over 200%) include pastry and bakery products in Spain, France, Canada, and Egypt; bread and flatbread in Spain, France, and Italy; and pie in Spain.

We observe smaller increases for other categories, including fresh produce (vegetable, fruit, salad, herb), meat and fish dishes (chicken, pork, fish, beef, lamb dishes), and wine, beer, liquor and cocktails, which saw an increase in some of the countries. These conclusions and the relative ranking between categories are robust to specific modeling choices (Table B.5).

In the Appendix B (Figure B.9), we additionally provide an alternative analysis where the outcome variable is the relative volume share (i.e., the fraction of the total weekly food interest that is allocated to the respective search queries), rather than absolute volume as analyzed above. This way, we control for the overall increased food interest. In terms of the share of interest, the most prominent increases again occurred for pastry and bakery products (over 50% increase in share fraction in 11 of the 18 countries) and bread and flatbread (over 50% increase in share fraction in six of the 18 countries), whereas the share of interest in other food categories remained robust or decreased slightly. The most prominent decreases in terms of the share of interest occurred for soft drinks, alcoholic drinks, and sandwiches—beverages and food presumably typically consumed in social contexts taking place outside of home.

Although most food categories saw increased interest in most countries, there are specific foods where interest decreased as mobility decreased, such as tapas and energy drinks (Table B.3).


We next measure the time it took for search interest to revert to normal, illustrated in Figure 6.2 for the example of Brazil. We measure how many weeks after the mobility decrease it takes until the modeled interest in 2020 is no longer significantly different from the counterfactual prediction based on 2019 (based on non-overlapping 95% confidence intervals). In addition to being large in amplitude, we observe that the numerous shifts in interests lasted for months. For instance, the shortest duration of increased interest in food prepared within the household and consumed at home was 6 weeks (in Egypt; Figure 6.7a), and the shortest duration of increased interest in specific food categories, nine weeks (wine, beer, and liquor in France, pastry and bakery products in Denmark, and sandwiches in Sweden; Figure 6.7b). Most of the changes in interest in specific groups of food entities are transient, and the interest returned to normal within 30 weeks.

In cases where interest did not return to normal within the 30 weeks after the mobility decrease, we measure (in Figure B.12) how elevated the interest remains at the end of the modeled period, 30 weeks after the mobility decrease, compared to the interest in the same week in 2019 (illustrated in Figure 6.2 using the example of Australia). While most interests come back to normal within 30 weeks, there are some notable exceptions of more permanent changes (Figure B.12): the interest in food prepared by third parties and consumed at home (e.g., food delivery) permanently increased in Italy, Canada, the US, Australia, Denmark, and Japan, while the interest in food prepared within the household and consumed at home (e.g., recipes) also remained elevated in Spain, the UK, the US, and Australia.

prepared within the household, consumed at home			27	29	19	26	20	29	27			29		29	6	17		
prepared by the third party, consumed at home			14															
prepared within the household, consumed outside					9												22	
prepared by the third party, consumed outside	18	18	14		16				29	29	14		18	14	26	10		
	ES	GB	FR	KE	IT	IN	ID	BR	CA	US	NG	MX	AU	DE	EG	DK	SE	JP

(a)

pastry and bakery product	28	29	27		19	12		25	27	29	26				9	16		20
bread and flatbread			21		20		29	30			28					26		
pie	27	28	25		25	14		28	23	29		30			12	18		
potato dish	19		26		22		18					29						
stew	14	26		10		20			25	28						20		
sauce	29	30	28	21		28	29	28				28			23	27		
cheese	28	28		24	24	27	23	28				28			25			
dessert	29	28	21		22		27	27	29	29					25	20		
chicken dish		29	27	24		27	12		29	30		29						
pasta, pizza and noodle dish	30		28	26	19	30	27	26	30	28		26			16	21		
pork dish	25	30	26					26				28						
sausage	19	28	25			27	19	28		29		28						
egg dish	29		27				29				27							
rice dish	26		28	22		28	29					29						
vegetable and legume	29	27	27	27		27		29	28	27		29						
fruit	28		22	26	24	28	28	28	29	30	26							
fish dish																		
soup	16	25	15			28			14	27	21	24						
spice	30	29	26			29				25								
beef dish	28					17		29				23						
snack			22			22		13	26									14
salad																		
herb				26		30			28	29							29	
sandwich						21											9	
wine, beer and liquor		26	9	17				23										
cocktail									22									
soft drink		27									19							
lamb dish			12					20										
	ES	GB	FR	KE	IT	IN	ID	BR	CA	US	NG	MX	AU	DE	EG	DK	SE	JP



(b)

Figure 6.7: Number of weeks until food interest goes back to normal. In (a), the number of weeks for ways of accessing food, and in (b) for food categories. The number of weeks is determined by measuring how many weeks after the mobility decrease, the modeled interest in 2020 is no longer significantly different from the counterfactual prediction based on 2019 interest, based on non-overlapping 95% CI. Gray marks that there were no significant differences in 2020 compared to 2019, and white marks that the interests did not go back to normal until the end of 2020.

6.3.5 Second wave had less impact on food interests

Finally, we explore the effect of the “second wave” of the pandemic, which occurred between October and December 2020 in the UK, Canada, Italy, US, France, and Spain. In Figure B.11a, we observe much smaller effects in the second wave compared to the first wave of mobility decrease. No significant increases in interest in food prepared within the household and consumed at home (e.g., recipes) are observed. While mobility saw large changes in some countries (Figure B.1), no drastic changes in food interests occurred.

Notable exception are France and Italy, where food prepared by third party and consumed at home saw large increases in interest in the second wave. We observe significant decreases in interest in restaurants in the UK, Italy, and France, but smaller effects compared to the first wave. Finally, no notable surges in interests in bread, pastry, baking, and desserts as in the first wave are observed in the second wave. The second wave brought on less drastic mobility decreases and was less of a disruption. Additionally, as populations adapt and acquire new skills, there might be less of a need for recipe searching over time.

6.4 Discussion

In order to formulate policies and allocate resources for mitigating the adverse nutritional impacts of the COVID-19 pandemic, governments, organizations, non-governmental organizations, and other stakeholders need reliable and timely data regarding the circumstances faced by affected populations. The results presented here document the impacts of confinement on nutritional interests. As the pandemic continues to unfold, warnings about potential public health issues emerge. In this study, we aim to point out prominent emerging behaviors by quantifying initial developments and providing a broad grounding for future studies.

6.4.1 Implications

Implications for public health. From a public-health perspective, the emerging surge in food interest during confinement is concerning. During the first wave of the COVID-19 pandemic, there was an overall surge in food interest, stronger and longer-lasting compared to the end-of-year holiday season of 2019 (Figure 6.3a). Since Christmas and Thanksgiving are known to be disruptive to dietary habits and a hazard to balanced diets [180], and the effects of confinement on food interests are comparable in amplitude and last longer, there is a pressing need to understand them.

In addition to the overall volume, the nature of the food interest also changed. After the shock of the mobility decrease, there was a large immediate increase in interest in consuming food at home and a decrease in consuming food outside of home. In 12 of the 18 studied countries, as mobility decreased, the interest in preparing and consuming food at home momentarily increased by more than 50% (Figure 6.5a), and the interest in baking and pastries more than

doubled (Figure 6.6). Since such modified interests persisted for a prolonged period (at least nine weeks, Figure 6.7) and since frequent consumption of meals prepared away from home is significantly associated with an increased risk of all-cause mortality [104], preparing more meals at home is a potentially positive side of the shifts in interest and should be understood further from a public-health perspective.

However, the sharply increased interest in potentially calorie-dense foods is worrisome. Overall, we find that the most drastic increases in interest are in carbohydrate-rich foods. These surges are not matched by proportional increases in interest in fresh produce, meat meals, vegetables, or fruit. Such shifts represent a danger of favoring processed and calorie-dense foods, at times when physical activity is reduced. This is particularly concerning from a population-scale well-being and mental-health point of view. These results call for developing a deeper understanding of the exact mechanisms in how stress, boredom, and emotional eating associated with the lockdown may have contributed to the observed effects [49, 255, 410]. It is also necessary to understand further how changes in product availability (e.g., a shortage of fresh ingredients including milk, eggs, and flour [240] in many countries) and changes in consumer behaviors (including the emergence of stockpiling [382]) are linked with the increases in interest in carbohydrate-rich foods.

Implications for consumer behavior. Figure B.12 hints at permanent small increases in interests in certain foods. While the interest in restaurants came back to normal in the studied countries except India, Indonesia, and Mexico within 30 weeks after the shock of the mobility decrease, interest in takeout remained increased in Italy, Canada, the US, Australia, Denmark, and Japan, and interest in recipes in Spain, the UK, the US, and Australia also remained increased. Future work should determine if these are new permanent habits brought on by the pandemic, or if they will fall back to normal in the future. These findings are particularly important to take into account in efforts to understand market readjustments.

Implications beyond COVID-19. Outside of the ongoing COVID-19 pandemic, spending more time at home due to enforced lockdowns is a naturally occurring implicit dietary intervention encouraging people to eat at home. By documenting the impacts on people's interests and measuring how lasting the effects are we learn something about the kinds of foods that people become interested in when staying at home in general. This has implications for designing interventions outside of COVID-19, and future work should compare effects on diet of staying at home due to COVID-19 lockdown measures to the impacts of staying at home due to other, more frequent external circumstances, such as extreme weather or air pollution.

6.4.2 Comparison to surveys

Our results confirm and refine what is known from survey-based research. A meta-analysis [413] of 12 preliminary articles studying the impact of COVID-19 confinement on dietary habits revealed a sharp rise of carbohydrate consumption, especially of foods with a high glycemic index (e.g., homemade pizza, bread, cake, and pastries), as well as more frequent

snacking. A high consumption of fruits and vegetables, as well as protein sources, particularly pulses, was also recorded, although there was no clear peak of increase in the latter. A decrease in alcohol intake and of fresh fish and seafood was further observed.

Whereas surveys are potentially a more accurate reflection of consumption, our findings, which were derived from passively sensed data, provide a complementary view. Search interest time series capture fine-grained temporal dynamics within the contrasted periods. Additionally, search interest time series are not subject to the reporting biases of surveys. By relying on them, we account for behavioral changes beyond subjective impressions. Finally, search interest time series provide insights at a population scale.

Contrary to previous concerns about the danger of alcohol abuse during confinement [72] triggered by stress, boredom, and emotional consumption, on a population level, we do not observe important surges in interest in alcohol consistent with these concerns. In fact, consistent with survey-based research [18, 413], we observe a significant negative correlation between seasonality-adjusted interest in alcoholic drinks and mobility in some of the studied countries (cocktails: -0.44 [$p = 0.002$] in Italy, -0.33 [$p = 0.02$] in Japan, -0.33 [$p = 0.02$] in Spain; wine, beer, and liquor: -0.61 [$p = 6 \times 10^{-6}$] in France, -0.38 [$p = 0.009$] in Japan, -0.34 [$p = 0.02$] in Denmark), meaning that more interest in alcohol is associated with more time spent outside of the home, not less. Additionally, the relative share of interest in alcoholic drinks (Figure B.9) decreased since the increase in other foods was not mirrored by the increase of interest in alcoholic drinks.

It is important to remember that these findings are based on aggregate population-level interests, and that specific subpopulations of users might still be susceptible to alcohol misuse. Future work should study search logs and alternative digital traces [11] of individual users in a longitudinal user-level study to understand what pre-pandemic user characteristics are predictive of behaviors emerging during confinement.

6.4.3 Limitations

When interpreting our results, several additional considerations should be kept in mind. First, searching for a food is not tantamount to consuming the food. Users may search but not consume, and vice versa. Also, search interest might not be an equally good sensor for real behavior in different countries. Note, however, that several factors nonetheless render our findings consequential. In other contexts, digital traces of nutritional behavior have been shown to be valid proxies of actual behavior [3]. Additionally, even if traces are imperfect proxies, major shifts in search interest have the potential to impact actual food consumption. In that sense, search interest, one of the few global signals that are publicly accessible to researchers and policymakers, can lead to consumption.

Second, while we make no claims of causal identification based on our statistical analyses, our regression discontinuity-based design alleviates the effect of unobserved covariates by

exploiting the sudden shock in mobility and accounting for seasonal variation. The observed dose-response relationship (Figure 6.5b) supports this, as does the fact that search interest in ways of accessing foods behaves as one would expect if those interests were causally affected by mobility.

Third, the data collection capacities limited the number of studied countries such that interest data could feasibly be collected. Our results may not be representative beyond the 18 countries studied here. Still, given the globally shifting interests in ways of accessing food (Figure 6.1), we believe the results from the countries studied here are indicative of shifts in interests in neighboring countries.

Finally, beyond people's shifting habits, interests, and emotional responses, other internal and external factors brought by the pandemic, most notably food product availability, price, and expected shelf life [240] or populations' present level of cooking skill and willingness and ability to learn to cook [287] can play a role, and should be kept in mind when interpreting the observed shifts in dietary interests.

6.4.4 Future work and conclusions

In this chapter, we study and document the impacts of a single event (COVID-19 crisis), but we observe similar impacts across culturally and geographically different countries. The observed impacts are therefore general to a certain extent, applying to different kinds of populations, in varying intensity depending on the intensity of the treatment.

When confined, people are interested in carbohydrates and calorie-dense foods (Figure 6.6), likely due to changes in preferences [49, 410] on the one hand and due to changes in accessibility and price of foods [240] on the other hand. These effects are consistent across countries, which is a demonstration that they occur across cultures and economic conditions.

While this study quantified initial developments during the pandemic, future studies aiming to understand the impacts of the pandemic and the related mobility restrictions on diet will continue to be important for designing policies and programs to tackle adverse health impacts.

Validity of studying dietary behaviors with digital traces

Part IV

7 Validity of dietary digital traces

7.1 Introduction

In order to be able to improve diets, researchers and stakeholders need to know what foods people consume, but monitoring diets at a population scale is challenging. Traditionally, nutritional studies rely on survey-based methods [66] employing questionnaires [399] and personal food journals [77, 78, 323], which are prone to biases, most notably social and cognitive biases, such as false recall and social desirability bias [46]. Traditional methods are also costly to organize. Researchers and practitioners might not have access to extensive surveying, and it might be hard to collect reliable statistics—even though a large and ever-growing [280] portion of the population has access to advanced technology including smartphones with Internet access [39].

In light of the challenges of the traditional methods on the one hand, and the opportunities afforded by widespread Internet access on the other hand, there is great promise in using passively collected digital data to estimate food consumption. In Chapters 4–6 we presented studies leveraging different types of digital trace data to study dietary behaviors. Such digital datasets are unmatched in terms of scale and immediacy [244, 315], do not rely on self-reports, and do not suffer from biases typical of traditional methods. Given this potential, researchers have been developing and applying their expertise to studying diets via passively collected digital data, whose tremendous potential for providing insights into food consumption has been showcased numerous times [3, 69, 319].

The promises of digital trace data notwithstanding, important methodological questions remain: Are researchers measuring what they aim to measure? Do digital traces reflect actual food consumption? Do effects estimated from online signals hold in the offline world? Are predictive models trained on online signals accurate in the offline world? In other words, the validity of studying diets with digital data remains opaque.

Online data is not primarily collected with scientific studies in mind and is therefore sometimes referred to as “found data” [315]. Found data overcomes several of the biases typical

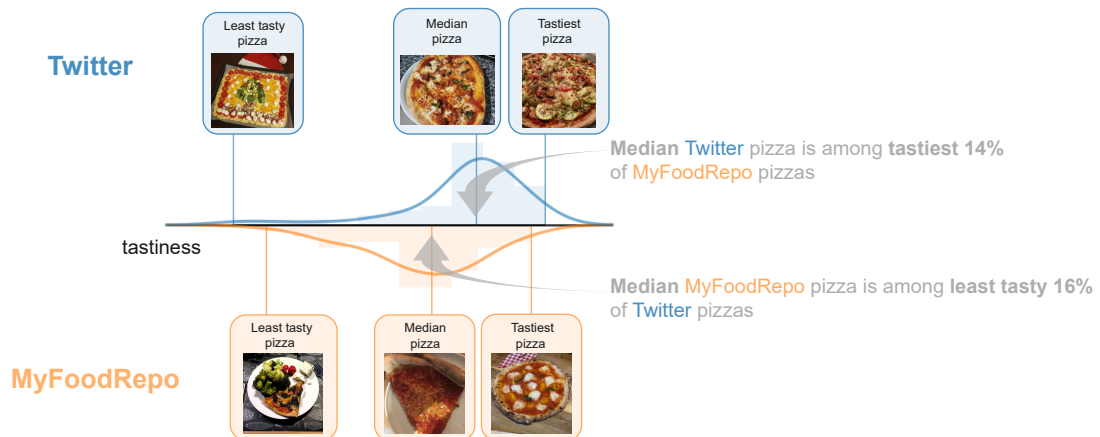


Figure 7.1: Perceived tastiness of tweeted food (*top*) vs. actually consumed and tracked food (*bottom*) of type “pizza”. Histograms summarize tastiness scores estimated in our crowd-sourcing framework. As illustrated, tweeted pizzas are perceived as considerably tastier than actually consumed and tracked pizzas.

of traditional methods, but may introduce new biases that threaten validity in their own ways [214, 264, 380]. Despite their potential, error-prone and unreliable data and methods may do more harm than good if handled without the required caution [76]. As our communities increasingly rely on large-scale digital data-sources, methods that offer insights into the validity of new measures thus become increasingly necessary [214].

7.1.1 Research questions

In this chapter, we focus on a specific digital trace data source—social media. Our overall research question asks: *Is social media a biased or a truthful mirror of actual food consumption, as measured via food tracking?* We focus on those dietary aspects in particular that researchers frequently study with social media: food type [88], nutritional properties [3], and appearance and subjective perception [243]. In order to establish a link between online and offline dietary behaviors at a population scale, we study images of food that people consumed and tracked, and contrast them with images of food posted on Twitter, addressing the following specific research questions:

- RQ1** *Bias of food-type distribution:* To what extent do food images posted on Twitter reflect the types (beef, bread, burger, etc.) of actually consumed food, as measured via food tracking?
- RQ2** *Biases within food types:* For a given food type, to what extent do food images posted on Twitter reflect the nutritional properties, tastiness, and appearance of actually consumed and tracked food of that type?

In order to address RQ1, we investigate whether food images posted on Twitter are a faithful reflection of the types of actually consumed and tracked food or not. Specifically, we envision the following potential outcomes:

- a) “Food images posted on Twitter are a faithful reflection of the types of actually consumed and tracked food, consistent with the demonstrated potential of Twitter to provide insight into dietary choices [3].”
- b) “Food images posted on Twitter are not a faithful reflection of the types of actually consumed and tracked food, given a variety of challenges in the practices of social media use for research [264].”

In order to address RQ2, we investigate whether or not food images posted on Twitter are a faithful reflection of actually consumed and tracked food in terms of how healthy, caloric, and tasty the food is. We investigate whether the two sources diverge, and if so, in what direction. Specifically, the potential outcomes that one could envision are the following:

- a) “Food images posted on Twitter are a faithful reflection of actually consumed and tracked food in terms of how healthy, caloric, and tasty the food is, consistent with the demonstrated potential of Twitter to provide insight into dietary choices [3].”
- b) “Tweeted food is healthier, less caloric, and less tasty than consumed and tracked food. Social media is increasingly used to promote trendy ingredients and recipes, and clean and healthy eating [68]. Social media inspires and connects people interested in healthy eating [245].”
- c) “Tweeted food is less healthy, more caloric, and tastier than consumed and tracked food, consistent with a documented fetishization of food online. Users share appetizing pictures of culinary experiences where exaggerated foods such as sugary desserts dominate over more standard local cuisines [243].”

Contributions. To the best of our knowledge, ours is the first attempt to investigate the link between online and offline dietary behaviors by studying food images, as measured via two platforms: Twitter and MyFoodRepo¹ food tracking app. We design and apply a novel crowdsourcing framework for estimating biases (Section 7.2), and we perform a case study of food consumption in Switzerland (Section 7.3). Controlling for location, period and food types, we contrast an extensive set of tweeted food images with images of consumed and tracked food.

¹<https://www.myfoodrepo.org/>

7.1.2 Summary of the main findings

We find that food type distributions among social media foods vs. among consumed and tracked foods diverge (RQ1, Section 7.3.1). Controlling for the discrepant food-type distributions by studying food types individually (RQ2, Section 7.3.2), we find that Twitter still provides a biased view of food consumption as measured via food tracking. Tweeted food is, on average across food types, perceived as more caloric, less healthy, less likely to have been consumed at home, and tastier (example in Figure 7.1), compared to actually consumed and tracked food. For example, on average across food types, a median-tasty Twitter dish is among the top 26% tastiest MyFoodRepo dishes, and a median-caloric Twitter dish is among the top 34% most caloric MyFoodRepo dishes. While social media traces can be a reasonable proxy of tracked consumption for certain foods types (Figure 7.6), we find that, overall, food shared on social media and consumed and tracked food significantly diverge from each other (Figure 7.4 and 7.6, Table 7.1). The fact that there is a divergence between food consumption measured via the two platforms—food tracking and social media—implies that at least one of the two is not a faithful representation of the true food consumption in the general Swiss population. Researchers should thus be attentive and try to establish evidence of validity before using either social media or tracking apps as a proxy for the true food consumption in the general population.

7.1.3 Implications

Measuring biases in digital traces is the first step toward correcting them and drawing valid conclusions despite their presence [385]. Through a case study of the Twitter and MyFoodRepo platforms in Switzerland, contrasting tweeted food images with consumed and tracked foods, we provide grounding and first insights by controlling for location, period, and food types. Our findings cannot be assumed to generalize globally, and future work should apply our framework to other populations, other social media platforms and Web traces, and other food tracking apps. Our study may serve the purpose of a “proof by counterexample”: we have identified one common setting where there is a bias between two types of digital trace data. Hence, we should assume that there can be bias in other populations, too. We conclude the chapter with a discussion (Section 7.4) of how the methods and findings reported here can inform researchers in their efforts to leverage digital traces for various applications, in the context of food and beyond.

7.2 Materials and methods

7.2.1 Food tracked via MyFoodRepo

To get as close as possible to capturing true food consumption, we use a novel dataset of food images collected via the MyFoodRepo mobile app [248]. By design, the food present in these images was actually consumed, for the purpose of the application is to track users’ personal

food consumption. Through the app, volunteer users from Switzerland provide images of their complete daily food intake, mainly in the context of being enrolled in a digital cohort called Food & You [110].² MyFoodRepo captures all foods that individuals consume, in any context.

The images are publicly available as part of the Food Recognition Challenge.³ The dataset has been annotated such that the individual foods are mapped onto an ontology of food types. Images were logged between 2017 and 2020. In our analyses, we study the training-set portion of the dataset, comprising 24,120 images, along with their corresponding 39,328 food-type annotations.

7.2.2 Food shared on Twitter

To answer the question of whether images shared on social media diverge from food consumption as measured via food tracking, we aim to contrast images of consumed and tracked food with images of food posted on social media. To this end, we curate a dataset of food images shared on Twitter in Switzerland during the same period spanned by the images collected via the food tracking app, this way controlling for location and time.

Since our goal is to investigate the validity of studying diets with social media, in our Twitter data collection strategy, we aim, first, to follow data collection methods present in the existing literature closely, to be able to make conclusions that can be relevant for researchers working in this area, as opposed to inventing novel strategies that would be less relevant. Our data collection pipeline is therefore similar to pipelines described in related work, extracting nutritional information from social media posts with keywords (Chapter 2). Note that, since we follow existing work, specific data collection decisions are not limitations per se. Instead, the impact of data collection based on user-specified keywords is intended to be measured, since this is how researchers usually collect Twitter posts to study food consumption.

Second, we aim to gather a complete dataset, i.e., to collect all food images posted on Twitter by the relevant population in the relevant time frame. To this end, we use the full-archive search endpoint, available to researchers via the Academic Research product track,⁴ which allows searching Twitter's complete archive going back until March 2006.

Third, we aim to find images posted on Twitter that actually contain food, as we are interested in studying the posted food itself, rather than only how it is described. To this end, we apply automated and manual annotation of the collected images. With these three goals in mind, we employ the following data collection pipeline.

Step 1: Twitter data collection. We start from the set of MyFoodRepo image annotations (e.g., “bread”, “banana”). We remove drinks and merge small types that are similar, obtaining 155 food types. We map each type to suitable handcrafted high-precision keywords, translate the

²<https://www.digitalepidemiologylab.org/projects/food-and-you>

³<https://www.aicrowd.com/challenges/food-recognition-challenge>

⁴<https://developer.twitter.com/en/docs/twitter-api/tweets/search/>

keywords from English to German, French, and Italian (the large Swiss national languages) via Google Translate, and use the disjunction (“OR”) of keywords pooled across languages to query Twitter’s full archive search API for the respective food type. We thus obtain all posts that in the text contain at least one of the keywords related to the food, in one of the four languages, in either singular or plural form (if relevant). For example, for the type “bread”, we retrieve all English tweets containing “bread” or “breads”, French tweets containing “pain” or “pains”, Italian tweets containing “pane” or “pani”, and German tweets containing “Brot” or “Brote”. Additional restrictions ensure that tweets were posted between 2017 and 2020 (the period when images of MyFoodRepo food were logged) from a location in Switzerland and contain at least one image. This step yields 33,425 unique images.

Step 2: Automated annotation. We are interested in studying the foods themselves, so we make sure that images indeed contain food. To that end, we perform detection of food in images with the ResNet50 model trained on ImageNet [164] and finetuned for food-vs.-not-food classification on the publicly available Food-5K food image dataset [336].

We obtain 98% recall and 96% precision on the task of detecting food images on the held out 20% test set (using a threshold of $p = 0.5$). Inspection of the images revealed that the images that do not contain food most frequently occur in food types where keywords have homonymous meanings. Two food types with the largest fraction of images that do not contain food are “date” (fruit “date” can have meaning “the day of a year”, or “a social appointment”) and “apple” (“apple” can relate to the company, and not the fruit). After this step, we keep 7,723 tweets with images that contain food.

Step 3: Manual annotation. We manually inspect the images to verify that an image contains the food that the user mentions in the tweet text, even if a small quantity. The visible food item needs to be edible, e.g., a silver pendant of lemon shape or a carved and decorated Halloween pumpkin does not qualify as such. Additionally, the image needs to contain a prepared dish, and not all the ingredients laid out separately, nor an uncooked caught fish. Finally, no explicit content can be present in the background for the image to be safe for crowdsourced participants.

In the end, after this step, we retain 3,692 images of food along with their corresponding 4,481 food-type annotations. Due to the completeness of Twitter’s full archive search and the manual inspection of collected images, at the end of this process, we obtain all tweets posted from Switzerland between 2017 and 2020 with images that contain a food that is mentioned in the tweet text.

Data summary. The two datasets we analyze contain 24,120 images of *consumed and tracked food* and 3,692 images of *tweeted food*. Images are mapped on the food-type level and contain foods that we can compare in order to address our research questions. See Figure 7.2 for examples of images of type “pizza”.

Having described the data, we continue by outlining our crowdsourcing framework for measuring biases. We then describe how we implement this framework on Amazon Mechanical Turk.

7.2.3 Crowdsourcing framework for estimating biases

Beyond food types, previous work has most notably used social media to estimate nutritional properties of food [3], as well as its appearance and perception [243]. Based on these themes, we operationalize four pertinent dimensions along which we contrast tweeted and consumed and tracked food, capturing how (1) healthy, (2) tasty, (3) caloric, and (4) likely to have been consumed at home the food is.

For each dimension, we aim to estimate a score for each image. Contrasting the scores of tweeted vs. consumed and tracked food then allows us to assess biases. In principle, we could obtain scores directly via human annotation by asking, e.g., “How tasty does this dish look, on a scale from 1 to 10?” It is, however, challenging for humans to place items on an interval scale that is consistent across individuals [8]. Based on the fact that judging between two alternatives is generally easier and more intuitive for humans [61], we instead adopt a pairwise paradigm, where we confront human raters with pairwise choices (e.g., “Which of these two dishes looks tastier?”) and later infer latent scores from the pairwise preferences.

Consider a given dimension (we use tastiness for concreteness in the following exposition) and a given food type. Then, for two images a and b showing food of the same type, we use the notation “ $a > b$ ” to express that a is preferred over b by a human rater. Note that human preference is a random variable: different raters may have different preferences with respect to a given pair. We assume, however, that certain images show inherently tastier dishes and are thus more likely to be preferred. More formally, following the Bradley–Terry (BT) model [48], we assume that each image i has a latent tastiness score $s(i)$ and that the probability that a rater will prefer image a over image b [image b over image a] in a pairwise comparison is proportional to the score of a [score of b]:

$$\Pr(a > b) = \frac{s(a)}{s(a) + s(b)}. \quad (7.1)$$

Given this setup, maximum likelihood estimation [235] can be used in order to infer the latent scores that best explain the empirically observed pairwise preferences. Thus, although only pairwise choices are made by humans, we can rank all images in a total order based on their latent scores s . The BT model is appropriate for our purposes, as it has a well-understood interpretation and is well-suited to model human preferences [61]. In practice, we fit a so-called Plackett–Luce model [235], a generalization of BT that does not require comparisons for all image pairs.

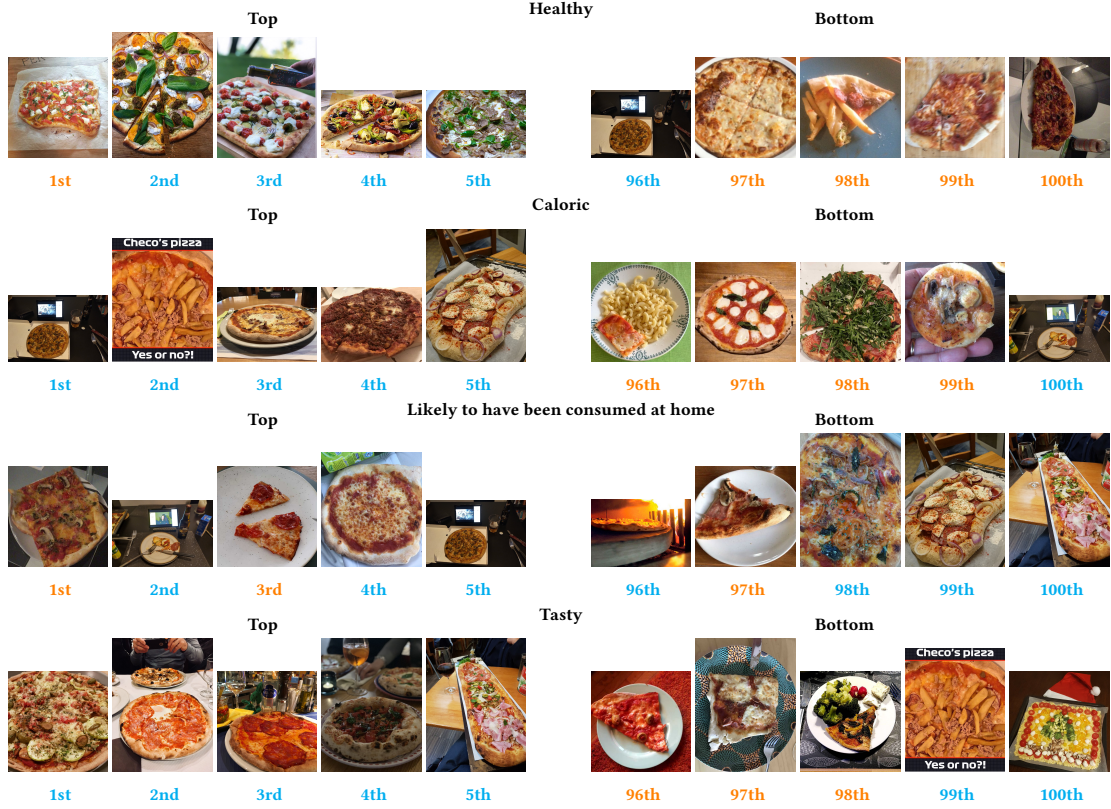


Figure 7.2: Example images in food type “pizza”. In rows, overall top five (Rank 1st-5th) and overall bottom five (Rank 96th-100th) images with respect to estimated rank. Twitter foods are marked in blue and MyFoodRepo foods in orange.

7.2.4 Implementation on Amazon Mechanical Turk

In the remainder of this section, we describe how we implemented the above-described framework on Amazon Mechanical Turk. To estimate the latent scores of images of MyFoodRepo and of Twitter food in terms of the four dimensions, we selected 24 well-represented food types (Figure 7.6). The types are selected such that each type has at least 50 tweeted, and at least 50 consumed and tracked food images. The different food types are considered as independent “tournaments”, so we obtained a separate ranking per type.

We first sampled the same number of tweeted food images and consumed and tracked food images per type, to account for potentially different food-type frequencies. Recall that location and period are already controlled in the data collection. We randomly sampled 100 images from each type, 50 tweeted foods and 50 consumed and tracked foods, resulting in 2,400 competing images in total.

We then did random sampling of comparison pairs. From each set of images for a given food type, we sampled K comparisons, at random taking one image of consumed and tracked food and one image of tweeted food, constrained such that each image participates in the

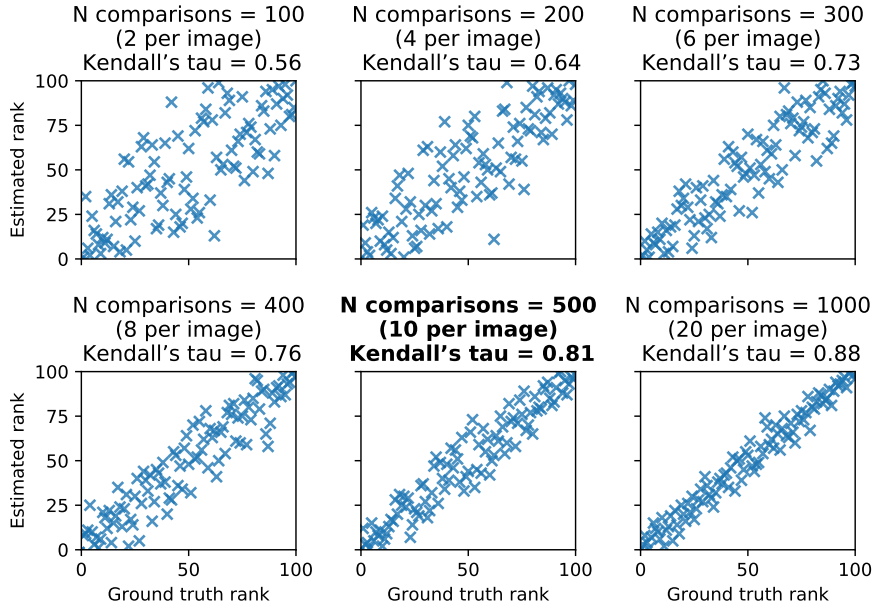


Figure 7.3: Estimation of image rank with simulated sampling of duels. We chose to perform ten comparisons per image.

same number of comparisons. We chose the number of “duels” per food type based on rank inference simulations, with the goal of ensuring that we can infer ranks accurately. For a unique finite maximum likelihood estimate to exist, the comparison graph must be fully connected, i.e. multiple comparisons are needed, but how many?

We assume 100 items divided into two groups with item quality sampled from the standard normal distribution. We rank the items and then randomly sample comparisons between items from the two groups. We decide the outcome of a duel based on the item quality, increasing the number of sampled comparisons. We then estimate quality and rank with a BT model (Equation 6.1). Figure 7.3 depicts how well the ranking can be recovered for a different number of comparisons. As more comparisons are performed, estimated rank (y-axis) correlates more strongly with ground truth rank (x-axis). We chose to perform ten comparisons per image ($K = 500$, Kendall's tau 0.81). That is, at $K = 500$, each of the 50 images is compared to 10 competitors and rank can be accurately inferred with Kendall's tau 0.81.

In every rating task, a participant was shown a random pair of images containing food of the same type (images in a pair were scaled to the same size and shown in randomized order), and asked to give a preference label for each of the four dimensions (healthiness, tastiness, caloric content, likelihood to be consumed at home). The pairwise comparison task had no neutral option. Participants were required to choose one image. As the order within pairs was randomized, this is a valid way of breaking ties, and recommended practice [274]. Additionally, we asked participants to explain how they perceived both images by providing between one and three free-form tags (e.g., “dull”, “greasy”). (Prior to data collection, we did not make

hypotheses about specific biases as revealed by the tags, but rather explore them post-hoc in order to gain insights about how people describe the appearance of tweeted vs. consumed and tracked food.) In total, we collected 12,000 pairwise comparisons (500 duels for each of 24 types) for each of the four dimensions.

Participants. Since the tasks require reading and writing text in English, participants were restricted to those residing in the United States, Canada, or the United Kingdom. To ensure high-quality answers, we admitted only workers with approval rates greater than 99% and with more than 1000 previously approved tasks. We collected the 12,000 pairwise preferences through 24 batches with 500 assignments each, over the course of five days. The task was performed by 595 distinct workers, who performed 20.2 pairwise comparisons, on average. One crowdsourced participants rated one pairwise comparison.

Compensation. We targeted a pay rate of \$9 per hour. Participants were paid \$0.15 per pairwise comparison. The mode of the time taken per comparison is 57 seconds, which corresponds to an estimated hourly rate of \$9.5 (U.S. federal minimum hourly wage in 2021 is \$7.25 per hour, for reference). Note that this is likely an underestimate of the hourly rate since crowd workers often use scripts that make it possible to automatically accept a task they are interested in, and hold it assigned while not actively working on it.

Instructions. To ensure reproducibility of our experiment, below we quote the instructions as they were displayed to the participants:

*Please take a look at the two images displayed below. Please focus on the food itself, and not the other contents of the image. Answer the questions about the pizza shown in the images by entering either **1** for Image 1, or **2** for Image 2. Additionally, please explain your preferences by adding **at least one word or short phrase to describe foods** appearing in each image. Write a word or a short phrase, and not full sentences.*

- 1: Which image contains pizza that appears more **tasty**?
- 2: Which image contains pizza that appears more **healthy**?
- 3: Which image contains pizza that appears more **caloric**?
- 4: Which image contains pizza that appears more likely to have been **consumed at home**?
- 5: Pizza shown in Image 1 is?
Add a word or a short phrase to describe food in Image 1
- 6: Pizza shown in Image 2 is?
Add a word or a short phrase to describe food in Image 2

7.3 Results

7.3.1 Bias of food-type distribution

To address RQ1, we start by comparing MyFoodRepo food images with images of foods posted on Twitter. We compare the prevalence of the 155 food types across the two sets (Figure 7.4).

First, we observe a significant positive correlation (Spearman's rank correlation coefficient $\rho = 0.49$, $p = 8.5 \times 10^{-11}$). The more frequent a food type is among consumed and tracked foods, the more frequent it tends to be among tweeted foods. That said, although food type frequencies are correlated, important deviations can be observed: For instance, bread and butter are more likely to be observed in MyFoodRepo foods compared to Twitter foods, i.e., these foods are underrepresented on Twitter. Bread is 2.5 times more frequent among MyFoodRepo foods, while butter is 5.5 times more frequent among MyFoodRepo foods. On the other hand, cake, soup, chocolate, raclette, burgers, etc., are more likely to be observed among Twitter foods, compared to MyFoodRepo foods, i.e., these foods are overrepresented on Twitter. Soup is 11.5 times, cake 12.0 times, burger 10.0 times, and raclette 9.5 times more frequent among Twitter foods.




Figure 7.4: Comparison of food type frequency among Twitter food (x-axis) vs. MyFoodRepo food (y-axis). Categories where the larger frequency is at least two times greater than the smaller frequency are annotated. Font size is proportional to pointwise KL divergence between the larger and the smaller frequency. Gray diagonal line marks identity, where two frequencies are equal.




Figure 7.5: Histogram of crowdsourced participants' preferences. On x-axis, fraction of duels where Twitter food image is selected when compared to a MyFoodRepo food image. On y-axis, fraction of workers with such preferences. Dashed vertical lines mark average fractions of wins, across participants. Full vertical line marks 0.5.

7.3.2 Biases within food types

Duel outcomes. Controlling for the different food type distributions, we address RQ2, which is concerned with biases within fixed food types. As an initial look into the duel outcomes, across all duels, we first check the fraction where the Twitter image won. Together with the fraction of duels where Twitter image won, we report p -values from two-sided binomial tests, where the null hypothesis is that the outcome of comparisons is random, i.e., that the Twitter image wins in 50% of duels. Across all duels, in 58.46% of duels Twitter image is chosen as more caloric ($p < 10^{-70}$), in 45.96% as more healthy ($p < 10^{-10}$), in 38.08% as more likely to have been consumed at home ($p < 10^{-140}$), and in 61.73% as more tasty ($p < 10^{-100}$).⁵

Bias measurement: score estimations. Next, in order to compare images, as opposed to the outcomes of duels, we fit the BT model (Equation 6.1) on the collected preferences and estimate a score that represents the latent quality of each competing image concerning how healthy, tasty, caloric, and likely to have been consumed at home the food appears.

Consider a given dimension (we again use tastiness for concreteness). Let each MyFoodRepo image $i \in \{1, \dots, N_M\}$ have an estimated tastiness score $s(i)$ sampled from the distribution of consumed and tracked foods by MyFoodRepo users, and each Twitter image $j \in \{1, \dots, N_T\}$ have an estimated tastiness score $s(j)$ sampled from the distribution of foods posted on Twitter. The tastiness bias $b(T, M)$ between food consumption measured with Twitter and

⁵We also examined the macro-average duel outcomes, where we check preference of each crowd worker, and average over the workers. There appears to be no rater bias where some workers overwhelmingly prefer one or the other (Figure 7.5) as the estimates are consistent, and no notable outlier crowd workers emerge.

food consumption measured with MyFoodRepo can be expressed as a difference in the average tastiness scores measured via the respective data sources, T and M :

$$b(T, M) = T - M = \frac{1}{N_T} \sum_{i=1}^{N_T} s(i) - \frac{1}{N_M} \sum_{j=1}^{N_M} s(j). \quad (7.2)$$

The measured bias⁶ $b(T, M)$ is 0.52 [0.46, 0.56] for how tasty the food is, 0.39 [0.35, 0.45] how caloric, -0.18 $[-0.23, -0.11]$ how healthy, and -0.58 $[-0.63, -0.52]$ for how likely to have been consumed at home the food is. These bias measurements indicate that, on average, food posted on Twitter is perceived as significantly tastier, more caloric, less healthy, and less likely to have been consumed at home compared to consumed and tracked foods.

Bias measurement: rank estimations. Once latent scores are estimated, images can also be ranked, either jointly, or separately, among Twitter and MyFoodRepo food. Given its intuitive interpretation, our main method of analysis is quantifying the shifts in distributions via ranks, as depicted in Figure 7.1. For concreteness, it is helpful to consider an example before studying images on a more aggregate level across food types. Figure 7.2 contains the top and bottom portions of the joint rankings (one ranking for each of the four dimensions) for food type “pizza”.

For each food type, we rank the two sets (MyFoodRepo food vs. Twitter food) separately and determine which percentile of MyFoodRepo food each percentile of Twitter food corresponds to. We focus on the median Twitter food image (also referred to as the “typical” Twitter image), and compute the percentile rank of its score, relative to scores of MyFoodRepo food images. The rank of the median Twitter image among MyFoodRepo food images is presented in Figure 7.6, across 24 food types.

Bias in nutritional properties. We find that Twitter foods are perceived as more caloric, less healthy and less likely to have been consumed at home. On average across food types, the median-caloric Twitter food is among the top 34% most caloric MyFoodRepo foods. The median-healthy Twitter food is among the bottom 42% most healthy MyFoodRepo foods. The median Twitter food is among the bottom 27% most likely home-consumed MyFoodRepo foods.

Examining food types separately, regarding how healthy the foods are estimated to be, we find no significant differences in 15 out of the 24 types. In eight food types, tweeted food is perceived as *less healthy*, and one food type (vegetables) is found to be more healthy on Twitter. With respect to perceived caloric content, we find no significant differences in 12 out of the 24 types. For 11 types, tweeted food is perceived as *more caloric*. Vegetables are again an exception, found to be less caloric on Twitter compared to MyFoodRepo foods. We find that in most of the types (16 out of the 24), tweeted foods are *less likely to have been consumed at*

⁶The corresponding 95% confidence intervals are obtained via bootstrap resampling of images.




Figure 7.6: Relative rank (with respect to estimated latent scores) of median tweeted food among consumed and tracked foods, where 100% corresponds to top score. Colored bars mark estimates that significantly differ from median (i.e., 50% relative rank). Yellow marks ranks higher, purple marks ranks lower than median, while gray marks non-significant differences. Error bars mark bootstrapped 95% confidence intervals obtained via bootstrap resampling of duels. Example: relative rank of median tweeted raclette with respect to tastiness is 90% (95% CI [79%,96%]); i.e., median tasty tweeted raclette is ranked among top 10% [4%,21%] most tasty consumed and tracked raclette images.

home. For eight food types, there are no significant differences in likelihood of having been consumed at home.

Although there are food types with no biases in nutritional properties, we observe large biases for certain foods, where a median-healthy Twitter food is among the bottom 20% of MyFoodRepo food images with respect to healthiness. For instance, median-healthy Twitter cheese is among the bottom 14% (95% confidence interval [6%,24%]) on MyFoodRepo; median-healthy Twitter chocolate, among the bottom 18% [10%,28%] on MyFoodRepo; and median-healthy eggs, among the bottom 19% [11%,29%] on MyFoodRepo.

We also measure large biases for food types where a median-caloric Twitter food is among the top 20% caloric MyFoodRepo food. For example, median-caloric chocolate is among the top 6% [2%,17%] on MyFoodRepo; median-caloric cheese, among the top 8% [4%,16%] on MyFoodRepo; median-caloric chicken, among the top 16% [8%,27%] on MyFoodRepo; and median-caloric cake, among the top 18% [12%,28%] on MyFoodRepo.

Bias in tastiness. Next, we find substantial bias in how tasty the foods are perceived to be. There are more discrepancies in perceived tastiness compared to nutritional properties (Figure 7.6). On average across food types, the median-tasty tweeted food is among top 26% most tasty consumed and tracked foods. A median-tasty tweeted food is ranked significantly higher than the median-tasty consumed and tracked food image in 20 out of the 24 types. In four types, there are no significant differences regarding tastiness (mushrooms, potato, egg, croissant).

We note that, for a number of foods, a median-tasty tweeted food is ranked as high as among the top 20% of consumed and tracked food. The median-tasty Twitter raclette is among the top 10% [4%,21%] on MyFoodRepo; the median-tasty soup, among top 10% [4%,16%]

on MyFoodRepo; the median-tasty cheese, among top 12% [4%, 23%] on MyFoodRepo; the median-tasty pizza, among top 14% [4%, 26%] on MyFoodRepo; the median-tasty sauce, among top 16% [18%, 24%] on MyFoodRepo; and the median-tasty Twitter chocolate, among top 18% [9%, 26%] on MyFoodRepo.

Correlations. Next, we inspect the correlation between tastiness and nutritional properties. Recall that we obtained 2,400 images (100 images for each of the 24 types), with four estimated quality scores. Computing Pearson’s correlation between the estimated scores (Figure 7.7), we observe a moderate positive correlation between how caloric and tasty ($\rho = 0.33$, $p < 10^{-60}$) foods are, and a negative correlation between how caloric and how healthy ($\rho = -0.39$, $p < 10^{-80}$) they are. While the correlation between how caloric and how healthy foods are is expected, the correlation between how tasty and how caloric they are might indicate that tweeted food might be perceived as overwhelmingly tastier because it is more caloric and exaggerated.

Complete rank comparisons. While so far we focused on the median Twitter image and studied its relative rank among MyFoodRepo images, in Figure 7.8 we present the full percentile rank comparisons, for completeness. Black diagonal line marks identity, where percentile distributions of Twitter food and MyFoodRepo food are equal. In the example of cheese (in the bottom-right), percentiles among tweeted food correspond to higher percentiles among consumed and tracked food regarding how caloric and tasty the cheese is (tweeted food ranks are above the diagonal line). Percentiles among tweeted food correspond to lower percentiles among consumed and tracked food regarding how healthy and home-consumed the cheese is (tweeted food ranks are below the diagonal line).




Figure 7.7: Correlation matrix between estimated qualities among 2,400 competing images.

Bias in appearance. Finally, we analyze the tags provided by crowd participants, in order to understand further how MyFoodRepo food and Twitter food differ in their appearance.⁷ Recall that crowd workers were asked to enter up to three tags per image (on average, participants entered 2.4 tags per image). Analyzing the 58,645 collected tags, we ask: *How do people perceive social media foods compared to consumed and tracked foods?*

Our main method of analysis is to compare tags across the two groups (Twitter vs. MyFoodRepo food). To capture the distinctiveness of a tag [408], we leverage pointwise Kullback–Leibler (KL) divergence between the two respective probabilities of observing the tag.

⁷Note that this is an exploratory post-hoc analysis to gain deeper insights into potential mechanisms that drive the observed biases. Our main analyses are related to perceived nutritional properties and tastiness, as estimated via pairwise comparisons.

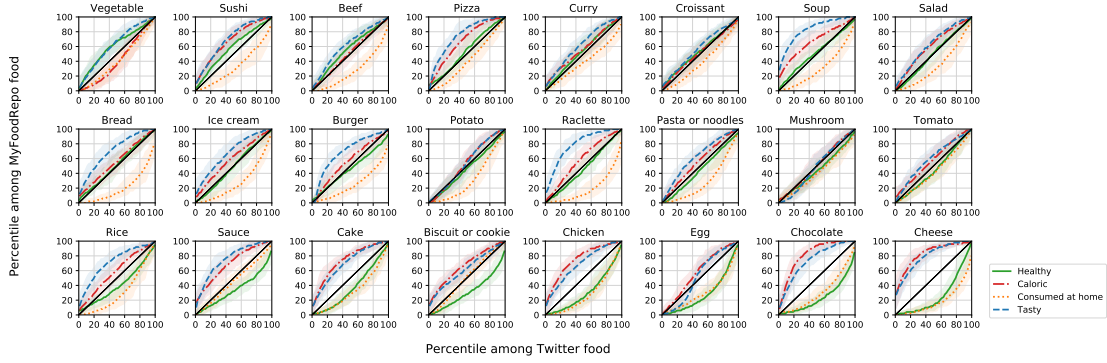


Figure 7.8: For percentiles among Twitter food (on the x-axis), the corresponding percentile among MyFoodRepo food (on the y-axis). Rank comparison is displayed for 24 types of food, across four estimated qualities: how healthy, caloric, likely to have been consumed at home, and tasty the food is. Food types are sorted according to the rank of median healthy tweeted food among consumed and tracked foods. Black diagonal line marks identity, where percentile distributions of Twitter food and MyFoodRepo food are equal.

We first performed normalization of tags provided by crowdsourced participants. We split commas, convert to lowercase, remove stop-words at the beginning of the text (e.g., “looks”, “seems”), and we map versions of words with a dash to a single form (e.g., mapping “mouth watering” and “mouthwatering” to “mouth-watering”).

A tag is typical for one set of images if used frequently within the set, but at the same time unlikely to be used in the other set. Additionally, it is not only the discrepancy between the two probabilities that matters; a tag should also appear frequently in a set to be considered typical of the set. This intuition is captured by the Kullback–Leibler (KL) divergence.

Based on the respective probabilities of observing the tag t among all tags, we then compute the pointwise KL divergence between the distributions of tags for Twitter food images and for MyFoodRepo food images. The distinctiveness of t with respect to MyFoodRepo food compared to Twitter food images is calculated as

$$D_{\text{KL}}(p_{\text{M}}(t) \| p_{\text{T}}(t)) = p_{\text{M}}(t) \log \frac{p_{\text{M}}(t)}{p_{\text{T}}(t)}, \quad (7.3)$$

where for each tag t , $p_{\text{M}}(t)$ is probability of observing a tag among MyFoodRepo food images, and $p_{\text{T}}(t)$ is probability of observing a tag among Twitter food images. On the other hand, since the KL divergence is not symmetric, the distinctiveness of Twitter food compared to MyFoodRepo food is captured as

$$D_{\text{KL}}(p_{\text{T}}(t) \| p_{\text{M}}(t)) = p_{\text{T}}(t) \log \frac{p_{\text{T}}(t)}{p_{\text{M}}(t)}. \quad (7.4)$$

In Table 7.1, we present the tags with the largest pointwise KL divergence, separately for tags distinctive of consumed and tracked, or tweeted food. For each tag, a chi-squared test on the two frequencies is used to measure significance, under the null hypothesis that the two groups do not differ in frequency. We now examine—first overall, then separately by type—how foods differ in their appearance. We see that, overall, the tags most indicative of consumed and tracked food are “plain”, “bland”, “simple”, “boring”, “small”, “thin”, “dry”, “healthy”, and “homemade”. On the contrary, tweeted food is more likely to be described as “fancy”, “fresh”, “colorful”, “delicious”, “tasty”, “raw”, “gourmet”, “flavorful”, and “large”. Zooming into specific food types, the exact differences for specific foods become apparent. For example, tweeted pizzas are more likely to be described as “large”, “fresh”, and “delicious”, whereas consumed and tracked pizzas are seen as “small”, “pepperoni”, or “saucy”.

Inspecting the most discriminative tags overall and separately across types of food, we identified the following four prominent themes in tags that are discriminative of consumed and tracked vs. tweeted food:

1. *Complexity.* Consumed and tracked food is described as simple and homemade (“plain”, “bland”, “simple” and “homemade”); tweeted food, as more elaborate (“fancy” and “gourmet”).
2. *Portion size.* Consumed and tracked food comes in small portions (“small”, “thin”), whereas tweeted food is exaggerated in portion size (“large”). Portion size differences are particularly evident for specific types of food, e.g., consumed and tracked burgers are described as “small size”; pizzas are “small” when consumed and tracked, and “large” when tweeted.
3. *Ways of preparing.* Tweeted food is perceived as “raw” and “fresh”, while consumed and tracked food is described as “dry”. The differences are evident when it comes to specific foods. Consumed and tracked beef is more likely to be “grilled”, whereas tweeted beef is more likely “raw”. Consumed and tracked vegetables are “chopped” and “overcooked”, whereas on Twitter, they are “fresh” and “raw”. Rice and chicken are “fried” on Twitter, and “dry” when consumed and tracked.
4. *Presentation.* Tweeted food is visually appealing, whereas consumed and tracked food is more likely to look repulsive. Tweeted food is said to be “colorful”, “delicious”, “flavorful”, and “tasty”, whereas consumed and tracked food can be usually less appealing, with “watery” soup, “greasy” and “unappetizing” raclette, and “watery” and “thin” sauce.

7.4 Discussion

Our goal has been to determine the validity of estimating food consumption from digital traces: from social media posts vs. from images of consumed and tracked foods. To this

Table 7.1: Tags most distinctive of MyFoodRepo food (*left*) and tags most distinctive of Twitter food (*right*), as determined by pointwise KL divergence (cf. Section 7.3.2, *Bias in appearance*); * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, **** $p < 0.0001$).

Food	Top tags typical for MyFoodRepo food			Top tags typical for Twitter food		
Overall	plain**** boring**** dry****	bland**** small**** healthy***	simple**** thin**** homemade****	fancy**** delicious gourmet***	fresh**** tasty*** flavorful****	colorful**** raw**** large****
Beef	grilled**	bland**	simple*	raw***	fancy*	flavorful
Biscuit or cookie	shortbread**	chocolate*	plain*	decorated***	festive**	chocolate chip**
Bread	white****	plain**	dry**	delicious**	fresh**	fancy*
Burger	simple***	small size*	fast food*	interesting**	attractive*	filling*
Cake	small****	chocolate****	chocolatey****	decorated****	fruity***	fancy**
Cheese	boring***	plain***	cold***	melted****	warm****	hot****
Chicken	plain****	bland****	dry****	fried***	tasty***	spicy**
Chocolate	dark****	bitter***	broken***	delicious***	flavoured**	fancy**
Croissant	delicious*	buttery*	stuffed*	warm*	fresh	healthy
Curry	rice**	white*	homestyle*	yummy*	healthy*	flavorful*
Egg	plain***	simple***	bland**	decorated**	raw**	colorful**
Ice cream	chocolate***	nice*	simple*	delicious	decadent	fancy
Mushroom	colorful	simple	pizza	raw**	creamy	large
Pasta or noodles	plain**	bland**	simple**	fancy***	delicious**	yellow*
Pizza	small****	pepperoni**	saucy*	large***	fresh**	delicious*
Potato	boiled****	peeled****	plain*	raw***	unpeeled**	whole*
Raclette	healthy***	greasy***	unappetizing*	sliced**	hot**	tasty*
Rice	bland***	plain**	dry*	flavorful**	mixed*	fried*
Salad	simple**	plain**	homemade**	filling*	great*	fresh
Sauce	thin****	watery***	light****	thick****	creamy****	red****
Soup	bland****	simple****	watery***	hearty***	colorful***	delicious**
Sushi	boring***	simple**	plain**	appetizing*	variety*	fresh
Vegetable	chopped*	overcooked*	mixed	fresh**	raw*	delicate
Tomato	small*	healthy	salad	mouth-watering*	juicy*	flavorful

end, we designed a crowdsourcing framework for measuring biases, contrasting tweeted food images with consumed and tracked foods images, and deployed it for the case of Switzerland.

7.4.1 Summary of main findings

We find that social media does not provide a faithful representation of food types of consumed and tracked food. Measuring biases in food-type distributions, we observe that cake, soup, chocolate, raclette, and burgers are among the most overrepresented foods on Twitter in Switzerland. Cake, soup, and burger are visually appealing food types suitable for sharing on social media, while chocolate and raclette are foods typical of Switzerland. Winter social and sport activities among residents might make them more likely to be shared online. On the other hand, bread and butter—among the most underrepresented on Twitter—are simple everyday foods that tend not to have a lot of potential to look particularly visually appealing.

Controlling for the discrepant food-type distributions, we find that tweeted foods are perceived as less healthy, more caloric, and less likely to have been consumed at home, compared to consumed and tracked food of the same type (Figure 7.6). We also find substantial bias in perceived tastiness. A median-tasty tweeted food is, on average across food types, ranked among the top 26% of consumed and tracked foods. Exploring free-form tags provided by

crowd workers reveals that these biases are likely mediated by differences in portion size, complexity, presentation, and different ways of preparing food. For example, tweeted food is 3.5 times more likely to be described as “large”, and 4 times more likely to be described as “fancy”.

These results provide evidence that food shared online tends to be exaggerated compared to tracked food. The most biased foods in terms of nutritional properties are foods that can be very caloric and high in fat and carbohydrates: chocolate, cheese, chicken, cake, and egg. On the other hand, we find that some of the foods are not skewed in terms of nutritional properties. For instance, for mushrooms and croissants, no significant difference is observed in any of the four dimensions (healthiness, tastiness, caloric content, likelihood to be consumed at home). This implies that some foods can still be validly studied via social media as a proxy for consumed and tracked foods.

7.4.2 Consumed food vs. tracked food vs. tweeted food

Our study attempts to establish a link between online and offline dietary behaviors by studying food images as measured via two platforms: Twitter and the MyFoodRepo food tracking app. In what follows, we consider the relationship between three distributions: all foods consumed by the general population (the actual phenomenon of interest), food consumption estimated via MyFoodRepo, and food consumption estimated via Twitter.

For concreteness, consider tastiness (but the following argument equally applies to all other dimensions studied here). Let T denote the average tastiness score estimated via Twitter, M the average tastiness score estimated via MyFoodRepo, and G the true (unobserved) average tastiness score of food actually consumed by the general Swiss population. As before (cf. Equation 7.2 and Section 7.3.2), let $b(T, M) = T - M$, and analogously, $b(T, G) = T - G$ and $b(G, M) = G - M$. Although G , $b(T, G)$, and $b(G, M)$ are not observed, we have:

$$\begin{aligned} b(T, G) + b(G, M) &= (T - G) + (G - M) \\ &= T - M \\ &= b(T, M), \end{aligned} \tag{7.5}$$

as illustrated in Figure 7.9 for the case $T > M$ (without loss of generality: if $T < M$, we may simply make the argument about $b(M, T)$ instead of $b(T, M)$).

The established substantial and significant bias $b(T, M)$ along the four studied dimensions (Section 7.3) therefore implies a lower bound on the unobserved biases, since at least one of $b(T, G)$ and $b(G, M)$ is at least $b(T, M)/2$. In other words, along all four studied dimensions, either Twitter or MyFoodRepo differs by at least $b(T, M)/2$ from the general population. Dividing the measured biases $b(T, M)$ by two, we find that the lower bound on the bias still corresponds to significant gaps in how tasty, caloric, healthy, and likely to have been consumed at home the food is. Therefore, at least one of Twitter and MyFoodRepo foods significantly differs from

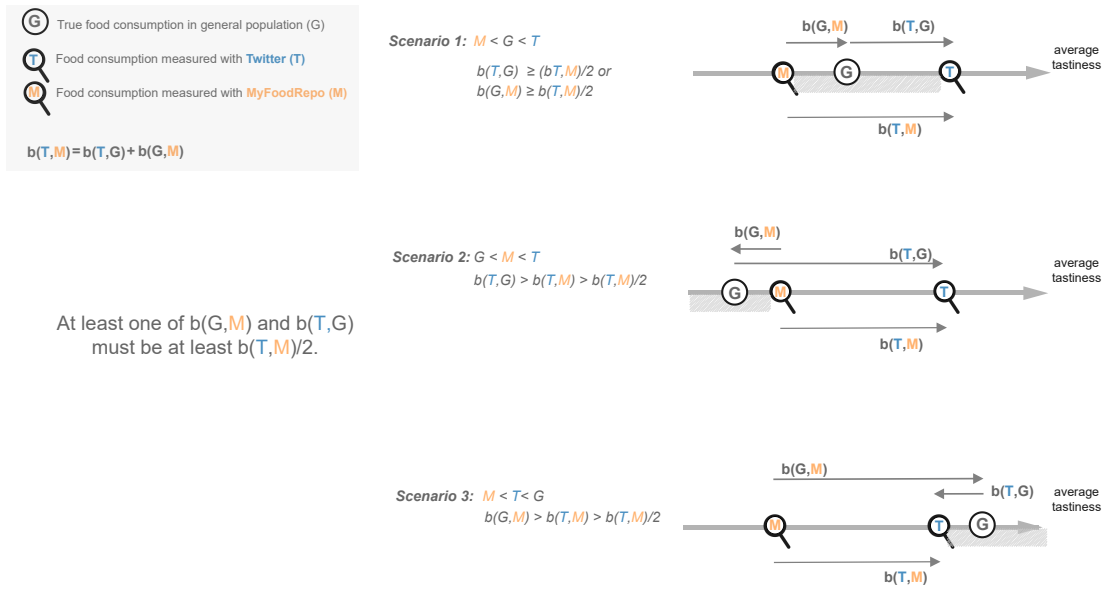


Figure 7.9: Illustration of the biases between the true food consumption, tracked food, and tweeted food, in the example of perceived tastiness. Consider the targeted true food consumption in the general population, the food consumption as measured with MyFoodRepo, and the food consumption as measured with Twitter. The bias between tweeted food and tracked food, $b(T, M)$, is characterized in our study, while $b(G, M)$ and $b(T, G)$ are unobserved. Although $b(T, G)$ and $b(G, M)$ are not observed, at least one of $b(T, G)$ and $b(G, M)$ is at least $b(T, M)/2$. The three illustrated possible scenarios depict considered situations where (1) $M < G < T$, (2) $G < M < T$, or (3) $M < T < G$.

the foods consumed by the general population. For example, the measured bias $b(T, M)$ is 0.52 [0.46, 0.56] for how tasty the food is (Section 7.3.2). The lower bound with the shifted corresponding 95% confidence intervals, $b(T, M)/2 = 0.26$ [0.20, 0.30], still corresponds to significant gaps in how tasty the food is.

Consequently, the fact that there is a divergence between food consumption as measured via food tracking and as measured via social media implies that at least one of the two is not a faithful representation of the true food consumption in the general population concerning how healthy, tasty, caloric, and likely to have been consumed at home the food is. Figure 7.9 illustrates possible scenarios in more detail. On the one hand, it might be that the true food consumption in the general population is somewhere between the food consumption as measured with MyFoodRepo, and the food consumption as measured with Twitter (i.e., $M < G < T$, Scenario 1). On the other hand, true food consumption in the general population might be skewed, and closer to either MyFoodRepo (i.e., $G < M < T$, Scenario 2) or Twitter (i.e., $M < T < G$, Scenario 3).

MyFoodRepo might be the main source of bias (i.e., G might be closer to T than to M) since individuals might not track all consumed foods, and food trackers are known to be situated

and contextualized [70, 225]. Similarly, consumed foods logged with the MyFoodRepo app are likely not representative of the full Swiss population. People who log food with the tracking app have access to a smartphone with an Internet connection and care about their diets. However, we argue that it is less likely that MyFoodRepo is the main source of bias since the majority of MyFoodRepo food images are collected from volunteers enrolled in a digital cohort called Food & You⁸ who are instructed and reminded to provide images of their complete daily food intake. By design, the food present in these images was actually consumed, and omissions are discouraged. Therefore, it appears less likely that MyFoodRepo could misrepresent the true food consumption in the general population to such an extent that it would fully explain the measured biases $b(T, M)$.

Nonetheless, given the absence of data about the general population and the fact that all the considered scenarios (cf. Figure 7.9) are not strictly impossible, we remain agnostic about the true source of bias. We argue that researchers should be attentive and aim to establish evidence of validity before using either social media or tracking apps as a proxy for true food consumption in the general population, since at least one of them differs by at least $b(T, M)/2$ from the general population. Future work should apply our framework for bias estimation to a representative sample of the overall population, with all food consumption recorded. At the present time, doing so remains challenging as the images logged with the tracking app by the volunteers are as good a peek into actual plates as we can currently get.

7.4.3 Implications

Implications for research studying food tracking as a proxy for offline behaviors. Based on our considerations of the two platforms (i.e., Twitter and MyFoodRepo), we argue that it is less likely that food tracking is the main source of bias when estimating food consumption in the general population. Nonetheless, researchers relying on food tracking should be attentive before implicitly assuming that the tracked consumption perfectly reflects the true consumption. Whenever possible, further contextualization of the tracked consumption data and investigation of alternative digital traces of the studied persons can be beneficial for examining the validity and establishing robustness. For instance, if there is a concern that users consume food systematically different from the logged one, future research should consider designing logging reminders and nudges within the tracking applications, targeted towards and specifically encouraging logging the true behaviors. Future research can also encourage users to assess the accuracy of logging through the tracking applications, for instance, by self-reporting the overall perceived truthfulness.

Implications for social media research studying online traces as a proxy for offline behaviors. Studies using passively collected digital traces as a proxy for real behaviors need to be valid in order to support public health research and have implications for the design of policies and interventions that can impact health outcomes. Based on our findings, we now highlight

⁸<https://www.digitalepidemiologylab.org/projects/food-and-you>

major potential pitfalls that can threaten validity of such applications and provide actionable implications for overcoming them.

Actionable implication 1: Addressing over- and underrepresentation of food types. If researchers were to estimate what foods general population consumes based on the number of tweets containing these foods, the estimates could be biased. We suggest triangulating social media behaviors with known government statistics whenever these are publicly available. For example, although bread is underrepresented while burgers are overrepresented compared with consumed and tracked foods (Figure 7.4), researchers could—even without access to logs of actual food consumption—identify the implausibly high prevalence of burgers on social media compared to bread by examining publicly available statistics [246]. For comparison, in 2019, the Swiss consumed 89 kg of products based on grains, compared to roughly half as many kilograms of meat (48 kg per person). When aggregate country-wide statistics are available for calibration, social media can still be used as a sensor for spatially and temporally fine-grained analyses (e.g., by neighborhood or during holidays). When no such statistics are available, researchers might be able to calibrate their methods on populations where statistics are available and adjust the final estimates. Domain knowledge about the populations being studied can also help in alleviating some of these disproportions. One could consider knowledge about foods that studied populations consume in a social context, foods frequently consumed by visitors and tourists, or only during special periods or occasions, and avoid studies being impacted by such idiosyncrasies.

Actionable implication 2: Foods with very biased nutritional properties. If researchers were to estimate nutritional properties of foods that the general population consumes based on the tweets containing images of foods, the estimates could be biased. Researchers should be careful about bias that stems from certain foods that appear particularly less healthy and more caloric compared to consumed and tracked food, such as chocolate, cheese, chicken, cake, and egg, although these might not necessarily be exactly the same foods types in populations beyond Switzerland. Researchers should also be aware of differences in portion sizes that likely mediate the difference in calories. We suggest detecting and examining images with an implausible amount of calories. For example, a single image might not be taken into account if the estimated amount of calories is not within reasonable bounds around the recommended daily 2,000–2,500 calories for an adult.

Similarly, when training machine learning models with datasets obtained from social media and aiming to generalize to the general population, samples should be adjusted such that the amount of calories more closely mirrors the amount of calories and portion sizes of real food. Otherwise, models trained on social media data to estimate the amount of calories will not make valid estimations outside of the context of exaggerated social media foods.

Actionable implication 3: Addressing systematic discrepancies in appearance. If researchers were to estimate appearance of foods that the general population consumes based on the tweets containing images of foods, the estimates could be biased. Foods that people consume

and track tend to appear significantly less tasty, simpler and less elaborate, prepared in different ways, and smaller in portion size, compared to tweeted food. These are challenging biases to overcome, as there is a need to use human annotation or computer vision models. Note, however, that the foods that are most biased in terms of nutritional properties and appearance are also precisely those that are overrepresented on Twitter (Figure 7.4). Therefore, adjusting the bias in the distribution of foods is likely to alleviate the bias in nutritional properties and appearance.

Implications for social media research studying online traces *per se*. Research studying online communities and online content not as a proxy for real behaviors but as a *phenomenon per se* need not necessarily worry about validity issues and potential pitfalls. Such typical applications include studies characterizing online communities and specific users, such as users self-reporting eating disorders online [56, 268] or online eating disorder support communities [57, 86]. The behaviors of interest in that case are precisely the online behaviors (i.e., the information that users choose to post). Similarly, studies developing machine learning models leveraging social media data that are *not* concerned with performance generalization beyond the platform and to the general populations are not necessarily impacted by these biases. Such applications might include social media food recognition [17, 45, 313, 403] or learning online food images embeddings [316].

Implications for food representation and users' well-being. Beyond the above implications for social media research, our results have implications for understanding the complex relationship between technology use and the well-being of social media users. In the case of food, Twitter users are exposed to unrealistic mirrors of reality [24], since foods that people actually consume and track are smaller, less “fancy”, and less visually appealing (Figure 7.6, Table 7.1). Such distortions might contribute to the high prevalence of social comparison [130], where for example, as much as one-fifth of Facebook users can recall recently seeing a post that made them feel worse about themselves [53]. A user exposed to the social media portrayal of food might therefore believe that other people consume food that is tastier than the food that they consume themselves. However, this would likely not be the case due to the discrepancy between social media foods and consumed and tracked foods. Social media might, in that case, promote an unhealthy relationship with food. Our findings have implications for research about the mechanisms of such social comparison.

7.4.4 Limitations

Next, we outline key limitations to be kept in mind when interpreting our results. In our main analyses, we study how foods are perceived by non-expert crowd workers. The extent to which expert nutritionists would agree with such participants is unknown. Furthermore, while the case study is focused on Switzerland, the crowdsourced workers are located in English-speaking countries, which might influence the food perception due to cultural factors. Although the tags provided by the participants (Table 7.1) provide insights about factors that

guide their ratings, future work should more deeply investigate the nutritional properties of the studied foods in collaboration with expert annotators. We did not include Swiss German dialect forms of keywords, as there is no written standard.

We note that the number of studied food images posted on Twitter—despite being the results of a best effort for completeness—is relatively small (around 3,700 Twitter photos of food, 2,400 of which were annotated). This number is small mostly due to the fact that we consider geolocated tweets only, and Switzerland is a relatively small country compared to the U.S., which has been studied in most related work [88, 244, 392].

We performed a case study of Twitter in Switzerland. Our findings cannot be assumed to generalize globally, and future work should apply our framework to other populations, other social media platforms and Web traces, and other food tracking apps. That said, this study may serve the purpose of a “proof by counterexample”: we have identified one common setting where there is a divergence between food consumption as measured via food tracking and as measured via social media, implying that at least one of the two is not a faithful representation of the true food consumption in general population. Hence, we should assume that there can be bias elsewhere, too. Researchers studying other populations should thus be attentive and aim to establish evidence of validity before using either social media or tracking apps as a proxy for the true food consumption in general population.

Potential sources of bias of social media traces as a proxy for true food consumption in general population. We provide grounding and first insights about validity of estimating food consumption from digital traces by contrasting consumed and tracked food with tweeted food, controlling for location, period, and food types. Revealing exact mechanisms that can lead to the biases of social media traces as a proxy for the true food consumption in general population is out of the scope of this work. However, in what follows, we consider the potential sources of bias. We postulate that biases are driven by both measurement error (related to operationalization of consumption via the concept of posting on Twitter) and population error (stemming from biases in subpopulations sharing food on Twitter) [327]. Sources of *measurement errors* might include these:

1. *Construct validity.* On the one hand, many foods that the Swiss consume are not posted on Twitter. Appealing food consumed in certain contexts is more likely to be shared, as positive and anticipated events are more likely to be disclosed on social media in general [310]. Furthermore, photos published on Twitter may be self-selected for higher quality, thus influencing how food is perceived by the annotators. On the other hand, not all foods shared on Twitter are necessarily consumed by the posting individual (especially not in their entirety, given the portion-size bias, Table 7.1). Conceivably, some of tweets may originate from promotions, restaurants, recipe sharing, all of which does not necessarily mirror actual consumption. In general, food images do not necessarily need to be related to consumption at all. They can mean something else entirely (e.g.,

a food can be a meme or a symbol of a political movement), although we did not find evidence of such biases in the studied data.

2. *Platform effects*. Numerous applications and platforms support improving image quality and editing with filters, which can all contribute to the food image being more visually appealing and appearing tastier [230].
3. *Community feedback*. Feedback received from other platform members influences the type of dietary content which a social media user posts [6], whereas negative feedback can lead to behavioral changes [63]. Biases in how food is represented online are implied by the design of online platforms.

Biases are also likely in part driven by *population error*. Users of social media platforms do not mirror the general population, neither demographically nor regarding other attributes such as behaviors and interest [214]. Users of public geotagged tweets are not randomly distributed over the general population [118, 229]. In the future, performing individual-level studies, as opposed to the population-level study reported here, will make it possible to disentangle measurement error from population error.

7.4.5 Future work

Beyond the already outlined future directions, the collected data can be used to further study patterns of sharing food online. Future work should further understand who shares food on Twitter (consumers, skilled individuals, but also non-individual agents, such as restaurants or caterers). We expect that further characterizing user types would not change our main conclusions, but would reveal how biases vary between different strata of Twitter users. Future work should further study where they share it from (residential vs. commercial areas), when, and in what context. Also, what are the predictors of engagement with food on Twitter?

Discussion and conclusion Part V

8 Discussion

8.1 Summary of contributions

The central argument of this thesis is that novel computational approaches powered by digital traces and causal science have an un-tapped potential to better understand and improve human behaviors related to nutrition. The results of our observational studies based on the analysis of digital traces—food purchase logs and information-seeking logs—demonstrate how the developed analysis frameworks make it possible to identify causal effects of interest, thereby enriching and refining the knowledge about human behaviors. However, considering that digital traces tend to be imperfect and are typically not collected with scientific research in mind, we also investigate the boundaries of what digital traces can tell us about true offline behaviors in the general population. In what follows, we summarize the main contributions.

We first investigate dietary behaviors with purchase logs, *in a situated on-campus context*. First, in Chapter 3, we introduce the transaction logs dataset and outline descriptive statistical analyses revealing, for instance, how academic schedules drive food consumption on campus on the yearly (lecture season vs. exam season vs. semester breaks) and the daily levels (lectures vs. breaks). In Chapters 4 and 5 we then address two specific research questions with the introduced dataset.

In Chapter 4, we focus on isolating the causal effect of an implicit naturally-occurring social intervention—social tie formation. We control confounds in a matched quasi-experimental design. Specifically, we identify focal users who initially do not have any regular eating partners but start eating with a fixed partner regularly. We match focal users into comparison pairs such that paired users are nearly identical with respect to covariates measured before acquiring the partner, but the two focal users' new eating partners diverge in the healthiness of their respective food choice. Contributing to the rich literature on social influences, the results of this study reveal significant shifts in the healthiness of focal users' food choice when they acquire their new eating partners.

In Chapter 5, we study social influence on campus in more detail, at the meal-level. We find evidence in favor of a specific behavioral mechanism for how dietary similarities between individuals occur—purchasing mimicry. We find significant mimicry of partners' purchases affecting all food types, and diminishing once the ordering of the purchasing queue is randomized. Moreover, we find that the effect is present across subpopulations, but strongest for students, implying that age is likely one of the key factors. The results of this study elucidate the behavioral mechanism of purchasing mimicry and have implications for understanding dietary behaviors among on-campus subpopulations.

In addition, Chapters 3, 4, and 5 demonstrate the utility of passively sensed food purchase logs for obtaining meaningful insights that can inform the design of public health interventions and food offerings, especially on campuses.

We then move on to study dietary behaviors with information-seeking logs, *in a global context*. The SARS-CoV-2 virus has altered people's lives around the world. Accordingly, in Chapter 6, we quantify the COVID-19-induced shifts in dietary interests, as revealed through time series of Google search volumes. We design quasi-experimental time-series analyses, isolating the effect of the 2020 discontinuity in mobility patterns on food interests. We find that, during the first wave of the COVID-19 pandemic, there was an overall surge in food interest, larger and longer-lasting than the surge during typical end-of-year holidays in Western countries. The largest (up to threefold) increases occurred for calorie-dense carbohydrate-based foods. These findings provide meaningful insights for governmental and organizational decisions regarding relevant measures to mitigate the effects of the COVID-19 pandemic on diet and nutrition worldwide.

In summary, the variety of studied contexts (ranging from campus-wide to worldwide) and types of traces (including purchase logs and search logs) demonstrate the wide spectrum of findings that can be derived using such causal and computational approaches. Finally, in Part III, we take a step back and rigorously investigate the boundaries of what digital traces can tell us about true offline behaviors of the general population. We contribute a novel crowdsourcing framework for estimating biases and conduct a case study of estimating food consumption in Switzerland with social media and food tracking (Chapter 7). While social media traces can be a reasonable proxy of tracked consumption for certain food types, we find that, overall, food shared on social media and actually consumed and tracked food significantly diverge from each other. This divergence implies that at least one of the two—the tracked food or the one on social media—is not a faithful representation of the actual food consumption in the general Swiss population. These findings warrant researchers' attention and highlight the need to establish evidence of validity before using digital traces as a proxy for true food consumption in the general population.

8.2 Significance

In this section, we summarize how our contributions can inform policy-making, and discuss and the implications of our work beyond dietary behaviors.

Implications for policies and behavioral interventions. Taken together, our results demonstrate that analyses of passively sensed digital traces can provide meaningful insights for policies, public health interventions, and decision-making. For instance, on-campus analyses of purchase logs can support efforts to promote sustainable and healthy habits by modifying food offering, e.g., during exam sessions, while being aware of complex spatio-temporal variations and differences among various subpopulations (Chapters 3 and 5). Similarly, analyses of purchase logs can anticipate the impact of planned interventions by estimating the effect of similar naturally occurring interventions (Chapters 4 and 5).

Furthermore, monitoring population-scale online information-seeking traces provides stakeholders with timely and detailed insights into the circumstances that individuals face during a crisis (Chapter 6). Similarly, analyses of purchase and information-seeking logs can provide a better understanding of customer behaviors and, accordingly, inform ways of improving corresponding services. For instance, the results of purchase log analyses can support efforts to optimize logistics, avoid overcrowding, and staffing decisions on campus (Part II). Along similar lines, the analyses of information-seeking logs can shed more light on global consumer preferences during a crisis (Part III).

Implications beyond food. To the best of our knowledge, observational studies reported in Chapters 4 and 5 are the first to use large-scale transactional data to retrospectively evaluate the impact of implicit dietary behavioral interventions on dietary behaviors. Along with showing how food purchase logs can be used as effective sensors to detect behavioral changes, our results also highlight how careful quasi-experimental comparisons can be used to measure complex interactions. For example, the methodology of isolating influence through matched pairwise comparisons (Chapters 4 and 5) makes it possible to detect behavioral changes in areas relevant to other scenarios, such as exercise habits, coaching or counseling, traveling, sleep patterns, online interactions, collaborative writing, or coding practice, among many others. Ambitious questions on how individuals impact others' behaviors or skills could all be tackled by identifying the onset of such interactions and comparing outcomes in similar matched persons, while controlling for confounds.

Finally, beyond food and dietary behaviors, our crowdsourcing bias estimation framework (Chapter 7) can be used to measure biases of digital traces for studying many types of behaviors, including but not limited to politics, activism, behaviors important for health and well-being that are frequently shared online, such as fitness and time spent in nature, as well as travel, fashion and aesthetics, socialization, or pet ownership. Addressing the questions of truthfulness and validity of digital traces beyond food is an important direction for future research.

8.3 Advantages and limitations of studying diets locally vs. globally

In this thesis, we present contributions observing behaviors in two drastically different contexts: while Chapters 4 and 5 study diets in a situated on-campus context, on an individual level, Chapter 6 explores behaviors in a global context, on a population level. Along with providing meaningful insights, both setups are also subject to their unique shortcomings. In what follows, we first consider the major benefits and drawbacks of leveraging situated behavioral traces, such as campus-wide purchase logs.

Strengths of campus-wide purchase logs. A prominent strength of studying individuals' purchase logs in a situated context is that such investigations make it possible to perform *detailed interpretable analyses* and consistently *track individuals*. Access to contextualized knowledge about specific circumstances that individuals are immersed in offers crucial advantages. Conversely, patterns present in aggregated online traces are typically far more cryptic and, due to the loss of contextualization, rarely allow for a definitive interpretation.

A major lesson learned through our studies of campus-wide behaviors is that the researcher analyzing the purchase logs cannot assume the role of a friendly outsider. Instead, the researcher should be an active part of the campus ecosystem and acquire domain knowledge. In what follows, we provide a couple of examples of how, in our studies, we benefited from having the respective information by being embedded into the campus.

- It proved beneficial to know the personnel who collected the purchase data, since they were able to share valuable insights about the nature of the data and its limitations. For instance, in the analyses presented in Chapters 4 and 5, we benefited from talking to campus food-providing administration and transaction system managers in order to understand the information about the food items and restaurants encoded in the dataset using the terms that students and staff do not commonly use.
- It was also helpful to be aware of the opening of shops that are near the campus, but absent in the purchase logs. Without such background knowledge about the campus food landscape, it would have been challenging to interpret variations across years, as discussed in Chapter 3, Section 3.2.4.
- We benefited from being familiar with the academic schedules, timetables, and past deviations and circumstances. Overall, yearly patterns reveal regularities that are consistent with academic schedules. However, some regularities tend vary between different years, as, for instance, Easter break occurs on varying dates in different years. We were aware of such specific campus schedules and holidays in Switzerland and could interpret the regularities.

Shortcomings of campus-wide purchase logs. At the same time, the most prominent limitation of characterizing behaviors in a situated context is the lack of *universality*. It remains

unclear to what extent the observations hold in other similar setups, or the general population. However, it should be noted that the transfer to the general population is not necessarily an issue. Since many people regularly consume foods on campuses, findings will typically have actionable implications, even if not necessarily relevant to the general population.

Strengths of worldwide information seeking traces. Similar trade-offs emerge in research leveraging global online traces, outside of a situated context of a campus. Online information-seeking traces offer certain advantages. Most notably, analyses of online logs benefit from *rich background knowledge, are not confounded by the local context*, and can *scale*. With regard to the first advantage, online traces are often curated, meaning that the collected information is typically already rich and conveniently structured. For instance, search logs can be linked to Wikidata knowledge graph entities, allowing researchers to leverage already available rich background knowledge about food. In Chapter 6, we use such knowledge to derive taxonomies of food concepts. Conversely, analyses of campus food purchases have to “start from scratch”. Item labels are unstructured, and the information derived from them is incomplete. Therefore, there is a need to work closely with nutrition experts to develop custom annotations and healthiness estimates, as we did in Chapter 4. Regarding the second advantage, online behaviors are generally not constrained by the situated context—indeed, people can search for the entire universe of food online. In contrast, on campus, people can only buy and consume what is offered, which makes availability a major confounding factor. Finally, the third and perhaps most important advantage of online information-seeking traces is that they enable monitoring a larger number of people. As a result, it becomes possible to scale the analyses, from 39,000 unique users present in the campus-wide purchase logs to billions of Internet users worldwide.

Shortcomings of worldwide information seeking traces. Yet, an important shortcoming of online information-seeking traces is *weak construct validity*. It should be acknowledged that online traces are distant proxies for the true behavior of interest: food consumption. Said differently, searching for food is not tantamount to consuming the food. Indeed, users may search but not consume, and vice versa. Admittedly, the purchase log data do not directly capture food consumption either, but rather provide indirect proxies via purchasing. In fact, there is no certainty that the items purchased on campus were consumed. Conversely, other food items that were not purchased might have been consumed (e.g., a soda brought from home). However, compared to information-seeking traces, purchase logs are more closely tied to offline behaviors. In the case of online traces, construct validity is more difficult to ensure and assess. Moreover, online information-seeking logs make it problematic to track individuals, as they are typically available to researchers only in aggregated form due to privacy concerns. In addition, event counts frequently require a careful calibration to allow comparisons (cf. Chapter 6).

Overall, observations made in situated contexts and observations made outside of a situated context are mutually complementary and equally valuable. In what follows, we discuss the challenges in further detail and identify future directions for addressing them. In particular,

we outline promising future avenues to explore in order to ensure that on-campus studies can generalize and that online traces can have stronger validity guarantees.

8.4 Future opportunities

The contributions of this thesis and the outlined limitations open the door for a number of future research directions. We outline some of the concrete research directions we envision.

8.4.1 Situated campus-wide dietary behaviors: The next frontiers

We identify the following four specific areas for future work: (1) deriving generalizable insights about on-campus dietary behaviors, (2) collecting more detailed on-campus food offering and consumption data, (3) engaging with all on-campus stakeholders to discover the needs and priorities regarding on-campus food offer and consumption, and (4) developing new principles and practices around ethics and privacy. For each of these four areas, we identify a set of concrete possible solutions and avenues for future work.

Deriving generalizable insights about on-campus dietary behaviors. Each campus is an independent eco-system, which makes it challenging to derive findings that hold between campuses. Estimates are typically produced at different times and different locations. Furthermore, cultural factors can, in important ways, alter behaviors—for instance, due to varying susceptibility to stress. What behaviors can generalize across campuses around the world?

Future work: Data sharing between campuses. Future efforts can include creating a network of partner institutions that would enable researchers to replicate the same set of analyses and then perform meta-analyses to discover universal behavioral patterns. In particular, we envision developing a system to enable processing the anonymized purchase logs (in a pre-determined format) and obtaining and sharing aggregated high-level insights with other campuses. An a priori designed meta-analysis would then be performed across a cohort of campuses. Such “megastudies” have the potential to improve the evidentiary value of behaviorally informed policy interventions [247]. A unified client-side analysis framework could be built to process the anonymized logs, following templates of studies that can be conducted worldwide in order to answer pre-agreed research questions, such as the ones outlined in Chapter 3, Section 3.3. In the future, we also envision organizing an event (e.g., as a workshop) where researchers studying on-campus behaviors can convene, share ideas and expertise, and agree on priorities.

Collecting more detailed on-campus food offering and consumption data. A major challenge of the efforts to study on-campus food consumption is the lack of transparency regarding food offer and consumption. On campuses, complex sets of factors beyond the knowledge and reach of individuals who consume food often determine the availability of options, the nature

of collected data, how it is used, and by whom, resulting in a lack of transparency regarding food on campus. Moving forward, there is a clear need for more robust and open policies.

Future work (a): More detailed passively sensed data collection. Since institutions are responsible for the health of everyone on campus, there is a pressing need to collect more data about offered food. Future efforts can involve collecting rich information about the food items, including but not limited to their origin, ingredients, nutrients, calories, sustainability metrics (e.g., carbon footprint), preparation methods, and food waste statistics. There is also a need to enrich purchase logs with detailed information about the sold items, with a particular focus on information relevant to health and sustainability.

Future work (b): Digital cohorts and active data collection. Another potential solution to the concern about the lack of transparency regarding food on campus is setting up digital cohorts where students and staff could share their food consumption and other health-related data for research purposes. Future avenues should involve letting individuals on campus take ownership of their data and share it with researchers, if they wish to do so. In particular, the individuals could share information about specific conditions closely linked with dietary habits (e.g., diabetes and heart diseases), as well as regarding behaviors assumed to be linked with food consumption, but not well understood. The latter could include, for instance, information about major life events, daily habits, social media usage and web browsing, menstrual cycle, and mental health. Active data collection efforts are required to make it possible to answer such ambitious research questions with implications for the health of the general population. Furthermore, such efforts can allow linking diets and other on-campus behaviors and outcomes, such as social interactions, mobility [251], and academic performance [198].

Further engaging with on-campus stakeholders to discover needs and priorities. Since campuses worldwide face a lack of unified effort towards health and sustainability, there is a need for a closer examination of structural processes of power at the campuses we are part of, and a need for a better understanding of factors that slow down current efforts towards health and sustainability. In addition, there is no principled and systematic understanding of what individuals on campuses worldwide want and need, as well as difficulties and challenges they face. Future efforts should explore experiences and perspectives from unheard subpopulations of students and essential workers who consume large amounts of food on campuses but traditionally do not play a role in the decision-making.

Future work: Participatory approaches. To tackle the challenges mentioned above, new qualitative and participatory approaches are needed. It is necessary to have a direct contact and engage with both food consumers and providers, i.e., local stakeholders (restaurants, vending machine operators, transaction managers, administration). It is equally necessary to engage with individuals (students and staff) who, in one way or another, participate in the campus food system, are impacted by any implemented policies and interventions, and would bene-

fit from prospective scientific findings. To this end, in future research efforts, participatory designs [344] and qualitative interviews would be particularly suitable methodologies.

Similarly, there is also a need to acknowledge that food offering, the collected data, and the derived insights embody and reproduce the values of those who designed the food offering systems and those who build and deployed data collection mechanisms to begin with [129]. An important point to make is that the values of thousands of individuals who are part of the campus should be discovered, not assumed [216]. Therefore, our responsibility as researchers is not only to strive to use such data to promote assumed shared values of social good (health and sustainability), but also to discover the priorities and values of the people who are part of the underlying system. For instance, it should be taken into account that students may prioritize economic concerns over sustainability concerns.

Overall, the goal of such participatory approaches is to enable individuals on campuses to identify their own priorities and make decisions about the food system and about the future use of digital traces they contribute. In the process, an overarching goal is to give a voice to all stakeholders and enable everyone involved to answer questions such as “Is the food offering optimal for our needs? How is our data being collected? Will the findings be of benefit to us? How will we use them?” [140].

Developing new principles and practices around ethics and privacy. Since the behavioral data and the corresponding analyses can be misused or abused, it is necessary to balance the potential to do good with the potential to harm. New norms for privacy and ethics and best practices are needed. At present, corresponding institutional approvals are needed to perform analyses within one campus. Such institutional processes help guide researchers through ethical and privacy concerns, as purchases can reveal potentially sensitive information about individuals [173]. However, new challenges emerge when there is a need to replicate the same analyses across several campuses so as to derive generalizable insights about dietary behaviors. Having all the data, which is potentially sensitive, at a single central point is a risk and a liability since the data can be compromised in the event of a breach. Privacy considerations are necessary to design future systems for processing the anonymized purchase logs. It is equally important to ensure that the high-level findings do not reveal anything about specific individuals at a given campus and to protect potentially sensitive information. In what follows, we briefly outline a variety of possible solutions to these concerns.

Future work (a): Processing transaction logs across silos. While it is reasonable to assume that the owners of a processing system cannot be trusted, if there is an agreement about data formats, one could locally run universal processing scripts. Future efforts can involve designing a web application that would allow locally running purchase log analysis scripts with embedded privacy mechanisms. Previously, decentralized data processing across silos has been successfully deployed in settings where silos corresponded to hospitals processing medical datasets [19]. Such processing is consistent with the paradigm of federated learning, a privacy-enhancing technique based on the idea of keeping and processing the data in centers

of origin. Such processing increases privacy and ensures that institutions keep control of the data. Furthermore, to match confidentiality requirements, federated learning can be used with other privacy-enhancing technologies, such as differential privacy (DP). Such a mechanism ensures that from the results of the analyses, one cannot infer additional information about the original data [106, 107].

Future work (b): Privacy-preserving synthetic purchasing traces. When the input dataset contains a large number of attributes, DP mechanisms require injecting a prohibitive amount of noise compared to the signal in the data, which can compromise the usefulness of the data [411]. To facilitate collaboration over sensitive data, an alternative approach is to take a sensitive dataset as input and generate a structurally and statistically similar synthetic dataset with strong privacy guarantees [193, 277]. In that case, data owners need not release their data, while collaborators can perform data analyses. Such synthetic traces can then be directly shared with other researchers to speed up analyses. The benefit of this solution is that there is no need to regenerate synthetic data for analyses that were not initially envisioned.

Future work (c): Trusted execution environments. Trusted environments [308] enable processing the logs on a dedicated secure subsystem, i.e., a trusted execution environment isolated from the main processor of the central server. The benefit of this approach is that it makes it easier to update the computation centrally, without needing to update the code run by each participating institution.

8.4.2 Worldwide online dietary behaviors: The next frontiers

Construct validity of online traces. As previously discussed, one of the most pressing challenges when studying diets with online traces is ensuring the construct validity of such methods. In Chapter 7, we proposed a framework for measuring the biases of digital traces *at population level*. The findings demonstrated that, in the case of estimating food consumption in Switzerland with social media images, lower bounds to the biases of social media traces are statistically significant. These findings highlight the need to further determine and improve the validity of digital traces. Moving forward, it is necessary to link the other distant proxies such as search queries studied in Chapter 6 with the actual food consumption, *on the individual level*. The methods and findings relying on search logs would have improved validity and reliability if it were possible to answer questions such as “What fractions of food searches are linked with the consumption of the searched food? What fraction of consumed foods are searched online?”.

Future work: Linking online traces and offline behaviors. To answer such questions, it is necessary to link the ground truth food consumption (measured, e.g., via reliable tracking or situated purchase logs) with an individual’s online traces (e.g., search history) and analyze them jointly. Although data sharing initiatives making it possible for individuals to share data with researchers exist, they are still in their infancy. Beyond linking online and offline behaviors, such individual-level online traces are valuable for scientific studies. In the future,

it is necessary to develop new ways to empower individuals and communities to contribute their data for scientific insights. Individuals could then participate in the scientific process and even in the selection of the research questions, so as to benefit from the findings based on data they contribute. Sensitive data may in that way be exchanged for services to individuals and the society. Interestingly, a recent study by Gefen et al. [136] demonstrates that individuals would be willing to contribute—99% of people were willing to contribute their online data in exchange for monetary compensation and an analysis of their data, while 53% were willing to pay to have their online data analyzed. Such data-sharing efforts can also help to change the balance of power between the public and tech companies that rely on implicit and explicit data contributions from the public [374].

8.4.3 Studying diets with digital traces: The next frontiers

Improved data representation is necessary to unlock new applications. In this thesis, foods were represented as items available on campus (cf. Chapters 5 and 4), concepts people can search for and read about (cf. Chapter 6), or food types that people log or post online (cf. Chapter 7). In these studies, the foods are mostly considered in isolation, as rich information about the ingredients or preparation is largely unavailable. However, in reality, concepts related to diets—e.g., foods, ingredients, recipes, restaurants, or purchases—interact in complex ways. For instance, people combine only certain foods within a meal; within a meal, only certain ingredients are combined, and, eventually, the ingredients are combinations of certain compounds.

Future work: Graph-based representations. Relying on models capable of accurately representing concepts related to diets and the complex interactions between such concepts can unlock new applications to help promote health and sustainability and change how our food is designed and prepared. For instance, Ahn et al. [10] studied the flavor network that captures the flavor compounds shared by culinary ingredients. Rich interactions between ingredients revealed novel insights about our tastes—Western cuisines tend to use ingredient pairs that share many flavor compounds, supporting the so-called food pairing hypothesis, while, in contrast, East Asian cuisines typically avoid compound-sharing ingredients. Furthermore, along with nutrients, food also contains bioactive molecules, some chemically similar to anti-cancer drugs. For instance, Veselkov et al. [372] used deep learning on protein-protein and drug-protein interaction graphs between food molecules and molecules in our body to identify which foods contain ingredients that can prevent diseases. Rich representations through modern machine learning techniques can help provide a more comprehensive view of our dietary behaviors, culinary practices, and design better products.

Predictive vs. explanatory approaches to studying traces. Finally, throughout the thesis, we make an argument for and use an *explanatory* data analysis approach—the studies presented here aim to identify and estimate causal effects of interventions (i.e., staying at home or starting to eating together with someone) on human dietary behaviors (i.e., purchasing or

searching for food). In contrast to this approach, an orthogonal path frequently taken in data-driven research is a *predictive* one, focused on making and testing predictions. A predictive data analysis approach would instead ask: “How predictable are human dietary behaviors under intervention?”. Such predictive approaches aim to develop models that can accurately generalize to unseen samples, while not necessarily being concerned with interpretability and explaining the mechanisms. However, answering the question of predictability brings additional information and can help to inform the design of policies and interventions.

Future work: Integrative modeling. Future directions to understand and improve dietary behaviors should consider integrating predictive and explanatory modeling—for instance, by measuring the extent to which specific causal estimates made in one domain transfer to another domain, or by verifying whether the predictive models fitted in one domain generalize to another, as Hofman et al. [174] propose. Such integrative modeling will likely facilitate more valid and impactful findings on human dietary behaviors.

8.5 Connections with further thesis work

The contributions presented in Chapters 3–7 are not exhaustive of the thesis work. In what follows, we briefly outline further thesis work (Figure 8.1). Although further thesis work is not specifically on the topic of dietary behaviors, there are rich connections with the contributions of the thesis. We summarize five further contributions and the connections with the thesis work.

COVID-19-induced shifts in general interests. First, in further thesis work [291], we study the impact of COVID-19-induced mobility restrictions on people’s interests in general, beyond foods. In this work, we demonstrate how the approach introduced in Chapter 6 can generalize to information-seeking behaviors beyond food, by studying different type of in-

Further thesis contributions	Digital traces	Methods
COVID-19-induced shifts in general interests <i>Sudden Attention Shifts on Wikipedia During the COVID-19 Crisis.</i> Manoel Horta Ribeiro*, Kristina Gilgoric*, Maxime Peyrard*, Florian Lemmerich, Markus Strohmaier, and Robert West. International AAAI Conference on Web and Social Media ICWSM, 2021. * Equal contributions.	Information seeking logs: Wikipedia access logs	Observational study: Difference-in-differences design
How platform design impacts linguistic features of social media content <i>How Constraints Affect Content: The Case of Twitter’s Switch from 140 to 280 Characters.</i> Kristina Gilgoric, Ashton Anderson and Robert West. International AAAI Conference on Web and Social Media ICWSM, 2018.	Social media posts	Observational study: Matched design
Anticipated versus actual effects of platform design <i>Anticipated versus Actual Effects of Platform Design Change: A Case Study of Twitter’s Character Limit.</i> Kristina Gilgoric, Justyna Czesochowska, Ashton Anderson, and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing CSCW, 2022.	Social media posts	Longitudinal statistical analyses
How brevity impacts success of social media content <i>Causal Effects of Brevity on Style and Success in Social Media.</i> Kristina Gilgoric, Ashton Anderson and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing CSCW, 2019.	Social media posts	Experimental crowdsourced study
How linguistic features impact success of news headlines <i>Linguistic effects on news headline success: Evidence from thousands of online field experiments (registered report protocol).</i> Kristina Gilgoric, George Lifchits, Robert West and Ashton Anderson. PLOS ONE, 2021.	Information seeking logs: Web browsing logs	Field experiments (A/B tests)

Figure 8.1: Outline of the further thesis contributions discussed in Section 8.5. For each contribution, we summarize the leveraged digital trace data and the main methods.

formation seeking logs—namely, Wikipedia access logs. We find that the impact of the crisis on information-seeking behavior persisted beyond mobility restrictions and that many of the topics with persistent increases relate not to basic needs pertaining to the pandemic, but to entertainment and self-actualization (e.g., VIDEO GAMES). These findings highlight the utility of online digital traces for research on how the pandemic has affected not only people's diets, but also their needs, interests, and concerns. This work also demonstrates the potential of various web access logs (beyond search logs studied in Chapter 6) for studying dietary behaviors at scale.

Message framing and the causal effects of linguistic properties. Moreover, further thesis work is connected with nudges toward healthy eating, which must be carefully framed and phrased such that they are maximally effective [23, 365]. In further thesis work on message effectiveness we aim to discover such strategies. In particular, we design observational and experimental studies to identify the causal impact of platform design and message framing on the success of textual content. The study designs and the derived scientific findings are connected with designing dietary interventions and answering questions such as *How to make the language of sustainable and healthy food appealing?*

First, studying how platform design impacts linguistic features of social media content [145], we take advantage of the platform change and find, when subject to the 140-character limit, users write more tersely. Designs taking advantage of such exogenous shocks to study human behaviors can be effective since they minimize the impact of confounding factors. Similarly, in Chapter 6 we took advantage of the shock of COVID-19-induced interventions on human behaviors. Moving forward, unexpected policy changes, availability shocks, and external events provide researchers with valuable opportunities to study determinants of dietary behaviors.

Second, further studying the anticipated and actual effects of the platform design change on the linguistic features of tweets [147], we find that the fusion of design decisions and human behaviors can lead to feedback loops and calls for more cautious approaches that aim to take into account the dynamic nature of people's responses. Beyond platform design, these findings have implications for the design of dietary interventions and policy changes (e.g., regarding how food is offered on campus). These findings highlight the fluidity of human behaviors and have direct implications for studies predicting the impacts of hypothetical policies attempting to change dietary behaviors.

Third, we focus specifically on identifying the causal effects of brevity on the style and success of social media messages [144]. Whereas most prior work has studied the effect of wording on style and success in observational setups, we conduct a controlled experiment, in which crowd workers shorten social media posts to prescribed target lengths and other crowd workers subsequently rate the original and shortened versions. In contrast with observational approaches leveraging passively-sensed data developed in this thesis, experimental crowdsourced studies have their advantages. Most notably, the researcher fully controls the setting in which the

behaviors are measured and can perform randomization. Moreover, crowdsourcing platforms make it possible to recruit participants worldwide and thus scale the analyses beyond the campus-wide context. In the future, controlled experimentation involving rating tasks and pairwise comparison tasks (as performed in Chapter 7) is a powerful way to elicit dietary preferences and study diets.

Finally, in a study aiming to study linguistic properties beyond social media [149], we conduct analyses of thousands of online A/B tests of headline variations performed by Upworthy (a once-popular social news outlet) to identify how linguistic features impact the success of news headlines. Our registered findings contribute to resolving competing hypotheses about the linguistic features that affect the success of text and provide avenues for research into the psychological mechanisms that are activated by those features. The analyses of data generated via A/B testing fuse the advantages of observational and experimental approaches—they allow researchers to randomize assignment, while the behaviors are still observed in an organic setting. Such an approach and the analysis framework for estimating linguistic effects can be used to understand how to improve dietary behaviors by modifying the language of food.

8.5.1 How to make the language of sustainable and healthy food appealing?

Text-based nudges and causal effects of their linguistic properties. A promising new avenue fusing the contributions of the further thesis work with the main argument of this thesis is to explore how to use the language of food and linguistic cues [195] to promote healthy and sustainable dietary behaviors on campus and beyond. Despite the impact of diet on the environment and health, sustainable and healthy food is often perceived as unappealing [284, 369]. To promote sustainable and healthy diets, it is necessary to address the following questions: *How to make the language of sustainable and healthy food appealing? How to name and word healthy and sustainable dishes to make them sound attractive and tasty?*

Future work: Causal analyses of the language of food. Novel computational approaches are necessary to isolate such causal links between linguistic properties and the success of behavioral nudges, e.g., between the dish name and its purchase probability. Research at the intersection of NLP and causal inference [108, 113, 178, 389] aims to approach the challenge of estimating causal effects in settings where text is used as an outcome, treatment, or as a means to address confounding [201]. Combining quasi-experimental methods that aim to isolate causal effects while controlling for confounding factors with modern text representations is a promising way forward. The ability to answer these questions would have clear implications for how dishes are labeled on-campus and for how to describe healthy and sustainable food in general such that it is perceived as appealing.

9 Conclusion

The contributions of this thesis are motivated by the idea that causal computational approaches leveraging digital behavioral traces can shed new light on human dietary behaviors. Closing the circle, we weigh up the general conclusions, lessons learned, and paths forward.

The results of presented studies offer evidence in favor of **feasibility** of behavioral changes via social interventions, and by leveraging disruptions due to crises and exogenous shocks. Beyond food and dietary behaviors, the methodological contributions—the bias estimation framework and quasi-experimental study designs—can be applied to studies of human behaviors in efforts to mitigate **measurement** and **identification** issues, respectively.

In particular, the presented studies make a clear case for the usefulness and potential of digital traces. Throughout, the connecting thread is that the idea of repurposing—repurposing the data not meant for scientific research and repurposing unpredictable variation in the world around us—can lead to powerful ways to study complex phenomena. The scientific insights we contribute have immediate implications and can inform policy-making and behavioral interventions. The direct future frontiers leveraging digital trace data include further establishing and improving the validity of traces, developing new paradigms for data sharing and deriving universal insights, further discovering the needs, values, and priorities regarding food offer and consumption, and developing new principles and practices around ethics and privacy.

The presented studies also made a clear case for adopting causal approaches to analyze digital traces. Causal approaches allow approximating the experiments we cannot conduct and are a particularly good fit for making sense of the noisy and imperfect digital traces. Fortunately, causal approaches tend to benefit from the strengths of digital data. Most importantly, being “big” allows the identification of strictly comparable units among a large pool. For instance, hundreds of strictly comparable individuals are identified among tens of thousands of individuals on campus to study the impact of tie formation. Moreover, the “always-on” property of digital data enables the identification of exogenous shocks and retroactive analyses of unintended randomizations. For example, real-time measurements enabled the identification

of the effect of COVID-19-induced shocks on dietary behaviors. Causal approaches let us minimize the impact of biasing factors and digital trace limitations while being transparent about the assumptions and the unobserved biases, carefully considered via causal diagrams and sensitivity analyses. Consequently, our trace analyses progress far beyond merely restating the “correlation is not causation” mantra.

Taking a glimpse beyond conducting a stream of empirical investigations, the final aspiration is to use the digital traces to inform future developments of new data-first, empirically grounded theories. For instance, regarding the campus-wide setting, there is a need for new theories of and insights into the social, psychological, behavioral, and organizational phenomena that surround dietary behaviors in campus contexts, be they educational, corporate, industrial, or medical campuses. This thesis merely scratches the surface of the range of future possibilities by demonstrating the potential of computational approaches based on trace data.

At present, much about human behaviors around diets and their relation to health and the environment remains unknown. We are largely in the dark regarding the precise, complex ways how our dietary behaviors interact with our mental health (e.g., what are the links between nutrient deficiencies, dietary inflammation, and mental illnesses [120, 213]), our genetics, pathogens and the gut microbiome (e.g., what are the links between genetic predispositions, microorganisms, and food preferences [226]).

Looking ahead, in one way or another, truly resolving causality is key to understanding the determinants and implications of our dietary behaviors. Evidence reported in this thesis demonstrates that we should aspire to reach complete answers to such ambitious questions by leveraging various forms of trace data we leave behind throughout our lives.

A Purchasing mimicry in food consumption on campus

A.1 Supplementary information and robustness checks

We provide supplementary information and perform analyses that support our main conclusions or provide complementary insights. In Tables A.1, A.3, and A.2, we present co-eating matrices that outline the dyad frequency among the subset of the studied situations with available demographic data. The tables illustrate a preference for eating with others of the same gender, age, and status.

In Figure A.1, we present the effect estimate among the sub-population with available demographic information. The estimated risk difference across the matched situations is shown separately, depending on the individuals' status, age, and gender. Lastly, in Figure A.2, we present the relative version of the main findings by measuring the relative risk.

A.2 Supplementary analyses

A.2.1 Effect estimate under varying assumptions

We consider how our estimation framework and the subsequent estimates vary as Assumption 1 is violated. The alternative DAGs capture the relaxed assumptions. In Figure 5.2, Figures (b), (c), and (d), we illustrate the variations of the assumed causal relationships where the Assumption 1 is violated such that the traits of the individuals can influence the observed purchasing behavior through factors not related to friendship strength $S_{a,b}$. For a set of plausible variations, we derive the minimal sufficient adjustment set according to the backdoor criterion [153]. In particular:

1. Allowing the partner's eating profile to influence the focal person's manifested behaviors through factors not related to friendship strength (Figure 5.2b), the minimal sufficient adjustment set of variables for estimating the total effect of $Y_a(t)$ on $Y_b(t)$ is $\{X_a, P(t)\}$.

2. Allowing the focal person's eating profile to influence partner's manifested behaviors through factors not related to friendship strength (Figure 5.2c), the minimal sufficient adjustment set of variables for estimating the total effect of $Y_a(t)$ on $Y_b(t)$ is $\{X_a, X_b, P(t)\}$.
3. Allowing both eating profiles to influence both manifested behaviors through factors not related to friendship strength (Figure 5.2d), the minimal sufficient adjustment set of variables for estimating the total effect of $Y_a(t)$ on $Y_b(t)$ is $\{X_a, X_b, P(t)\}$.

Since the scenario depicted in Figure 5.2b is already addressed by our main analysis, we investigate how robust estimates are when situations are additionally matched on the focal person's identity to control for X_b (necessary in variations depicted in Figure 5.2c and Figure 5.2d). When additionally matching on focal person identity, we obtain the overall risk difference of 13.35% and findings qualitatively similar findings to our main analyses. Overall, the conclusions are robust to this choice.

A.2.2 The impact of social tie strength

We further investigate the impact of social tie strength. Social tie strength $S_{a,b}$ is operationalized by calculating the fraction of instances when the pair is eating together, out of all instances when either one is observed eating with someone. In Figure A.3, we demonstrate that risk difference and risk ratio estimates are the greatest for the highest values of social tie strength. However, the estimates are significant in all strata of social tie strength, and, at minimum, focal persons are estimated to be +10% more likely to purchase the food item when the partner purchases vs. not.

A.2.3 Coordination hypothesis

An alternative hypothesis explaining the observed similarities between adjacent persons in the purchasing queue is that the two persons coordinated to go for a meal together and agreed on the food choice before lining up in the purchasing queue. We investigate the presence of such coordination.

There are 226 pairs of persons A and B such that that are at least ten matched pairs of situations in order A-B and at least ten matched pairs of situations in order B-A. For each pair, we independently test the coordination null hypothesis, namely, that the order A-B or B-A does not matter since similarities come from coordination before making a choice. Under the null hypothesis, people agree on what to eat before lining up in the queue, so the order of how they go (A-B or B-A) is arbitrary. When the partner purchases an item, the focal's probability of purchasing is the same in two orders since the persons pre-agreed, i.e., purchasing probability does not depend on the order.

Concretely, there is a set of matched situations in A-B order and a set of matched situations in B-A order. We calculate the purchasing probability of the focal person in the two sets

Appendix A. Purchasing mimicry in food consumption on campus

and test the null hypothesis that they are the same. Using a two-sided t-test, we reject the coordination null hypothesis in 66 pairs out of 226 (29%), with a $p < 0.05$ threshold. Based on this investigation, we argue that it is unlikely that pre-purchase coordination can entirely explain the measured effect in all pairs. Since situations are matched, differences in ordering likely stem from different mimicry exhibited by person A and person B when they are the focal person vs. the partner.

Table A.1: Gender co-eating matrix with condition frequency among the subset of the studied situations with available demographic data. In rows, the gender of the focal person, and in columns, the gender of the partner.

Partner Focal person	Female	Male
Female	57.00%	10.65%
Male	12.67%	19.68%

Table A.2: Status co-eating matrix with condition frequency among the subset of the studied situations with available demographic data. In rows, the status of the focal person, and in columns, the status of the partner.

Partner Focal person	Staff	Student
Staff	30.05%	9.63%
Student	9.83%	50.49%

Table A.3: Age co-eating matrix with condition frequency among the subset of the studied situations with available demographic data. In rows, the age of the focal person. In columns, the age of the partner.

Partner Focal person	<=22	23-32	>32
<=22	25.66%	2.34%	1.24%
23-32	2.67%	14.21%	9.45%
>32	0.02%	7.01%	37.40%

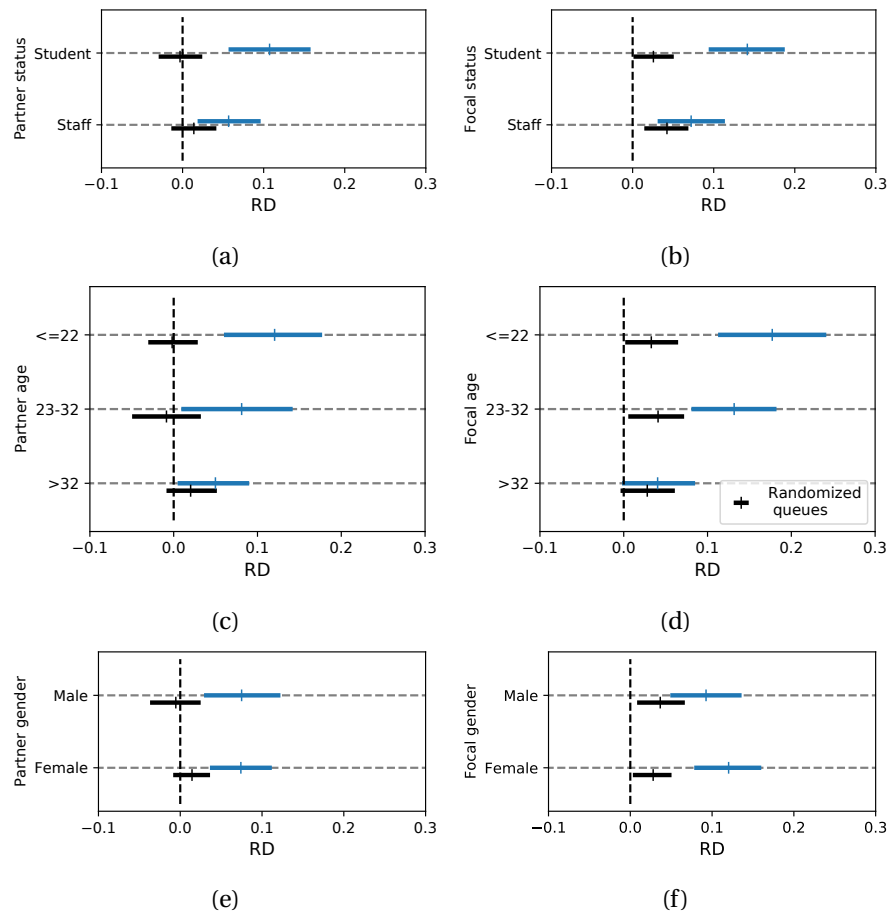


Figure A.1: The estimated risk difference across the matched situations (on the x-axis), depending on the individuals' status, age, and gender (on the y-axis). The error bars mark 95% bootstrapped CI. Risk difference estimates are presented in blue. The randomized baseline is presented in black. In (a) for partner's and in (b) for focal person's status. In (c) for partner's and in (d) for focal person's age. In (e) for partner's and in (f) for focal person's gender.

Appendix A. Purchasing mimicry in food consumption on campus

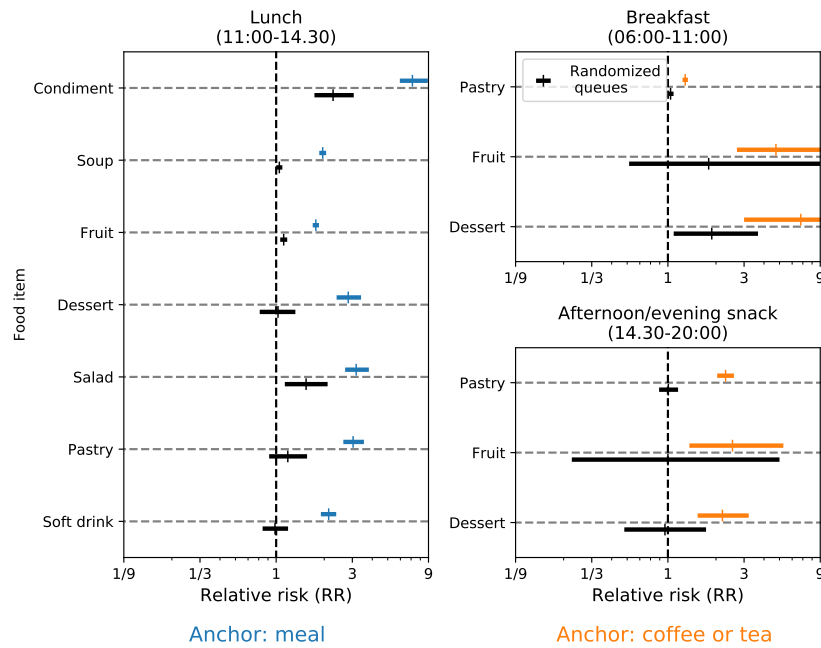


Figure A.2: Separately for lunch, breakfast, and afternoon or evening snack, the estimated risk ratio (on the x-axis), for the different food item additions (on the y-axis). The error bars mark 95% bootstrapped CI. Relative risk estimates are colored (blue for lunch where the anchor is the meal, orange for breakfast and afternoon or evening snack where the anchor is a beverage). The randomized baseline is presented in black. Note the logarithmic x-axis.




Figure A.3: In (a), the histogram of the social tie strength between the focal person and the partner. In (b), the risk difference estimate within the subset of matched situations (on the y-axis), with the given social tie strengths (on the x-axis). In (c), the risk ratio estimate within the subset of matched situations (on the y-axis), with the given social tie strengths (on the x-axis). The shaded areas mark 95% bootstrapped CI.

B COVID-19-induced shifts in dietary interests

B.1 Supplementary information and robustness checks

In Figure B.1, we present the detected mobility decreases and increases in the 18 countries. Table B.1 summarizes the descriptions of food categories and contains examples of popular foods in each category. We list all fitted coefficients and statistics of our main model in Table B.4. We provide correlation plots with Pearson correlation coefficient, instead of Spearman rank correlation coefficient in Figure B.2. Next, we provide our main results obtained with the RDD model with varying design choices and confirm that the qualitative interpretations of the effects remain stable under a number of robustness checks.

B.1.1 The impact of model order

We show our main results with a linear model in Figure B.3, and in Figure B.4 with a constant model, instead of a quadratic model. While quadratic and linear models let us estimate the short-term effect, with the constant model we estimate the average effect in the entire period, from the discontinuity, until the bandwidth (K_2) weeks after discontinuity. The estimates of the effect with the constant model are then lower because the weeks when the effect diminishes are taken into account to calculate the average (see illustration in Figure B.7). While the nature and the amplitude of the estimated effect vary (i.e., whether the immediate short-term boost of average boost is captured), most of the conclusions are robust to this choice.

Additionally, for each category of food items, we fit a slightly different model pulling the interest volume across different countries, but with an added country-specific offset, that lets us measure effect across all countries. In Table B.5, as a robustness check, we show the food categories ranked by effect size pulled across countries, estimated with a constant, linear, and quadratic model. The rank between categories is strongly correlated (Spearman rank correlation 0.95 ($p = 3.7 \times 10^{-14}$) between constant and linear, 0.89 ($p = 3.2 \times 10^{-10}$) between constant and quadratic, and 0.94 ($p = 8.7 \times 10^{-14}$) between linear and quadratic models).

In Figure B.5, we show how the quadratic model fits the temporal evolution in the case of pastry and bakery category, in 18 countries. We also show linear (Figure B.6) and constant fit (Figure B.7) for comparison.

B.1.2 The impact of bandwidth

In Figure B.8, we study the impact of the choice of the bandwidth $\max(t_{\min}, t_{\max}) = 30$ and the choice of the degree of the model. We observe that for a sufficiently large bandwidth, all four models estimate a similar effect, and the choice of bandwidth does not matter as the estimates converge.

B.1.3 Modeling interest share

We show our main results with the same model, but the dependent variable being the weekly share of interest in Figure B.9. This way, we control for overall increased interest in all categories. This analysis provides an alternative view. We see that the share of volume decreases significantly for foods whose growth is not proportional to the growth of the foods that experience major surges of interest.

B.2 Supplementary analyses

We perform supplementary analyses that support our main conclusions or provide complementary insights. In Figure B.10 we show short-term effects estimated with quadratic model, grouped by country. In each country, the gray line represents the overall country-specific short-term effect, that is the increase in interest in all food entities. Finally, we explore the effect of the second mobility decrease in Figure B.11, and we present the long-term effects in Figure B.12.

Table B.1: Summary of the 28 food entity categories. For each category, we present the category description, the number of entities in the category (Number), and category size (Size) that is the fraction of search interest covered by the category, on average, in the 18 studied countries in 2019 and 2020. Additionally, for each category, we show top 10 individual entities by the rank of the volume in average across 18 studied countries, in 2019 and 2020.

Category	Description	Number	Size
beef dish	food preparation based on beef	51	3.6%
Top 10 entities:	Hamburger, Beef, Steak, Meatball, Meatloaf, Beef Stroganoff, Fajita, Beef mince, Sirloin steak, Big Mac		
chicken dish	food preparation based on chicken	37	3.3%
Top 10 entities:	Chicken meat, Chicken nugget, Fried chicken, Chicken curry, Chicken soup, Hendl, Butter chicken, Chicken tikka masala, Barbecue chicken, Cordon bleu		
pork dish	food preparation based on pork	45	2.1%
Top 10 entities:	Pork, Ham, Bacon, Hot dog, Pork chop, Gyro, Pork tenderloin, Pork belly, Pulled pork, Schnitzel		
lamb dish	food preparation based on lamb	17	0.5%
Top 10 entities:	Lamb and mutton, Shawarma, Doner kebab, Mechoui, Sfiha, Rogan josh, Cig kofte, Kokoretsi, Pasanda, Arrosticini		
fish dish	type of dish comprised of fish	57	1.8%
Top 10 entities:	Tuna, Caviar, Salmon, Cod, Squid, Sardine, Catfish, Crayfish, Tempura, Smoked salmon		
sausage	food usually made from ground meat with a skin around it	16	0.6%
Top 10 entities:	Sausage, Salami, Chorizo, Bratwurst, Mortadella, Black pudding, 'Nduja, Sujuk, Boudin, Andouille		
pasta, pizza and noodle dish	Italian food made from flour and water and shaped in different forms, usually cooked and served with a sauce, or a dish made with pasta, or other type of staple food made from some type of unleavened dough	95	6.9%
Top 10 entities:	Pizza, Pasta, Spaghetti, Lasagne, Noodle, Carbonara, Gnocchi, Macaroni, Penne, Ravioli		
potato dish	type of food based on potatoes	27	1.0%
Top 10 entities:	French fries, Mashed potato, Gratin, Baked potato, Tortilla de patatas, Potato, Potato pancake, Sunday roast, Tater Tots, Patatas bravas		
rice dish	a type of dish made of rice	49	3.8%
Top 10 entities:	Rice, Sushi, Risotto, Fried rice, Basmati, Paella, Biryani, Bento, Pilaf, White rice		
egg dish	a type of dish made of eggs	22	2.6%
Top 10 entities:	Egg, Boiled egg, Omelette, Quiche, Scrambled eggs, Poached egg, Frittata, Eggs Benedict, Deviled egg, Egg roll		
stew	combination of solid food ingredients that have been cooked in liquid and served in the resultant gravy	24	0.4%
Top 10 entities:	Stew, Ratatouille, Jambalaya, Dolma, Gumbo, Sambar, Cassoulet, Blanquette de veau, Irish stew, Bigos		
soup	primarily liquid food	55	2.7%
Top 10 entities:	Soup, Broth, Ramen, Miso, Pho, Hot pot, French onion soup, Goulash, Cream of mushroom soup, Minestrone		
bread and flatbread	staple food prepared from a dough	31	2.7%
Top 10 entities:	Bread, Pita, Bagel, Baguette, Sourdough, Naan, Pretzel, Focaccia, Bruschetta, White bread		
sandwich	two slices of bread with filling in between them	20	0.5%
Top 10 entities:	Sandwich, Panini, Corn dog, Croque-monsieur, Banh mi, BLT, Tuna fish sandwich, Peanut butter and jelly sandwich, Filet-O-Fish, Bocadillo		
salad	dish consisting of a mixture of small pieces of food, usually vegetables or fruit	24	1.7%
Top 10 entities:	Salad, Lettuce, Potato salad, Caesar salad, Pasta salad, Tabbouleh, Greek salad, Romaine lettuce, Insalata Caprese, Olivier salad		
cheese	yellow or white, creamy or solid food made from the pressed curds of milk	90	2.9%
Top 10 entities:	Cheese, Mozzarella, Cream cheese, Parmigiano-Reggiano, Ricotta, Feta, Cheddar cheese, Fondue, Mascarpone, Cottage cheese		
sauce	liquid, creaming or semi-solid food served on or used in preparing other foods	60	3.5%
Top 10 entities:	Sauces, Mayonnaise, Pesto, Dip, Mustard, Tomato sauce, Bechamel sauce, Bolognese sauce, Soy sauce, Gravy		
snack	portion of food, often smaller than a regular meal	20	1.8%
Top 10 entities:	Peanut, Popcorn, Hummus, Cashew, Tapas, Pistachio, Guacamole, Nachos, Cracker, Edamame		
vegetable and legume	edible plant or part of a plant, involved in cooking	85	9.5%
Top 10 entities:	Vegetable, Tomato, Sweet potato, Onion, Cucumber, Spinach, Eggplant, Cauliflower, Cabbage, Asparagus		

Appendix B. COVID-19-induced shifts in dietary interests

fruit	food, edible in the raw state	63	9.9%
Top 10 entities:	Apple, Lemon, Pineapple, Avocado, Grape, Mango, Watermelon, Cherry, Strawberry, Banana		
herb	plant part used for flavoring, food, medicine, or perfume	29	2.2%
Top 10 entities:	Lavender, Basil, Herb, Rosemary, Celery, Coriander, Parsley, Eucalyptus, Peppermint, Liquorice		
spice	dried seed, fruit, root, bark, or vegetable substance primarily used for flavoring, coloring or preserving food	38	5.1%
Top 10 entities:	Garlic, Table salt, Chili pepper, Ginger, Spice, Turmeric, Vanilla, Cinnamon, Common Fig, Black pepper		
soft drink	non-alcoholic drink, often carbonated (sparkling)	27	2.1%
Top 10 entities:	Coca-Cola, Juice, Soft drink, Cola, Lemonade, Orange juice, Tonic water, Energy drink, Iced tea, Apple juice		
wine, beer and liquor	alcoholic drink, alcoholic drink typically made from grapes, or alcoholic beverage that is produced by distilling	46	6.6%
Top 10 entities:	Wine, Vodka, Beer, Alcoholic beverage, Rum, Gin, Champagne, Tequila, Red Wine, Sake		
cocktail	alcoholic mixed drink	142	1.3%
Top 10 entities:	Cocktail, Mojito, Martini, Sour, Margarita, Gin and tonic, Piña colada, Mimosa, Spritz, Bloody Mary		
pie	baked dish usually made of a pastry dough casing, containing a filling of various sweet or savoury ingredients	20	1.3%
Top 10 entities:	Pie, Tart, Apple pie, Cottage pie, Pumpkin pie, Borek, Tarte Tatin, Meat pie, Banoffee pie, Lemon meringue pie		
pastry and bakery product	various baked products made of dough	40	1.5%
Top 10 entities:	Baking powder, Pastry, Baker's yeast, Puff pastry, Brioche, Samosa, Filo, Ice cream cone, Choux pastry, Cannoli		
dessert	course that concludes a meal; usually very sweet	202	18.2%
Top 10 entities:	Cake, Chocolate, Ice cream, Honey, Pancake, Biscuit, Cookie, Doughnut, Cupcake, Chocolate brownie		

Table B.3: Entity-level Spearman rank correlation between interest and mobility. For modes all entities are shown. For foods, top 10 entities most and least correlated across countries on average (* marks $p < 0.05$ between seasonality adjusted interest and mobility according to two-sided test with no correlation null hypothesis, - marks too low search interest in a country). All entities related to consuming food at home are correlated positively on average, all entities related to consuming food outside of home are correlated negatively on average, except barbecue.

Mode entity (N=16)	Average	AU	BR	CA	DE	DK	EG	ES	FR	GB	ID	IN	IT	JP	KE	MX	NG	SE	US
On average positively correlated																			
Recipe	0.70	0.88*	0.87*	0.73*	0.74*	0.73*	0.45*	0.81*	0.72*	0.85*	0.76*	0.82*	0.85*	0.04	0.66*	0.95*	0.47*	0.61*	0.68*
Baking	0.62	0.88*	0.8*	0.71*	0.55*	0.28	0.41*	0.64*	0.43*	0.88*	0.71*	0.91*	0.46*	0.17	0.77*	0.79*	0.47*	0.55*	0.77*
Cooking	0.58	0.75*	0.67*	0.6*	0.39*	0.23	0.08	0.87*	0.61*	0.88*	0.63*	0.84*	0.75*	0.26	0.68*	0.83*	0.37*	0.34*	0.62*
Take-out	0.50	0.8*	0.66*	0.64*	0.56*	0.73*	0.07	0.1	0.62*	0.33*	0.42*	0.31*	0.79*	0.08	0.49*	0.03	0.63*	0.88*	0.88*
Grocery store	0.28	0.14	0.32*	0.64*	0.41*	0.32*	0.07	0.12	-0.29*	0.64*	0.38*	0.8*	0.46*	0.15	0.36*	-0.28	0.35*	0.16	0.39*
Supermarket	0.17	-0.22	0.58*	0.38*	-0.02	0.25	0.24	0.47*	0.4*	0.03	0.12	0.02	0.63*	0.04	0.23	0.16	0.23	0.29	0.29
Barbecue	0.11	0.04	0.61*	0.23	-0.01	-0.2	0.15	0.11	-0.01	0.4*	0.07	-0.81*	-0.05	0.39*	0.51*	0.31*	0.15	0.25	-0.2
Food delivery	0.09	0.51*	-0.12	0.4*	0.07	-0.25	0.04	-0.08	-0.05	0.27	0.1	-0.12	0.36*	0.27	-0.02	-0.33*	0.04	-0.08	0.52*
Drive-in	0.0003	-0.44*	0.25	-0.14	0.42*	0.51*	0.05	-0.24	0.47*	0.09	-0.23	-0.27	-0.32*	-0.04	-0.17	0.01	0.19	0.09	0.09
On average negatively correlated																			
Picnic	-0.18	-0.1	-0.62*	-0.18	0.48*	0.08	-0.13	-0.46*	-0.28	0.09	0.06	-0.61*	-0.65*	-0.47*	-0.06	-0.11	-0.18	0.04	-0.11
Lunchbox	-0.21	-0.45*	0.12	-0.09	-0.49*	-0.2	0.08	-0.08	-0.14	-0.57*	0.15	-0.56*	-0.33*	-0.15	-0.18	-	-0.23	-0.04	-0.33*
Diner	-0.24	-0.34*	-0.03	-0.6*	-0.65*	-0.13	0.1	0.14	-0.21	-0.78*	-0.46*	-0.34*	-0.33*	0.0	-0.03	0.08	0.32*	-0.72*	-0.72*
Food festival	-0.26	-0.34*	-0.59*	-0.24	-0.63*	-0.02	0.06	-0.24	-0.24	-0.56*	-0.21	-0.32*	-0.3*	0.1	-0.04	-0.04	-0.11	-0.43*	-0.48*
Cafeteria	-0.28	-0.34*	-0.52*	-0.33*	-0.69*	-0.28	-0.2	-0.15	-0.83*	-0.04	-0.2	-0.26	-0.03	-0.24	-0.14	0.36*	-0.39*	-0.17	-0.59*
Cafe	-0.72	-0.69*	-0.74*	-0.92*	-0.86*	-0.67*	-0.67*	-0.83*	-0.71*	-0.87*	-0.74*	-0.9*	-0.75*	-0.78*	-0.3*	-0.76*	-0.32*	-0.71*	-0.74*
Restaurant	-0.80	-0.72*	-0.83*	-0.94*	-0.86*	-0.76*	-0.85*	-0.87*	-0.81*	-0.87*	-0.84*	-0.75*	-0.97*	-0.83*	-0.4*	-0.85*	-0.66*	-0.77*	-0.76*
Food entity (N=1432)	Average	AU	BR	CA	DE	DK	EG	ES	FR	GB	ID	IN	IT	JP	KE	MX	NG	SE	US
Top 10, on average positively correlated																			
Pancake (dessert)	0.61	0.73*	0.87*	0.62*	0.64*	0.53*	0.47*	0.67*	0.61*	0.56*	0.54*	0.74*	0.8*	0.12	0.59*	0.72*	0.58*	0.53*	0.62*
Baking powder (pastry and bakery product)	0.58	0.61*	0.73*	0.76*	0.6*	0.16	0.33*	0.57*	0.61*	0.82*	0.76*	0.91*	0.81*	0.61*	0.3*	0.7*	0.29*	0.31*	0.63*
Bread (bread and flatbread)	0.58	0.81*	0.74*	0.69*	0.5	0.37*	0.58*	0.54*	0.58*	0.77*	0.63*	0.92*	0.82*	-0.14	0.57*	0.56*	0.56*	0.29	0.71*
Baker's yeast (pastry and bakery product)	0.58	0.79*	0.84*	0.71*	0.56*	0.33*	0.24	0.68*	0.7*	0.78*	0.56*	0.92*	0.54*	0.35*	0.5*	-	0.54*	0.05	0.78*
Cookie (dessert)	0.57	0.67*	0.83*	0.68*	0.51*	0.53*	0.31*	0.66*	0.73*	0.88*	0.34*	0.82*	0.53*	0.3*	0.33*	-	0.67*	0.42*	0.64*
Chocolate brownie (dessert)	0.56	0.65*	0.89*	0.63*	0.44*	0.23	0.53*	0.43*	0.73*	0.84*	0.57*	0.7*	0.6*	0.49*	0.32*	0.78*	0.3*	0.35*	0.65*
Chicken meat (chicken dish)	0.55	0.65*	0.83*	0.63*	0.47*	0.26	0.32*	0.76*	0.61*	0.75*	0.65*	0.73*	0.41*	0.31*	0.57*	0.83*	0.4*	-0.02	0.75*
Chocolate cake (dessert)	0.55	0.65*	0.8*	0.5*	0.41*	0.62*	0.28	0.62*	0.63*	0.8*	0.4*	0.86*	0.69*	0.75*	0.62*	-	0.06	0.04	0.55*
Biscuit (dessert)	0.54	0.7*	0.7*	0.75*	0.39*	0.05	0.46*	0.73*	0.7*	0.82*	0.56*	0.75*	0.68*	0.28	0.25	0.77*	0.33*	-0.0	0.73*
Pasta (pasta, pizza and noodle dish)	0.53	0.72*	0.51*	0.6*	0.4*	0.19	0.34*	0.68*	0.61*	0.74*	0.63*	0.77*	0.83*	0.25	0.36*	0.79*	0.05	0.45*	0.69*
Top 10, on average negatively correlated																			
Tapas (snack)	-0.34	-0.51*	-0.05	-0.49*	-0.76*	-0.45*	0.02	-0.89*	-0.83*	-0.77*	-0.38*	0.09	-0.18	-0.11	-0.24	0.41*	0.24	-0.46*	-0.7*
Energy drink (soft drink)	-0.27	0.02	-0.71*	-0.32*	-0.22	-0.2	-0.14	-0.42*	-0.13	-0.46*	-0.13	-0.48*	-0.24	-0.08	-0.22	-	-0.08	-0.29*	-0.52*
Korean barbecue (beef dish)	-0.26	-0.62*	0.02	-0.69*	-0.5*	-0.17	0.05	-0.17	-0.47*	-0.6*	-0.48*	0.09	-0.06	-0.02	-	0.03	-	0.12	-0.73*
Campanelle (pasta, pizza and noodle dish)	-0.22	-0.19	-0.22	-0.1	-	-	-	-0.09	-0.39*	-0.48*	-0.18	-	-0.14	-	-0.22	-	-0.22	-0.23	-0.23
Sushi (rice dish)	-0.19	-0.33*	-0.1	-0.36*	-0.05	-0.15	-0.02	0.02	-0.07	-0.12	-0.57*	-0.16	-0.43*	-0.71*	-0.04	0.15	0.22	0.05	-0.42*
Gelato (dessert)	-0.19	-0.21	-0.25	-0.24	0.05	-0.24	-0.23	0.02	0.04	-0.11	-0.57*	-0.29	-0.29*	-0.48*	-0.32*	-0.14	0.23	-0.07	-0.31*
Chewing gum (dessert)	-0.17	-0.14	0.03	-0.09	-0.39*	-0.19	-0.06	-0.25	-0.16	-0.48*	0.02	-0.21	-0.37*	0.02	-0.02	-	-0.05	-0.07	-0.54*
Ramen (soup)	-0.17	-0.31*	-0.22	-0.48*	-0.46*	-0.04	0.26	-0.44*	-0.14	0.07	0.02	-0.18	-0.18	-0.65*	0.15	-	0.02	-0.28	-0.1
Aglio e olio (pasta, pizza and noodle dish)	-0.17	-0.19	-0.21	-0.32*	-0.25	0.03	-	-0.28	-0.15	-0.58*	-0.26	0.14	0.2	-	-	-	-	0.18	-0.44*
Burrata (cheese)	-0.16	-0.02	0.18	-0.36*	-0.15	0.25	-0.12	-0.48*	-0.28	-0.27	-0.29*	0.06	-0.33*	-0.02	-	-0.12	-0.0	-0.2	-0.63*

Table B.4: Main quadratic model: fitted coefficient alpha and R^2 statistic. Entities Recipe, Food delivery, Restaurant, and Picnic mark the sets of entities described in Figure 6.5a.

Access mode	AU	BR	CA	DE	DK	EG	ES	FR	GB	ID	IN	IT	JP	KE	MX	NG	SE	US
Recipe	0.34(0.9)	0.46(0.9)	0.57(0.87)	0.43(0.77)	0.29(0.77)	0.28(0.74)	0.84(0.9)	0.73(0.81)	0.78(0.94)	0.64(0.69)	1.03(0.95)	0.77(0.87)	0.08(0.93)	0.73(0.88)	0.58(0.93)	0.41(0.79)	0.09(0.78)	0.61(0.91)
Food delivery	0.78(0.9)	0.77(0.93)	0.97(0.95)	0.15(0.85)	1.05(0.9)	-0.46(0.15)	-0.09(0.66)	0.96(0.57)	0.37(0.86)	1.40(0.34)	0.37(0.24)	0.95(0.7)	1.38(0.95)	-0.34(0.18)	0.4(0.65)	0.02(0.18)	-0.35(0.74)	1.11(0.96)
Restaurant	-0.78(0.86)	-0.23(0.88)	-0.68(0.87)	-0.98(0.84)	-0.77(0.75)	-0.75(0.89)	-1.65(0.89)	-1.70(92)	-0.95(0.89)	-0.71(0.89)	-1.45(0.97)	-1.45(0.98)	-0.62(0.92)	-0.54(0.93)	-0.99(0.67)	-0.13(0.76)	-0.39(0.82)	-0.13(0.82)
Picnic	0.11(0.7)	0.16(0.85)	0.24(0.97)	0.44(0.7)	0.19(0.55)	-0.19(0.36)	-0.25(0.77)	-0.58(0.77)	0.92(0.87)	0.26(0.59)	-0.36(0.92)	-0.6(0.72)	-0.21(0.78)	0.92(0.35)	0.11(0.49)	0.97(0.35)	0.51(0.7)	-0.06(0.84)
Food categories	AU	BR	CA	DE	DK	EG	ES	FR	GB	ID	IN	IT	JP	KE	MX	NG	SE	US
beef dish	0.16(0.79)	0.35(0.82)	0.24(0.63)	0.28(0.4)	0.16(0.39)	-0.0(0.28)	0.59(0.7)	0.38(0.68)	0.55(0.77)	0.20(0.3)	0.23(0.64)	0.17(0.39)	-0.28(0.92)	0.21(0.48)	0.2(0.66)	0.32(0.37)	0.12(0.63)	0.41(0.66)
bread and flatbread	0.39(0.91)	0.6(0.96)	1.08(0.9)	0.48(0.93)	0.69(0.88)	0.21(0.39)	1.41(0.91)	1.36(0.9)	0.67(0.96)	0.66(0.84)	0.64(0.8)	1.37(0.86)	-0.16(0.91)	0.42(0.71)	0.13(0.92)	0.61(0.76)	-0.04(0.85)	0.84(0.92)
cheese	0.37(0.74)	0.45(0.85)	0.35(0.68)	0.18(0.8)	0.26(0.52)	0.56(0.58)	0.64(0.65)	0.44(0.6)	0.62(0.76)	0.62(0.73)	0.6(0.89)	0.41(0.64)	0.04(0.87)	1.0(0.51)	0.62(0.66)	-0.33(0.45)	0.12(0.57)	0.27(0.66)
chicken dish	0.25(0.84)	0.28(0.91)	0.3(0.81)	0.31(0.77)	0.19(0.65)	0.74(0.31)	0.77(0.87)	0.52(0.67)	0.57(0.82)	0.59(0.43)	0.38(0.79)	0.5(0.81)	0.04(0.92)	0.76(0.6)	0.34(0.91)	0.5(0.61)	0.11(0.72)	0.29(0.83)
cocktail	0.25(0.73)	0.35(0.73)	0.29(0.8)	0.06(0.84)	-0.18(0.79)	-0.52(0.45)	0.09(0.74)	0.26(0.77)	0.7(0.72)	0.33(0.53)	-0.23(0.4)	0.5(0.79)	0.13(0.56)	-0.56(0.39)	-0.1(0.51)	-1.19(0.59)	0.02(0.66)	0.31(0.66)
dessert	0.44(0.79)	0.41(0.89)	0.56(0.83)	0.34(0.71)	0.32(0.69)	0.36(0.63)	0.99(0.9)	0.74(0.81)	1.01(0.85)	0.42(0.77)	0.51(0.76)	0.66(0.85)	0.06(0.63)	0.46(0.87)	0.27(0.88)	0.14(0.91)	-0.02(0.66)	0.37(0.72)
egg dish	0.82(0.74)	0.24(0.61)	0.23(0.58)	0.29(0.66)	0.14(0.47)	-0.02(0.56)	0.73(0.79)	0.58(0.81)	0.55(0.81)	0.68(0.79)	0.55(0.9)	0.33(0.73)	0.05(0.89)	0.27(0.66)	0.34(0.82)	0.55(0.75)	-0.37(0.55)	0.21(0.56)
fish dish	0.12(0.44)	0.36(0.38)	0.14(0.66)	0.24(0.5)	0.33(0.22)	0.52(0.15)	0.65(0.55)	0.27(0.57)	0.29(0.79)	0.19(0.69)	0.08(0.68)	0.11(0.31)	0.03(0.94)	0.5(0.25)	0.38(0.61)	0.35(0.46)	0.35(0.33)	0.16(0.85)
fruit	0.36(0.87)	0.21(0.8)	0.33(0.9)	0.2(0.79)	0.05(0.61)	0.21(0.77)	0.6(0.83)	0.64(0.77)	0.65(0.94)	0.21(0.63)	0.28(0.83)	0.33(0.78)	-0.11(0.9)	0.55(0.91)	0.35(0.93)	0.27(0.79)	0.09(0.78)	0.19(0.89)
herb	0.03(0.78)	0.05(0.95)	0.21(0.91)	0.16(0.93)	0.45(0.66)	0.39(0.48)	0.16(0.81)	0.19(0.87)	0.58(0.91)	-0.03(0.54)	0.17(0.92)	0.27(0.76)	0.05(0.89)	0.57(0.57)	-0.18(0.71)	0.25(0.75)	0.31(0.84)	0.14(0.92)
lamb dish	0.18(0.82)	0.61(0.71)	-0.04(0.35)	0.05(0.75)	0.13(0.22)	-0.4(0.54)	-0.31(0.08)	-0.32(0.4)	0.31(0.59)	0.09(0.21)	-0.14(0.47)	-0.37(0.5)	-0.32(0.9)	2.51(0.28)	-0.11(0.18)	0.29(0.65)	0.17(0.44)	0.11(0.34)
pasta, pizza & noodle dish	0.22(0.8)	0.43(0.78)	0.28(0.82)	0.17(0.63)	0.25(0.4)	0.43(0.76)	0.66(0.77)	0.34(0.65)	0.37(0.91)	0.58(0.88)	0.5(0.86)	0.61(0.72)	-0.17(0.91)	0.52(0.76)	0.3(0.77)	0.44(0.69)	0.16(0.61)	0.27(0.68)
pastry & bakery product	0.52(0.86)	0.91(0.93)	1.26(0.81)	0.96(0.87)	0.43(0.55)	1.14(0.51)	1.45(0.96)	1.32(0.82)	1.1(0.93)	0.34(0.49)	1.05(0.74)	1.08(0.84)	0.72(0.63)	1.08(0.6)	-0.0(0.32)	0.66(0.75)	0.2(0.44)	1.06(0.83)
pie	0.37(0.88)	0.49(0.87)	0.8(0.79)	0.3(0.45)	0.58(0.5)	0.6(0.36)	1.14(0.79)	0.74(0.77)	0.75(0.83)	0.04(0.54)	0.31(0.44)	0.84(0.84)	0.13(0.75)	0.59(0.4)	0.65(0.63)	0.57(0.64)	0.13(0.86)	0.66(0.75)
pork dish	0.19(0.85)	0.32(0.88)	0.35(0.59)	0.23(0.43)	0.31(0.28)	0.7(0.15)	0.42(0.69)	0.34(0.76)	0.56(0.85)	0.58(0.52)	0.19(0.83)	0.4(0.53)	0.04(0.92)	1.08(0.45)	0.53(0.35)	0.08(0.63)	0.34(0.61)	0.34(0.61)
potato dish	0.33(0.72)	0.46(0.87)	0.53(0.58)	0.47(0.77)	0.01(0.59)	0.88(0.18)	0.61(0.72)	0.59(0.82)	0.66(0.78)	0.57(0.67)	0.63(0.93)	0.47(0.83)	0.05(0.82)	0.34(0.48)	1.11(0.67)	2.26(0.29)	0.04(0.7)	0.44(0.73)
rice dish	0.17(0.75)	0.4(0.74)	0.16(0.66)	0.1(0.73)	0.24(0.5)	0.34(0.43)	0.63(0.86)	0.37(0.75)	0.28(0.87)	0.42(0.6)	0.61(0.89)	0.21(0.54)	-0.06(0.92)	0.49(0.56)	0.32(0.85)	0.4(0.68)	0.15(0.67)	0.23(0.85)
salad	0.18(0.84)	0.11(0.61)	0.04(0.9)	0.16(0.81)	-0.08(0.78)	0.32(0.37)	0.54(0.89)	0.41(0.85)	0.6(0.81)	0.44(0.48)	0.35(0.39)	0.24(0.8)	-0.02(0.93)	0.14(0.23)	0.24(0.64)	-0.67(0.29)	0.28(0.8)	0.02(0.85)
sandwich	0.25(0.48)	0.04(0.4)	-0.13(0.5)	0.17(0.31)	-0.09(0.28)	0.09(0.38)	0.26(0.25)	0.25(0.46)	0.14(0.49)	0.47(0.39)	0.33(0.69)	0.34(0.41)	0.09(0.91)	-5.08(0.28)	0.18(0.5)	0.53(0.13)	0.62(0.31)	-0.18(0.36)
sauce	0.38(0.84)	0.43(0.86)	0.31(0.8)	0.28(0.72)	0.36(0.6)	0.56(0.56)	0.71(0.81)	0.49(0.69)	0.6(0.86)	0.62(0.69)	0.71(0.9)	0.52(0.57)	0.19(0.93)	0.91(0.62)	0.46(0.84)	-0.42(0.59)	0.09(0.76)	0.33(0.85)
sausage	0.33(0.72)	0.29(0.77)	0.18(0.54)	-0.01(0.47)	0.15(0.22)	0.16(0.12)	0.77(0.73)	0.36(0.75)	0.55(0.8)	0.44(0.56)	0.81(0.56)	0.38(0.44)	0.16(0.77)	2.41(0.22)	0.27(0.75)	0.35(0.45)	-0.32(0.33)	0.34(0.77)
snack	0.28(0.56)	0.3(0.75)	0.26(0.78)	0.02(0.64)	-0.11(0.34)	-0.04(0.56)	0.13(0.45)	0.37(0.67)	0.57(0.72)	0.01(0.38)	0.43(0.8)	0.18(0.58)	0.32(0.86)	-0.29(0.49)	0.08(0.39)	0.06(0.46)	0.28(0.24)	0.32(0.72)
soft drink	0.27(0.54)	-0.05(0.59)	-0.02(0.48)	0.09(0.82)	-0.04(0.46)	0.22(0.54)	0.47(0.61)	0.04(0.57)	0.32(0.72)	0.18(0.23)	-0.13(0.64)	0.53(0.41)	0.09(0.86)	-0.06(0.3)	0.15(0.49)	0.36(0.51)	-0.02(0.45)	0.02(0.56)
soup	-0.11(0.92)	0.5(0.68)	0.27(0.91)	0.12(0.87)	0.22(0.85)	0.19(0.2)	0.5(0.74)	0.34(0.88)	0.29(0.94)	0.46(0.38)	0.34(0.63)	0.23(0.83)	-0.33(0.91)	-0.31(0.3)	0.31(0.56)	0.44(0.67)	-0.0(0.8)	0.25(0.96)
spice	0.28(0.84)	0.14(0.94)	0.27(0.85)	0.09(0.78)	0.03(0.57)	0.17(0.5)	0.49(0.78)	0.45(0.77)	0.49(0.85)	0.2(0.84)	0.42(0.95)	0.09(0.34)	0.07(0.9)	0.26(0.88)	0.16(0.89)	0.34(0.85)	0.03(0.52)	0.29(0.82)
stew	0.23(0.86)	0.55(0.58)	0.48(0.89)	0.41(0.52)	0.59(0.54)	-1.67(0.08)	0.57(0.78)	0.08(0.84)	0.43(0.87)	0.37(0.75)	0.55(0.83)	0.39(0.3)	0.18(0.8)	1.71(0.29)	0.49(0.25)	0.26(0.47)	0.77(0.28)	0.61(0.94)
vegetable and legume	0.29(0.9)	0.21(0.94)	0.28(0.86)	0.22(0.71)	0.34(0.46)	-0.02(0.62)	0.51(0.86)	0.69(0.79)	0.57(0.89)	0.4(0.74)	0.53(0.92)	0.27(0.73)	0.11(0.93)	0.57(0.85)	0.4(0.81)	0.26(0.6)	0.12(0.66)	0.3(0.84)
wine, beer and liquor	0.13(0.68)	0.19(0.85)	0.1(0.59)	0.02(0.82)	-0.04(0.68)	0.19(0.82)	0.06(0.64)	-0.29(0.78)	0.33(0.78)	-0.09(0.52)	0.18(0.55)	-0.16(0.53)	0.04(0.86)	0.87(0.56)	0.33(0.52)	0.3(0.27)	-0.13(0.7)	0.13(0.45)





Figure B.1: Mobility in 18 studied countries. Mobility changepoints (mobility decrease, mobility increase, and the second mobility decrease in case it occurs) are marked with vertical dashed lines.


Table B.5: The impact of model order. Food categories, ranked by short-term effect sizes in decreasing order, estimated with a constant, linear, and quadratic model.

Rank	Constant model	Linear model	Quadratic model
1	pastry and bakery product	pastry and bakery product	pastry and bakery product
2	pie	pie	bread and flatbread
3	dessert	bread and flatbread	potato dish
4	sauce	potato dish	pie
5	potato dish	dessert	dessert
6	bread and flatbread	sauce	cheese
7	chicken dish	chicken dish	sauce
8	stew	cheese	chicken dish
9	egg dish	vegetable and legume	pork dish
10	vegetable and legume	egg dish	sausage
11	cheese	pork dish	stew
12	fruit	pasta, pizza and noodle dish	pasta, pizza and noodle dish
13	herb	fruit	egg dish
14	spice	stew	vegetable and legume
15	sausage	rice dish	fruit
16	rice dish	spice	rice dish
17	pasta, pizza and noodle dish	herb	fish dish
18	fish dish	sausage	beef dish
19	pork dish	fish dish	spice
20	salad	snack	herb
21	snack	beef dish	soup
22	beef dish	salad	snack
23	sandwich	sandwich	salad
24	soft drink	soup	lamb dish
25	soup	soft drink	soft drink
26	lamb dish	wine, beer and liquor	wine, beer and liquor
27	wine, beer and liquor	cocktail	cocktail
28	cocktail	lamb dish	sandwich

Appendix B. COVID-19-induced shifts in dietary interests



(a)



(b)

Figure B.2: Pearson's correlation coefficient between mobility and interest volume. In (a), correlation for categories of food entities, and in (b), for ways of accessing food. For each group, $n = 18$ values represent correlation coefficient (calculated based on $n = 46$ samples corresponding to weeks of 2020). The boxplot summarizes the value across 18 countries. Significant correlations ($p < 0.05$), according to a two-sided hypothesis test whose null hypothesis is that interest and mobility are uncorrelated, are marked in blue, and not significant in orange. No adjustments for multiple comparisons are made. Boxplots represent the 50th (center line), 25th and 75th percentile (box limits). The whiskers extend to the minimum and maximum values but no further than 1.5 times IQR.

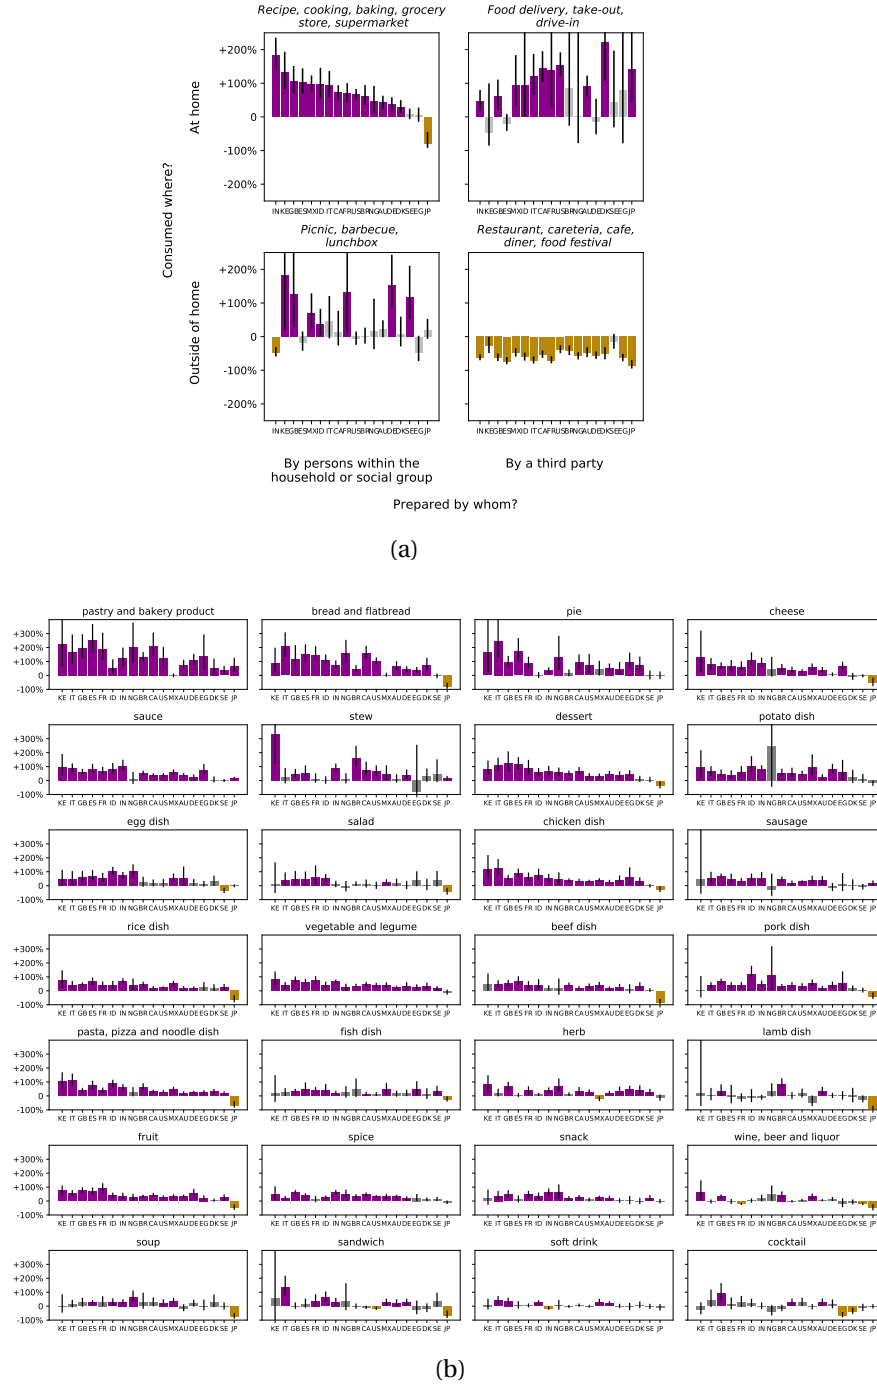
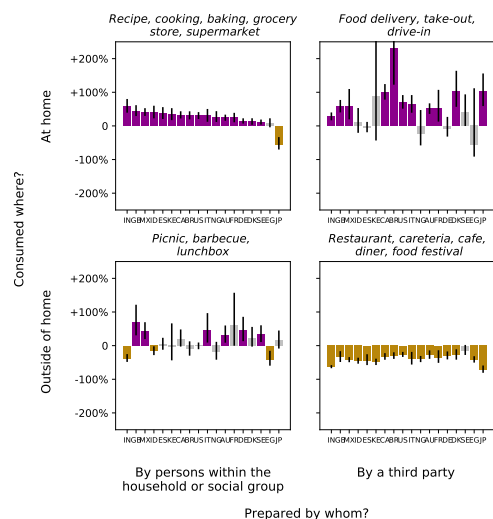


Figure B.3: Short-term effects estimated with a linear model. For each country ($n = 18$), for each food access mode (in (a), $n = 4$) and food category (in (b), $n = 28$), model (Equation 6.1) is fitted on $n = 82$ samples. Bars represent effect estimates (coefficient α estimated with our RDD-based model). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects.

Appendix B. COVID-19-induced shifts in dietary interests



(a)



(b)

Figure B.4: Short-term effects estimated with a constant model. For each country ($n = 18$), for each food access mode (in (a), $n = 4$) and food category (in (b), $n = 28$), model (Equation 6.1) is fitted on $n = 82$ samples. Bars represent effect estimates (coefficient α estimated with our RDD-based model). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects.

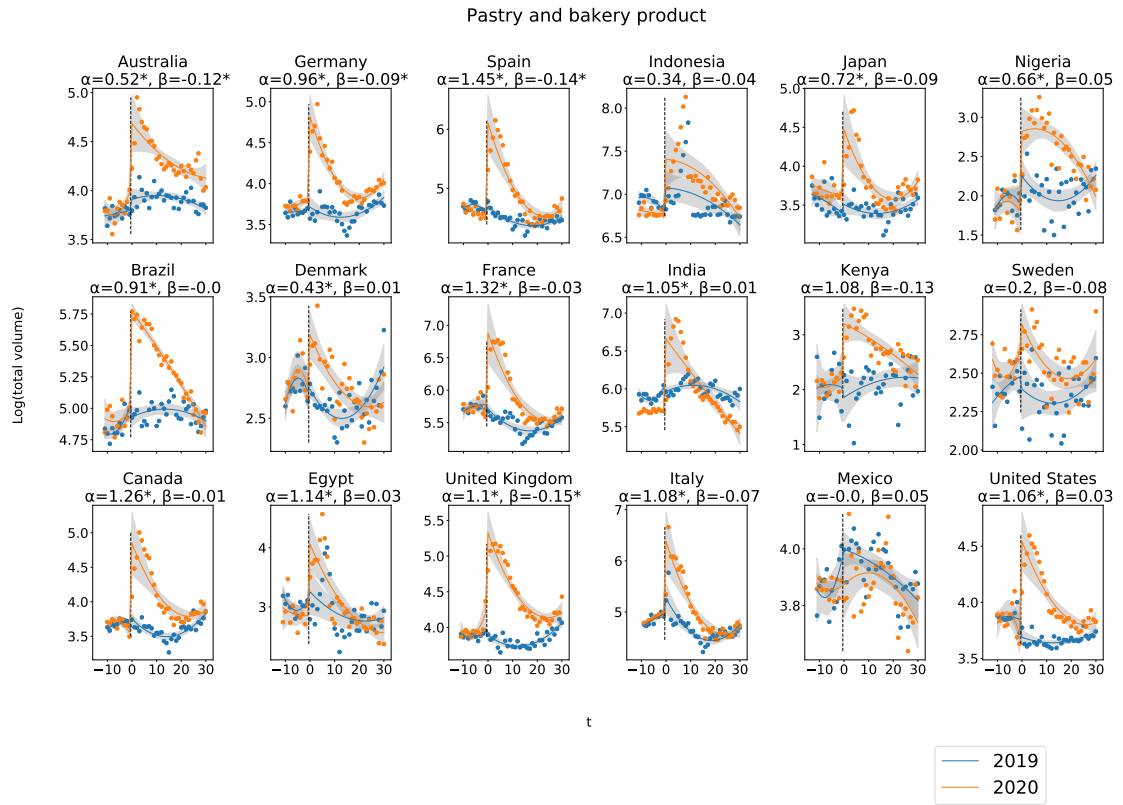


Figure B.5: Example of the quadratic model fit. On x-axis weeks, relative to the week of mobility decrease, on y-axis the interest volume. Error bars mark 95% confidence intervals of the model fit. α and β are fitted coefficients. Note the varying y-scales.

Appendix B. COVID-19-induced shifts in dietary interests

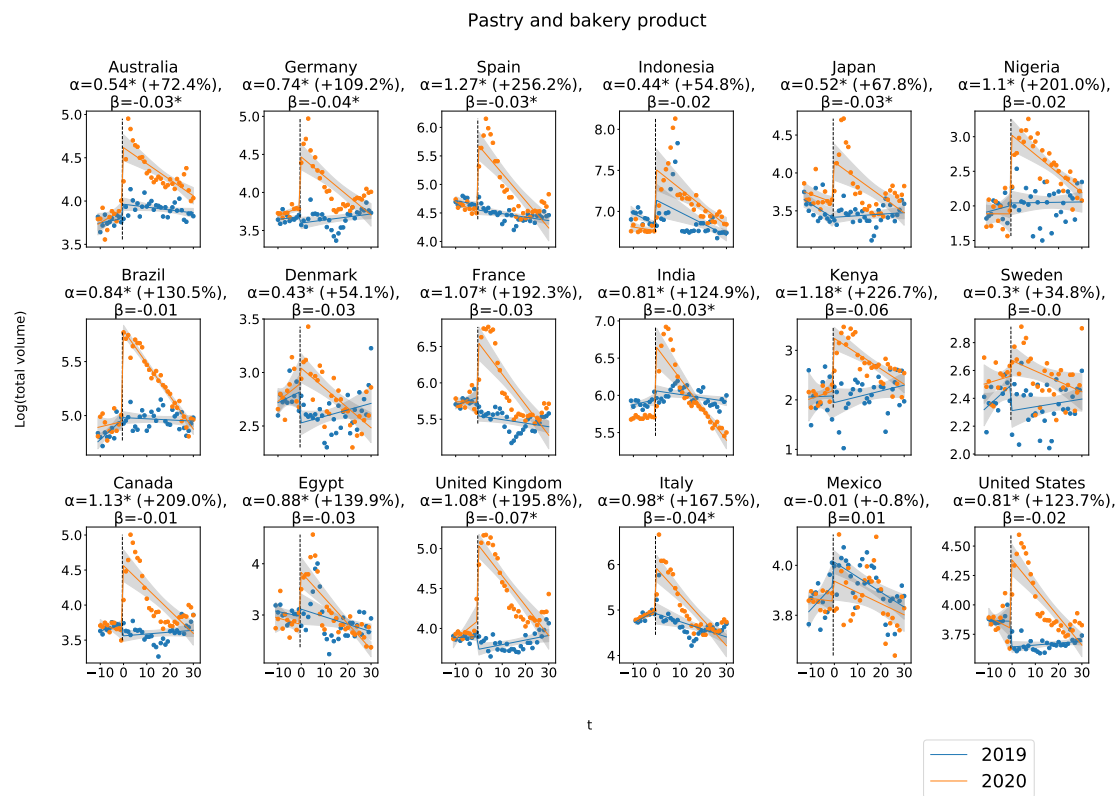


Figure B.6: Example of the linear model fit. On x-axis weeks, relative to the week of mobility decrease, on y-axis the interest volume. Error bars mark 95% confidence intervals of the model fit. α and β are fitted coefficients. Note the varying y-scales.

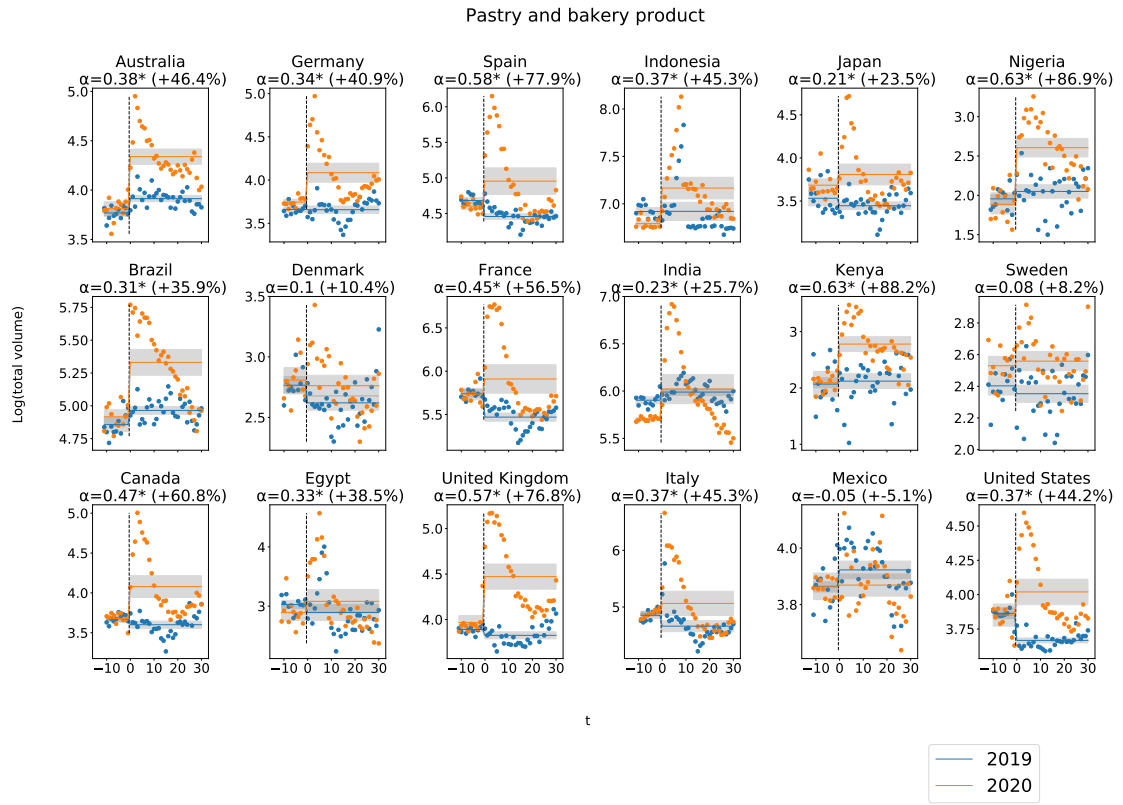


Figure B.7: Example of the constant model fit. On x-axis weeks, relative to the week of mobility decrease, on y-axis the interest volume. Error bars mark 95% confidence intervals of the model fit. α and β are fitted coefficients. Note the varying y-scales.

Appendix B. COVID-19-induced shifts in dietary interests

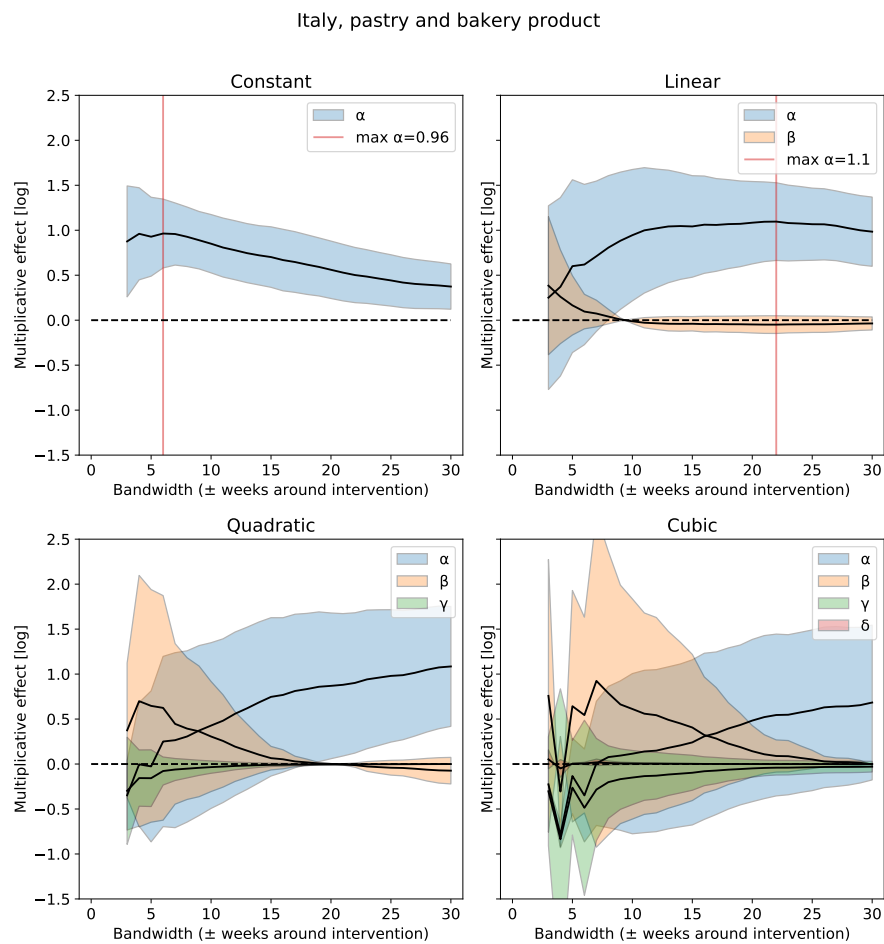


Figure B.8: Estimating the impact of the bandwidth (on x axis) on the fitted coefficients for constant, linear, quadratic, and cubic model. Error bars mark 95% confidence intervals of the fitted coefficients.

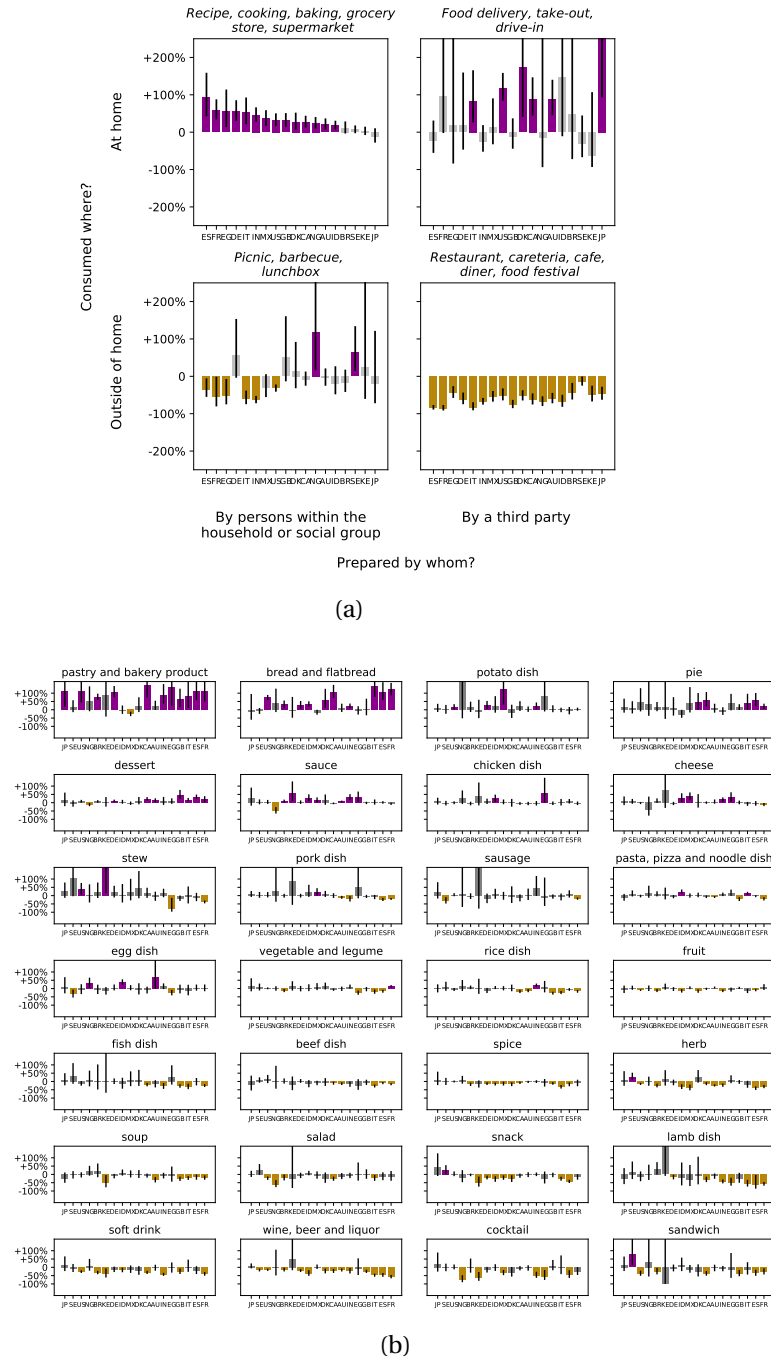


Figure B.9: Short-term effects on the share of interest, estimated with a quadratic model. For each country ($n = 18$), for each food access mode (in (a), $n = 4$) and food category (in (b), $n = 28$), model (Equation 6.1) is fitted on $n = 82$ samples. Bars represent effect estimates (coefficient α estimated with our RDD-based model). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects.

Appendix B. COVID-19-induced shifts in dietary interests

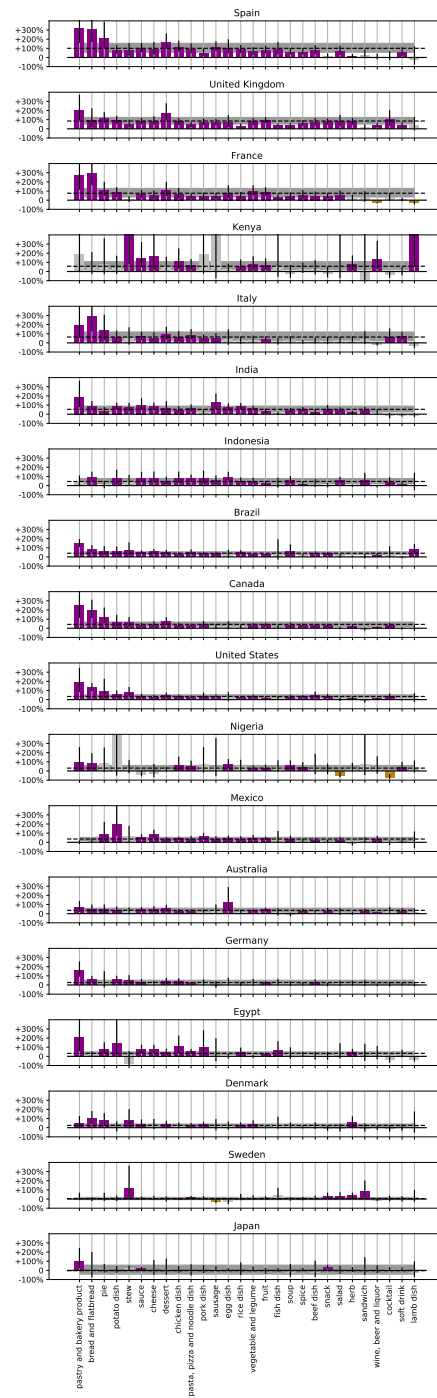


Figure B.10: Short-term effects across food categories, grouped by country. For each country ($n = 18$), for each food category ($n = 28$), model (Equation 6.1) is fitted on $n = 82$ samples. Bars represent effect estimates (coefficient α estimated with our RDD-based model). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects. The gray band marks the 95% CI of the effect on the country-specific total interest in all food entities.

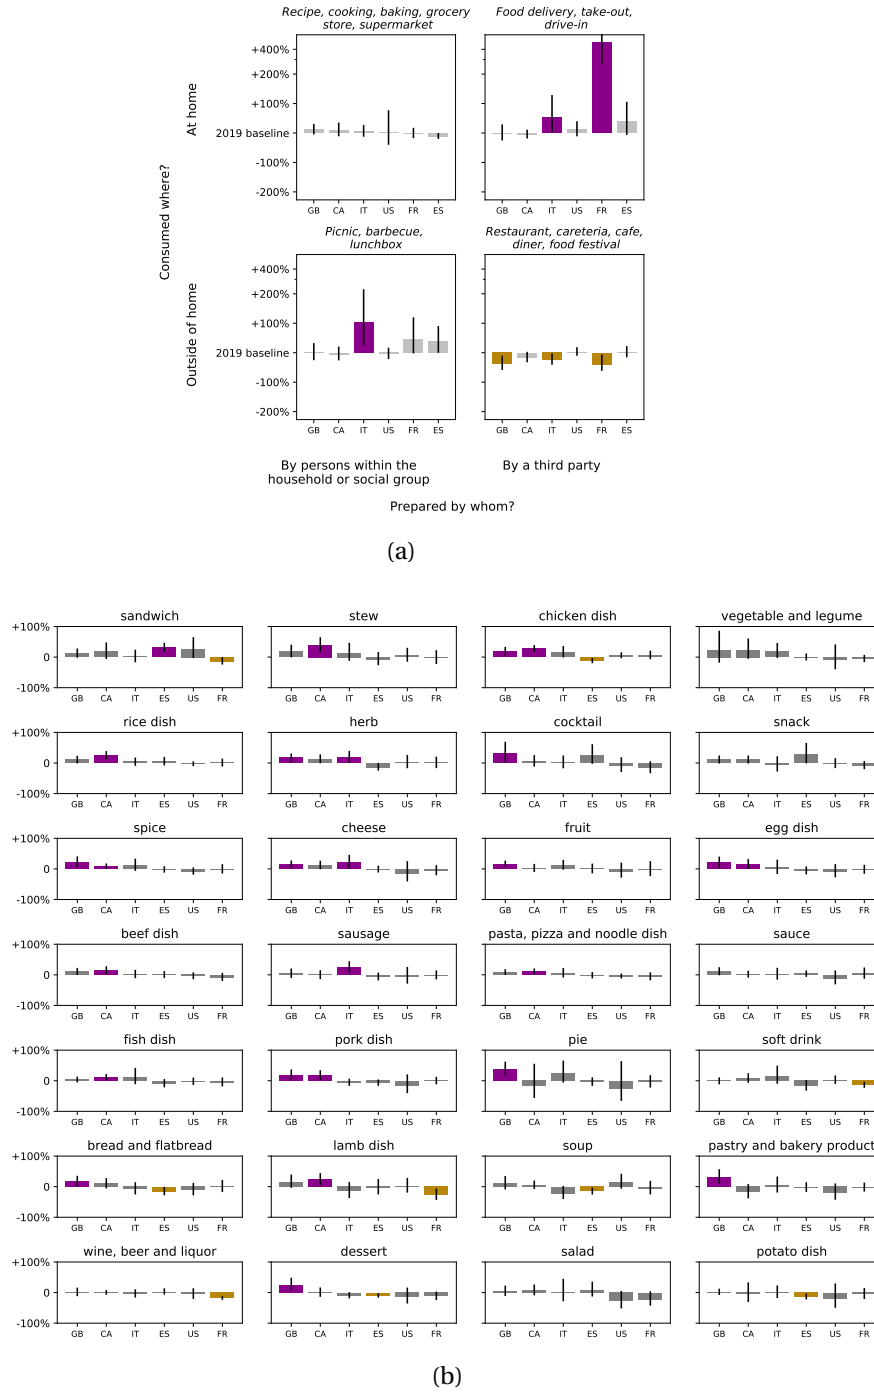




Figure B.11: Short-term effects estimated with a quadratic model, in the second wave. For each country ($n = 6$), for each food access mode (in (a), $n = 4$) and food category (in (b), $n = 28$), model (Equation 6.1) is fitted on $n = 82$ samples. Bars represent effect estimates (coefficient α estimated with our RDD-based model). Error bars mark 95% confidence intervals. Purple marks significant positive ($p < 0.05$), yellow significant negative ($p < 0.05$), and grey marks non-significant effects.

Appendix B. COVID-19-induced shifts in dietary interests



(a)



(b)

Figure B.12: Long-term effect of mobility decrease on food interests. In case the interest did not go back to normal within the 30 weeks after the mobility decrease, we measure how elevated the interest remains at the end of the modelled period, 30 weeks after mobility decrease, compared to the interest in 2019. White marks absence of long term effect when the interest eventually comes back to normal.

Bibliography

- [1] Kibrom A Abay, Kibrom Tafere, and Andinet Woldemichael. “Winners and Losers from COVID-19: Global Evidence from Google Search”. In: *World Bank Policy Research* (2020).
- [2] Sofiane Abbar, Carlos Castillo, and Antonio Sanfilippo. “To Post or Not to Post: Using Online Trends to Predict Popularity of Offline Content”. In: *Proceedings of the 29th Conference on Hypertext and Social Media*. 2018.
- [3] Sofiane Abbar, Yelena Mejova, and Ingmar Weber. “You tweet what you eat: Studying food consumption through Twitter”. In: *Proc. of the 33rd Conference on Human Factors in Computing Systems (CHI)*. 2015.
- [4] Carmen Aceijas, Sabrina Waldhäusl, Nicky Lambert, Simon Cassar, and Rafael Bello-Corassa. “Determinants of health-related lifestyles among university students”. In: *Perspectives in public health* 137.4 (2017).
- [5] Palakorn Achananuparp and Ingmar Weber. “Extracting food substitutes from food diary via distributional similarity”. In: (2016).
- [6] David Ifeoluwa Adelani, Ryota Kobayashi, Ingmar Weber, and Przemyslaw A Grabowicz. “Estimating community feedback effect on topic choice in social media with predictive modeling”. In: *EPJ Data Science* 9.1 (2020).
- [7] Rakesh Agrawal, Heikki Mannila, Ramakrishnan Srikant, Hannu Toivonen, A Inkeri Verkamo, et al. “Fast discovery of association rules.” In: *Advances in knowledge discovery and data mining* 12.1 (1996).
- [8] Alan Agresti. *Categorical data analysis*. Vol. 482. 2003.
- [9] Wasim Ahmed, Josep Vidal-Alaball, Joseph Downing, and Francesc López Seguí. “COVID-19 and the 5G conspiracy theory”. In: *Journal of Medical Internet Research (JMIR)*. 2020.
- [10] Yong-Yeol Ahn, Sebastian E Ahnert, James P Bagrow, and Albert-László Barabási. “Flavor network and the principles of food pairing”. In: *Scientific Reports* 1.1 (2011).
- [11] Luca Maria Aiello, Daniele Quercia, Rossano Schifanella, and Lucia Del Prete. “Tesco Grocery 1.0, a large-scale dataset of grocery purchases in London”. In: *Scientific Data* 7.1 (2020).

Bibliography

- [12] Luca Maria Aiello, Rossano Schifanella, Daniele Quercia, and Lucia Del Prete. “Large-scale and high-resolution analysis of food purchases and health outcomes”. In: *EPJ Data Science* 8.1 (2019).
- [13] Ahmet Aktay et al. *Google COVID-19 Community Mobility Reports: Anonymization Process Description*. 2020.
- [14] Tim Althoff, Pranav Jindal, and Jure Leskovec. “Online actions with offline impact: How online social networks influence online and offline user behavior”. In: *Proc. of the 10th ACM International Conference on Web Search and Data Mining (WSDM)*. 2017.
- [15] Tim Althoff, Hamed Nilforoshan, Jenna Hua, and Jure Leskovec. “Large-scale diet tracking data reveal disparate associations between food environment and diet”. In: *Nature Communications* 13.1 (2022).
- [16] Ana Rita Amaral, Eugénio Rodrigues, Adélio Rodrigues Gaspar, and Álvaro Gomes. “A review of empirical data of sustainability initiatives in university campus operations”. In: *Journal of Cleaner Production* 250 (2020).
- [17] Giuseppe Amato, Paolo Bolettieri, Vinicius Monteiro de Lira, Cristina Ioana Muntean, Raffaele Perego, and Chiara Renso. “Social media image recognition for food trend analysis”. In: *Proc. of the 40th International ACM Conference on Research and Development in Information Retrieval (SIGIR)*. 2017.
- [18] Achraf Ammar, Michael Brach, Khaled Trabelsi, Hamdi Chtourou, Omar Boukhris, Liwa Masmoudi, Bassem Bouaziz, Ellen Bintlage, Daniella How, Mona Ahmed, et al. “Effects of COVID-19 home confinement on eating behaviour and physical activity: results of the ECLB-COVID19 international online survey”. In: *Nutrients* 12.6 (2020).
- [19] Mathieu Andreux, Jean Ogier du Terrail, Constance Beguier, and Eric W Tramel. “Siloed federated learning for multi-centric histopathology datasets”. In: *Domain Adaptation and Representation Transfer, and Distributed and Collaborative Learning*. 2020.
- [20] Joshua D Angrist and Jörn-Steffen Pischke. *Mostly harmless econometrics: An empiricist’s companion*. 2009.
- [21] Sinan Aral, Lev Muchnik, and Arun Sundararajan. “Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks”. In: *Proceedings of the National Academy of Sciences (PNAS)* 106.51 (2009).
- [22] Dan Ariely and Jonathan Levav. “Sequential choice in group settings: Taking the road less traveled and less enjoyed”. In: *Journal of Consumer Research* 27.3 (2000).
- [23] Linda Bacon, Jonathan Wise, Sophie Attwood, and Daniel Vennard. “The language of sustainable diets: A field study exploring the impact of renaming vegetarian dishes on UK Cafe menus”. In: (2019).
- [24] Chris Bail. *Breaking the Social Media Prism*. 2021.

-
- [25] Kylie Ball, Robert W Jeffery, Gavin Abbott, Sarah A McNaughton, and David Crawford. "Is healthy behavior contagious: associations of social norms with physical activity and healthy eating". In: *International Journal of Behavioral Nutrition and Physical Activity* 7.1 (2010).
- [26] Kieron J Barclay, Christofer Edling, and Jens Rydgren. "Peer clustering of exercise and eating behaviours among young adults in Sweden: a cross-sectional study of egocentric network data". In: *BMC Public Health* 13.1 (2013).
- [27] Laura Batlle-Bayer, Rubén Aldaco, Alba Bala, Rita Puig, Jara Laso, Maria Margallo, Ian Vázquez-Rowe, Josep Maria Antó, and Pere Fullana-i-Palmer. "Environmental and nutritional impacts of dietary changes in Spain during the COVID-19 lockdown". In: *Science of the Total Environment* 748 (2020).
- [28] Kayleigh M Beaudry, Izabella A Ludwa, Aysha M Thomas, Wendy E Ward, Bareket Falk, and Andrea R Josse. "First-year university is associated with greater body weight, body composition and adverse dietary changes in males than females". In: *PloS one* 14.7 (2019).
- [29] Matthew J Belanger, Michael A Hill, Angeliki M Angelidi, Maria Dalamaga, James R Sowers, and Christos S Mantzoros. "Covid-19 and disparities in nutrition and obesity". In: *New England Journal of Medicine* 383.11 (2020).
- [30] Brooke M Bell, Donna Spruijt-Metz, George G Vega Yon, Abu S Mondol, Ridwan Alam, Meiyi Ma, Ifat Emi, John Lach, John A Stankovic, and Kayla De la Haye. "Sensing eating mimicry among family members". In: *Translational Behavioral Medicine* 9.3 (2019), pp. 422–430.
- [31] Tarek Ben Hassen, Hamid El Bilali, and Mohammad S Allahyari. "Impact of COVID-19 on Food Behavior and Consumption in Qatar". In: *Sustainability* 12.17 (2020).
- [32] Jonah Berger and Chip Heath. "Where consumers diverge from others: Identity signaling and product domains". In: *Journal of consumer research* 34.2 (2007), pp. 121–134.
- [33] Greyce Luci Bernardo, Manuela Mika Jomori, Ana Carolina Fernandes, and Rossana Pacheco Da Costa Proença. "Food intake of university students". In: *Revista de Nutrição* 30 (2017).
- [34] Mehrab Bin Morshed, Samruddhi Shreeram Kulkarni, Richard Li, Koustuv Saha, Leah Galante Roper, Lama Nachman, Hong Lu, Lucia Mirabella, Sanjeev Srivastava, Munmun De Choudhury, Kaya de Barbaro, Thomas Ploetz, and Gregory D Abowd. "A Real-Time eating detection system for capturing eating moments and triggering ecological momentary assessments to obtain further context: system development and validation study". In: *JMIR mHealth and uHealth* 8.12 (2020).
- [35] Leann Lipps Birch. "Effects of peer models' food choices and eating behaviors on preschoolers' food preferences". In: *Child Development* (1980).

Bibliography

- [36] Bodil Stilling Blichfeldt and Malene Gram. “Lost in transition? Student food consumption”. In: *Higher Education* 65.3 (2013).
- [37] Jason P Block, Matthew W Gillman, Stephanie K Linakis, and Roberta E Goldman. “”If it tastes good, I’m drinking it”: Qualitative study of beverage consumption among college students”. In: *Journal of Adolescent Health* 52.6 (2013).
- [38] DJ Blok, Pepijn van Empelen, FJ Van Lenthe, Jan Hendrik Richardus, and SJ De Vlas. “Unhealthy behaviour is contagious: An invitation to exploit models for infectious diseases”. In: *Epidemiology & Infection* 141.3 (2013).
- [39] Joshua Blumenstock, Gabriel Cadamuro, and Robert On. “Predicting poverty and wealth from mobile phone metadata”. In: *Science* 350.6264 (2015).
- [40] Stefania Boccia, Walter Ricciardi, and John PA Ioannidis. “What other countries can learn from Italy during the COVID-19 pandemic”. In: *JAMA internal medicine* 180 (2020).
- [41] Kurt Bollacker, Colin Evans, Praveen Paritosh, Tim Sturge, and Jamie Taylor. “Freebase: a collaboratively created graph database for structuring human knowledge”. In: *Proc. of the 2008 ACM International Conference on Management of Data (SIGMOD)*. 2008.
- [42] Giovanni Bonaccorsi et al. “Economic and social consequences of human mobility restrictions under COVID-19”. In: *Proceedings of the National Academy of Sciences (PNAS)* (2020).
- [43] Sadie Boniface, James Kneale, and Nicola Shelton. “Drinking pattern is more strongly associated with under-reporting of alcohol consumption than socio-demographic factors: evidence from a mixed-methods study”. In: *BMC public health* 14.1 (2014).
- [44] Benjamin Born, Alexander M Dietrich, and Gernot J Müller. “The lockdown effect: A counterfactual for Sweden”. In: *Plos one* 16.4 (2021).
- [45] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. “Food-101–mining discriminative components with random forests”. In: *European Conference on Computer Vision*. Springer. 2014.
- [46] Ann Bowling. “Mode of questionnaire administration can have serious effects on data quality”. In: *Journal of Public Health* 27.3 (2005).
- [47] Janet M Box-Steffensmeier, Jean Burgess, Maurizio Corbetta, Kate Crawford, Esther Duflo, Laurel Fogarty, Alison Gopnik, Sari Hanafi, Mario Herrero, Ying-yi Hong, et al. “The future of human behaviour research”. In: *Nature Human Behaviour* 6.1 (2022).
- [48] Ralph Allan Bradley and Milton E Terry. “Rank analysis of incomplete block designs: I. The method of paired comparisons”. In: *Biometrika* 39.3/4 (1952).
- [49] Abel Brodeur, Andrew E Clark, Sarah Fleche, and Nattavudh Powdthavee. “COVID-19, lockdowns and well-being: Evidence from Google Trends”. In: *Journal of public economics* 193 (2021).

-
- [50] Tamara Bucher, Clare Collins, Megan E Rollo, Tracy A McCaffrey, Nienke De Vlieger, Daphne Van der Bend, Helen Truby, and Federico JA Perez-Cueto. “Nudging consumers towards healthier choices: a systematic review of positional influences on food choice”. In: *British Journal of Nutrition* 115.12 (2016).
- [51] David L Buckeridge, Katia Charland, Alice Labban, and Yu Ma. “A method for neighborhood level surveillance of food purchasing”. In: *Annals of the New York Academy of Sciences* 1331.1 (2014).
- [52] Bekir Bulut, Nazım Ercüment Beyhun, Murat Topbaş, and Gamze Can. “Energy drink use in university students and associated factors”. In: *Journal of community health* 39.5 (2014).
- [53] Moira Burke, Justin Cheng, and Bethany de Gant. “Social comparison and Facebook: Feedback, positivity, and opportunities for comparison”. In: *Proceedings of the 2020 Conference on Human Factors in Computing Systems (CHI)*. 2020.
- [54] Karen Camelo and Marta Elliott. “Food insecurity and academic achievement among college students at a public university in the United States”. In: *Journal of College Student Development* 60.3 (2019).
- [55] Nicoletta Cavazza, Margherita Guidetti, and Fabrizio Butera. “Portion size tells who I am, food type tells who you are: Specific functions of amount and type of food in same-and opposite-sex dyadic eating contexts”. In: *Appetite* 112 (2017).
- [56] Stevie Chancellor, Zhiyuan Lin, Erica L. Goodman, Stephanie Zerwas, and Munmun De Choudhury. “Quantifying and predicting mental illness severity in online pro-eating disorder communities”. In: *Proc. of the 2016 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2016.
- [57] Stevie Chancellor, Jessica Annette Pater, Trustin Clear, Eric Gilbert, and Munmun De Choudhury. “# thyghgapp: Instagram content moderation and lexical variation in pro-eating disorder communities”. In: *Proc. of the 2016 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2016.
- [58] Tanya L Chartrand and John A Bargh. “The chameleon effect: The perception–behavior link and social interaction.” In: *Journal of personality and social psychology* 76.6 (1999), p. 893.
- [59] Tanya L Chartrand and Jessica L Lakin. “The antecedents and consequences of human behavioral mimicry”. In: *Annual review of psychology* 64 (2013), pp. 285–308.
- [60] Tanya L Chartrand and Rick Van Baaren. “Human mimicry”. In: *Advances in experimental social psychology* 41 (2009), pp. 219–274.
- [61] Xi Chen, Paul N Bennett, Kevyn Collins-Thompson, and Eric Horvitz. “Pairwise ranking aggregation in a crowdsourced setting”. In: *Proc. of the 6th ACM International Conference on Web Search and Data Mining (WSDM)*. 2013.
- [62] Lauren Chenarides, Carola Grebitus, Jayson L Lusk, and Iryna Printezis. “Food consumption behavior during the COVID-19 pandemic”. In: *Agribusiness* 37.1 (2021).

Bibliography

- [63] Justin Cheng, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. "How community feedback shapes user behavior". In: *Proc. of the eighth International AAAI Conference on Weblogs and Social Media (ICWSM)*. 2014.
- [64] Hyunyoung Choi and Hal Varian. "Predicting the present with Google Trends". In: *Economic record* 88 (2012).
- [65] Martin J Chorley, Luca Rossi, Gareth Tyson, and Matthew J Williams. "Pub crawling at scale: tapping untapped to explore social drinking". In: *Proc. of the 10th International AAAI Conference on Web and Social Media (ICWSM)*. 2016.
- [66] Nicholas A Christakis and James H Fowler. "The spread of obesity in a large social network over 32 years". In: *New England Journal of Medicine (NEJM)* 357.4 (2007).
- [67] Chelsea D Christie and Frances S Chen. "Vegetarian or meat? Food choice modeling of main dishes occurs outside of awareness". In: *Appetite* 121 (2018).
- [68] Chia-Fang Chung, Elena Agapie, Jessica Schroeder, Sonali Mishra, James Fogarty, and Sean A Munson. "When personal tracking becomes social: Examining the use of Instagram for healthy eating". In: *Proc. of the 2017 Conference on Human Factors in Computing Systems (CHI)*. 2017.
- [69] Chia-Fang Chung, Jonathan Cook, Elizabeth Bales, Jasmine Zia, and Sean A Munson. "More than telemonitoring: health provider use and nonuse of life-log data in irritable bowel syndrome and weight management". In: *Journal of Medical Internet Research* 17.8 (2015).
- [70] Chia-Fang Chung, Qiaosi Wang, Jessica Schroeder, Allison Cole, Jasmine Zia, James Fogarty, and Sean A Munson. "Identifying and planning for individualized change: Patient-provider collaboration using lightweight food diaries in healthy eating and irritable bowel syndrome". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* 3.1 (2019).
- [71] Çiler Çilingiroğlu. "The concept of "Neolithic package": considering its meaning and applicability". In: *Documenta Praehistorica* 32 (2005).
- [72] James M Clay and Matthew O Parker. "Alcohol use and misuse during the COVID-19 pandemic: a potential public health crisis?" In: *The Lancet Public Health* 5.5 (2020).
- [73] William Colglazier. "Sustainable development agenda: 2030". In: *Science* 349.6252 (2015).
- [74] Rachel C Colley, Tracey Bushnik, and Kellie Langlois. "Exercise and screen time during the COVID-19 pandemic". In: *Health Reports* (2020).
- [75] Emily IM Collins, Jason M Thomas, Eric Robinson, Paul Aveyard, Susan A Jebb, C Peter Herman, and Suzanne Higgs. "Two observational studies examining the effect of a social norm and a health message on the purchase of vegetables in student canteen settings". In: *Appetite* 132 (2019).

-
- [76] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. "Algorithmic decision making and the cost of fairness". In: *Proceedings of the 23rd ACM International Conference on Knowledge Discovery and Data Mining (KDD)*. 2019.
- [77] Felicia Cordeiro, Elizabeth Bales, Erin Cherry, and James Fogarty. "Rethinking the mobile food journal: Exploring opportunities for lightweight photo-based capture". In: *Proc. of the 2015 Conference on Human Factors in Computing Systems (CHI)*. 2015.
- [78] Felicia Cordeiro, Daniel A. Epstein, Edison Thomaz, Elizabeth Bales, Arvind K. Jagannathan, Gregory D. Abowd, and James Fogarty. "Barriers and negative nudges: Exploring challenges in food journaling". In: *Proc. of the 2015 Conference on Human Factors in Computing Systems (CHI)*. 2015.
- [79] David J Crandall, Lars Backstrom, Dan Cosley, Siddharth Suri, Daniel Huttenlocher, and Jon Kleinberg. "Inferring social ties from geographic coincidences". In: *Proceedings of the National Academy of Sciences (PNAS)* 107.52 (2010).
- [80] Tegan Cruwys, Kirsten E. Bevelander, and Roel C.J. Hermans. "Social modeling of eating: A review of when and why social influence affects food intake and choice". In: *Appetite* 86 (2015).
- [81] Andrea Cuttone, Sune Lehmann, and Marta C Gonzalez. "Understanding predictability and exploration in human mobility". In: *EPJ Data Science* 7 (2018), pp. 1–17.
- [82] Cristian Danescu-Niculescu-Mizil, Robert West, Dan Jurafsky, Jure Leskovec, and Christopher Potts. "No country for old members: User lifecycle and linguistic change in online communities". In: *Proc. of the 22nd International Conference on World Wide Web (TheWebConf)*. 2013.
- [83] Peter Dannenberg, Martina Fuchs, Tim Riedler, and Cathrin Wiedemann. "Digital transition by COVID-19 pandemic? The German food online retail". In: *Tijdschrift voor economische en sociale geografie* 111.3 (2020).
- [84] Vedant Das Swain, Jiajia Xie, Maanit Madan, Sonia Sargolzaei, James Cai, Munmun De Choudhury, Gregory D Abowd, Lauren N Steimle, and B Aditya Prakash. "Empirical networks for localized COVID-19 interventions using WiFi infrastructure at university campuses". In: *medRxiv* (2021).
- [85] Martin P Davoren, Jakob Demant, Frances Shiely, and Ivan J Perry. "Alcohol consumption among university students in Ireland and the United Kingdom from 2002 to 2014: a systematic review". In: *BMC public health* 16.1 (2016).
- [86] Munmun De Choudhury. "Anorexia on tumblr: A characterization study". In: *Proc. of the 5th International Conference on Digital Health*. 2015.
- [87] Munmun De Choudhury, Meredith Ringel Morris, and Ryen W White. "Seeking and sharing health information online: Comparing search engines and social media". In: *Proc. of the 2014 Conference on Human Factors in Computing Systems (CHI)*. 2014.

Bibliography

- [88] Munmun De Choudhury, Sanket Sharma, and Emre Kıcıman. "Characterizing dietary choices, nutrition, and language in food deserts via social media". In: *Proc. of the 2016 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2016.
- [89] Kayla De La Haye, Garry Robins, Philip Mohr, and Carlene Wilson. "Adolescents' intake of junk food: Processes and mechanisms driving consumption similarities among friends". In: *Journal of Research on Adolescence* 23.3 (2013).
- [90] Kayla De La Haye, Garry Robins, Philip Mohr, and Carlene Wilson. "Homophily and contagion as explanations for weight similarities among adolescent friends". In: *Journal of Adolescent Health* 49.4 (2011), pp. 421–427.
- [91] Kayla de la Haye, Garry Robins, Philip Mohr, and Carlene Wilson. "How physical activity shapes, and is shaped by, adolescent friendships". In: *Social Science and Medicine* 73.5 (2011).
- [92] Kayla de la Haye, Garry Robins, Philip Mohr, and Carlene Wilson. "Obesity-related behaviors in adolescent friendship networks". In: *Social Networks* 32.3 (2010).
- [93] Mary Delaney and Mary McCarthy. "Food choice and health across the life course: A qualitative study examining food choice in older Irish adults". In: *Journal of Food Products Marketing* 17.2-3 (2011).
- [94] Tom Deliens, Peter Clarys, Ilse De Bourdeaudhuij, and Benedicte Deforche. "Determinants of eating behaviour in university students: a qualitative study using focus group discussions". In: *BMC public health* 14.1 (2014).
- [95] Tom Deliens, Benedicte Deforche, Ilse De Bourdeaudhuij, and Peter Clarys. "Determinants of physical activity and sedentary behaviour in university students: a qualitative study using focus group discussions". In: *BMC public health* 15.1 (2015).
- [96] Tom Deliens, Rob Van Crombruggen, Sofie Verbruggen, Ilse De Bourdeaudhuij, Benedicte Deforche, and Peter Clarys. "Dietary interventions among university students: A systematic review". In: *Appetite* 105 (2016).
- [97] Anneliese Depoux, Sam Martin, Emilie Karafillakis, Raman Preet, Annelies Wilder-Smith, and Heidi Larson. *The pandemic of social media panic travels faster than the COVID-19 outbreak*. 2020.
- [98] Melanie Deschasaux-Tanguy, Nathalie Druesne-Pecollo, Younes Esseddik, Fabien Szabo de Edelenyi, Benjamin Alles, Valentina A Andreeva, Julia Baudry, Helene Charreire, Valerie Deschamps, Manon Egnell, et al. "Diet and physical activity during the COVID-19 lockdown period (March-May 2020): Results from the French NutriNet-Sante cohort study". In: *MedRxiv* (2020).
- [99] Stephen Devereux, Christophe Béné, and John Hoddinott. "Conceptualising COVID-19's impacts on household food security". In: *Food Security* 12.4 (2020).

-
- [100] Laura Di Renzo, Paola Gualtieri, Francesca Pivari, Laura Soldati, Alda Attinà, Giulia Cinelli, Claudia Leggeri, Giovanna Caparello, Luigi Barrea, Francesco Scerbo, et al. “Eating habits and lifestyle changes during COVID-19 lockdown: an Italian survey”. In: *Journal of translational medicine* 18 (2020).
 - [101] Neele Dijkstra. “Cross-modal recipe analysis for fine-grained geographical mapping of food consumption from images and supermarket sales”. MA thesis. 2022.
 - [102] Annie Doubleday, Youngjun Choe, Tania Busch Isaksen, Scott Miles, and Nicole A Errett. “How did outdoor biking and walking change during COVID-19?: A case study of three US cities”. In: *PLOS one* 16.1 (2021).
 - [103] Erin M Driver, Adam Gushgari, Jing Chen, and Rolf U Halden. “Alcohol, nicotine, and caffeine consumption on a public US university campus determined by wastewater-based epidemiology”. In: *Science of the Total Environment* 727 (2020).
 - [104] Yang Du, Shuang Rong, Yangbo Sun, Buyun Liu, Yuxiao Wu, Linda G. Snetselaar, Robert B. Wallace, and Wei Bao. “Association Between Frequency of Eating Away-From-Home Meals and Risk of All-Cause and Cause-Specific Mortality”. In: *Journal of the Academy of Nutrition and Dietetics* (2021).
 - [105] Elizabeth Dunford, Helen Trevena, Chester Goodsell, Ka Hung Ng, Jacqui Webster, Audra Millis, Stan Goldstein, Orla Hugueniot, and Bruce Neal. “FoodSwitch: a mobile phone app to enable consumers to make healthier food choices and crowdsourcing of national food composition data”. In: *JMIR mHealth and uHealth* (2014).
 - [106] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. “Calibrating noise to sensitivity in private data analysis”. In: *Theory of Cryptography Conference*. Springer. 2006.
 - [107] Cynthia Dwork, Aaron Roth, et al. “The algorithmic foundations of differential privacy”. In: *Foundations and Trends® in Theoretical Computer Science* 9.3–4 (2014).
 - [108] Naoki Egami, Christian J Fong, Justin Grimmer, Margaret E Roberts, and Brandon M Stewart. “How to make causal inferences using texts”. In: *arXiv preprint arXiv:1802.02163* (2018).
 - [109] Marla E. Eisenberg, Dianne Neumark-Sztainer, Mary Story, and Cheryl Perry. “The role of social norms and friends’ influences on unhealthy weight-control behaviors among adolescent girls”. In: *Social Science & Medicine* 60.6 (2005).
 - [110] Douae El Fatouhi, Harris Héritier, Chloé Allémann, Laurent Malisoux, Nasser Laouali, Jean-Pierre Riveline, Marcel Salathé, and Guy Fagherazzi. “Associations Between Device-Measured Physical Activity and Glycemic Control and Variability Indices Under Free-Living Conditions”. In: *Diabetes Technology & Therapeutics* 24.3 (2022).
 - [111] Gunther Eysenbach. “Infodemiology: The epidemiology of (mis) information”. In: *The American journal of medicine* 113.9 (2002).

Bibliography

- [112] Thomas Favale, Francesca Soro, Martino Trevisan, Idilio Drago, and Marco Mellia. “Campus traffic and e-Learning during COVID-19 pandemic”. In: *Computer Networks* (2020).
- [113] Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, and Margaret E Roberts. “Causal inference in natural language processing: Estimation, prediction, interpretation and beyond”. In: *arXiv preprint arXiv:2109.00725* (2021).
- [114] Charles Feldman, Heather Harwell, and Joseph Brusca. “Using student opinion and design inputs to develop an informed university foodservice menu”. In: *Appetite* 69 (2013).
- [115] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Eric Pujol, Igmar Poesse, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, et al. “Implications of the COVID-19 Pandemic on the Internet Traffic”. In: *Broadband Coverage in Germany*. 2021.
- [116] Ana Carolina Fernandes, Renata Carvalho de Oliveira, Vanessa Mello Rodrigues, Giovanna Medeiros Rataichesk Fiates, and Rossana Pacheco da Costa Proença. “Perceptions of university students regarding calories, food healthiness, and the importance of calorie information in menu labelling”. In: *Appetite* 91 (2015).
- [117] Gerda I.J. Feunekes, Cees de Graaf, Saskia Meyboom, and Wija A. van Staveren. “Food choice and fat intake of adolescents and adults: Associations of intakes within social networks”. In: *Preventive Medicine* 27.5 (1998).
- [118] Casey Fiesler, Michaelanne Dye, Jessica L Feuston, Chaya Hiruncharoenvate, Clayton J Hutto, Shannon Morrison, Parisa Khanipour Roshan, Umashanthi Pavalanathan, Amy S Bruckman, Munmun De Choudhury, et al. “What (or who) is public? Privacy settings and social media content sharing”. In: *Proc. of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2017.
- [119] Tara Finnerty, Sue Reeves, Jaqueline Dabinett, Yvonne M Jeanes, and Claus Vögele. “Effects of peer influence on dietary intake and physical activity in schoolchildren”. In: *Public Health Nutrition* 13.3 (2010).
- [120] Joseph Firth, Brendon Stubbs, Scott B Teasdale, Philip B Ward, Nicola Veronese, Nitin Shivappa, James R Hebert, Michael Berk, Alison R Yung, and Jerome Sarris. “Diet as a hot topic in psychiatry: A population-scale study of nutritional intake and inflammatory potential in severe mental illness”. In: *World Psychiatry* 17.3 (2018).
- [121] Brianna Fjeldsoe, Maike Neuhaus, Elisabeth Winkler, and Elizabeth Eakin. “Systematic review of maintenance of behavior change following physical activity and dietary interventions.” In: *Health Psychology* 30.1 (2011).
- [122] Emily W Flanagan, Robbie A Beyl, S Nicole Fearnbach, Abby D Altazan, Corby K Martin, and Leanne M Redman. “The impact of COVID-19 stay-at-home orders on health behaviors in adults”. In: *Obesity* 29.2 (2021).

-
- [123] Seth Flaxman et al. *Report 13: Estimating the number of infections and the impact of non-pharmaceutical interventions on COVID-19 in 11 European countries*. Report. Imperial College London, 2020.
 - [124] Adam Fletcher, Chris Bonell, and Annik Sorhaindo. “You are what your friends eat: systematic review of social network analyses of young people’s eating behaviours and bodyweight”. In: *Journal of Epidemiology & Community Health* 65.6 (2011).
 - [125] David R Foxcroft, Maria Teresa Moreira, Nerissa ML Almeida Santimano, and Lesley A Smith. “Social norms information for alcohol misuse in university and college students”. In: *Cochrane database of systematic reviews* 12 (2015).
 - [126] Joshua Freedman and Dan Jurafsky. “Authenticity in America: Class distinctions in potato chip advertising”. In: *Gastronomica: The journal of food and culture* 11.4 (2011).
 - [127] Marjorie R Freedman and Rebecca J Rubinstein. “Obesity and food choices among faculty and staff at a large urban university”. In: *Journal of American College Health* 59.3 (2010).
 - [128] Charles NC Freitas, Filipe R Cordeiro, and Valmir Macario. “Myfood: A food segmentation and classification system to aid nutritional monitoring”. In: *2020 33rd SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*. 2020.
 - [129] Batya Friedman. “Value-sensitive design”. In: *interactions* 3.6 (1996).
 - [130] Eline Frison and Steven Eggermont. “Browsing, posting, and liking on Instagram: The reciprocal relationships between different types of Instagram use and adolescents’ depressed mood”. In: *Cyberpsychology, Behavior, and Social Networking* 20.10 (2017).
 - [131] Emmanuela Gakidou, Ashkan Afshin, Amanuel Alemu Abajobir, Kalkidan Hassen Abate, Cristiana Abbafati, Kaja M Abbas, Foad Abd-Allah, Abdishakur M Abdulle, Semaw Ferede Abera, Victor Aboyans, et al. “Global, regional, and national comparative risk assessment of 84 behavioural, environmental and occupational, and metabolic risks or clusters of risks, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016”. In: *The Lancet* 390.10100 (2017).
 - [132] Ruth Garcia-Gavilanes, Anders Mollgaard, Milena Tsvetkova, and Taha Yasseri. “The memory remains: Understanding collective memory in the digital age”. In: *Science advances*. 2017.
 - [133] Venkata Rama Kiran Garimella, Abdulrahman Alfayad, and Ingmar Weber. “Social media image analysis for public health”. In: *Proc. of the 2016 Conference on Human Factors in Computing Systems (CHI)*. 2016.
 - [134] Emma E Garnett, Andrew Balmford, Chris Sandbrook, Mark A Pilling, and Theresa M Marteau. “Impact of increasing vegetarian availability on meal selection and sales in cafeterias”. In: *Proceedings of the National Academy of Sciences (PNAS)* 116.42 (2019).
 - [135] Emma E Garnett, Theresa M Marteau, Chris Sandbrook, Mark A Pilling, and Andrew Balmford. “Order of meals at the counter and distance between options affect student cafeteria vegetarian sales”. In: *Nature Food* 1.8 (2020).

Bibliography

- [136] Gilie Gefen, Omer Ben-Porat, Moshe Tennenholtz, and Elad Yom-Tov. “Privacy, altruism, and experience: Estimating the perceived value of Internet data for medical uses”. In: *Companion Proceedings of the Web Conference 2020*. 2020.
- [137] L Gemming and C Ni Mhurchu. “Dietary under-reporting: what foods and which meals are typically under-reported?”. In: *European journal of clinical nutrition* 70.5 (2016).
- [138] Nicholas Generous et al. “Global disease monitoring and forecasting with Wikipedia”. In: *PLOS computational biology*. 2014.
- [139] Sarah Gerritsen, Victoria Egli, Rajshri Roy, Jill Haszard, Charlotte De Backer, Lauranna Teunissen, Isabelle Cuykx, Paulien Decorte, Sara Pabian Pabian, Kathleen Van Royen, et al. “Seven weeks of home-cooked meals: changes to New Zealanders’ grocery shopping, cooking and eating during the COVID-19 lockdown”. In: *Journal of the royal society of new zealand* (2020).
- [140] Sucheta Ghoshal and Amy Bruckman. “The role of social computing technologies in grassroots movement building”. In: *ACM Transactions on Computer-Human Interaction (TOCHI)* 26.3 (2019).
- [141] Davide Giacalone, Michael Bom Frøst, and Celia Rodriguez-Pérez. “Reported changes in dietary habits during the Covid-19 lockdown in the Danish population: the Danish COVIDiet study”. In: *Frontiers in nutrition* 7 (2020).
- [142] Jeremy Ginsberg, Matthew H Mohebbi, Rajan S Patel, Lynnette Brammer, Mark S Smolinski, and Larry Brilliant. “Detecting influenza epidemics using search engine query data”. In: *Nature* 457.7232 (2009).
- [143] Joel Gittelsohn, Megan Rowan, and Preety Gadhoke. “Interventions in small food stores to change the food environment, improve diet, and reduce risk of chronic disease”. In: *Preventing Chronic Disease* 9 (2012).
- [144] Kristina Gligorić, Ashton Anderson, and Robert West. “Causal effects of brevity on style and success in social media”. In: *Proc. of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* (2019).
- [145] Kristina Gligorić, Ashton Anderson, and Robert West. “How Constraints Affect Content: The Case of Twitter’s Switch from 140 to 280 Characters”. In: *Proc. of the International AAAI Conference on Web and Social Media (ICWSM)*. 2018.
- [146] Kristina Gligorić, Arnaud Chiolero, Emre Kıcıman, Ryen W White, and Robert West. “Population-scale dietary interests during the COVID-19 pandemic”. In: *Nature Communications* 13.1 (2022).
- [147] Kristina Gligorić, Justyna Czestochowska, Ashton Anderson, and Robert West. “Anticipated versus Actual Effects of Platform Design Change: A Case Study of Twitter’s Character Limit”. In: *Proc. of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* (2022).

-
- [148] Kristina Gligorić, Irena Djordjević, and Robert West. “Biased Bytes: On the Validity of Estimating Food Consumption from Digital Traces”. In: *Proc. of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* (2022).
- [149] Kristina Gligorić, George Lifchits, Robert West, and Ashton Anderson. “Linguistic effects on news headline success: Evidence from thousands of online field experiments (Registered Report Protocol)”. In: *Plos one* 16.9 (2021).
- [150] Kristina Gligorić, Manoel Horta Ribeiro, Martin Müller, Olesia Altunina, Maxime Peyrard, Marcel Salathé, Giovanni Colavizza, and Robert West. “Experts and authorities receive disproportionate attention on Twitter during the COVID-19 crisis”. In: *arXiv preprint arXiv:2008.08364* (2020).
- [151] Kristina Gligorić, Ryen W. White, Emre Kiciman, Eric Horvitz, Arnaud Chiolero, and Robert West. “Formation of Social Ties Influences Food Choice: A Campus-Wide Longitudinal Study”. In: *Proc. of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* (2021).
- [152] Amanda Grech, Lana Hebden, Rajshri Roy, and Margaret Allman-Farinelli. “Are products sold in university vending machines nutritionally poor? A food environment audit”. In: *Nutrition & Dietetics* 74.2 (2017).
- [153] Sander Greenland, Judea Pearl, and James M Robins. “Causal diagrams for epidemiologic research”. In: *Epidemiology* (1999), pp. 37–48.
- [154] Samuel L Groseclose and David L Buckeridge. “Public health surveillance systems: recent advances in their use and evaluation”. In: *Annual Review of Public Health* 38 (2017).
- [155] Dabo Guan, Daoping Wang, Stephane Hallegatte, Steven J Davis, Jingwen Huo, Shuping Li, Yangchun Bai, Tianyang Lei, Qianyu Xue, D’Maris Coffman, et al. “Global supply-chain effects of COVID-19 control measures”. In: *Nature Human behaviour* (2020).
- [156] Venkat N Gudivada, Ricardo Baeza-Yates, and Vijay V Raghavan. “Big data: Promises and problems”. In: *Computer* 48.03 (2015).
- [157] Sumedha Gupta et al. *Effects of social distancing policy on labor market outcomes*. Tech. rep. National Bureau of Economic Research, 2020.
- [158] Suzan A Haidar, NK De Vries, Mirey Karavetian, and Rola El-Rassi. “Stress, anxiety, and weight gain among university and college students: a systematic review”. In: *Journal of the Academy of Nutrition and Dietetics* 118.2 (2018).
- [159] Thomas Hale, Sam Webster, Anna Petherick, Toby Phillips, and Beatriz Kira. “Oxford COVID-19 Government Response tracker”. In: <https://bit.ly/2ZF3Q7h> (2020). Accessed: 2020-06-04.
- [160] Karen S Hamrick and Ket McClelland. *Americans’ eating patterns and time spent on food: The 2014 eating & health module data*. Tech. rep. 2016.

Bibliography

- [161] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. “Bias in online freelance marketplaces: Evidence from taskrabbit and fiverr”. In: *Proc. of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2017.
- [162] Brook E Harmon, Melinda Forthofer, Erin O Bantum, and Claudio R Nigg. “Perceived influence and college students’ diet and physical activity behaviors: an examination of ego-centric social networks”. In: *BMC Public Health* 16.1 (2016).
- [163] Jennifer L Harris, Sarah E Speers, Marlene B Schwartz, and Kelly D Brownell. “US food company branded advergames on the Internet: Children’s exposure and effects on snack consumption”. In: *Journal of Children and Media* 6.1 (2012).
- [164] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. “Deep residual learning for image recognition”. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2014.
- [165] Derek Headey, Rebecca Heidkamp, Saskia Osendarp, Marie Ruel, Nick Scott, Robert Black, Meera Shekar, Howarth Bouis, Augustin Flory, Lawrence Haddad, et al. “Impacts of COVID-19 on childhood malnutrition and nutrition-related mortality”. In: *The Lancet* 396.10250 (2020).
- [166] Kristina Heinonen and Tore Strandvik. “Reframing service innovation: COVID-19 as a catalyst for imposed service innovation”. In: *Journal of Service Management* (2020).
- [167] Lisa Henry. “Understanding food insecurity among college students: Experience, motivation, and local solutions”. In: *Annals of Anthropological Practice* 41.1 (2017).
- [168] Roel CJ Hermans, Junilla K Larsen, C Peter Herman, and Rutger CME Engels. “Effects of social modeling on young women’s nutrient-dense food intake”. In: *Appetite* 53.1 (2009).
- [169] Roel CJ Hermans, Anna Lichtwarck-Aschoff, Kirsten E Bevelander, C Peter Herman, Junilla K Larsen, and Rutger CME Engels. “Mimicry of food intake: The dynamic interplay between eating companions”. In: *PloS one* 7.2 (2012), e31027.
- [170] Marion M Hetherington, Annie S Anderson, Geraldine NM Norton, and Lisa Newson. “Situational effects on meal intake: A comparison of eating alone and eating with others”. In: *Physiology & Behavior* 88.4-5 (2006).
- [171] Kyle S Hickmann, Geoffrey Fairchild, Reid Priedhorsky, Nicholas Generous, James M Hyman, Alina Deshpande, and Sara Y Del Valle. “Forecasting the 2013–2014 influenza season using Wikipedia”. In: *PLoS computational biology* 11.5 (2015).
- [172] Suzanne Higgs and Jason Thomas. “Social influences on eating”. In: *Current Opinion in Behavioral Sciences* 9 (2016).
- [173] Kashmir Hill. “How target figured out a teen girl was pregnant before her father did”. In: *Forbes, Inc* (2012).

-
- [174] Jake M Hofman, Duncan J Watts, Susan Athey, Filiz Garip, Thomas L Griffiths, Jon Kleinberg, Helen Margetts, Sendhil Mullainathan, Matthew J Salganik, Simine Vazire, et al. "Integrating explanation and prediction in computational social science". In: *Nature* 595.7866 (2021).
 - [175] Tanya M Horacek, Maria B Erdman, Carol Byrd-Bredbenner, Gale Carey, Sarah M Colby, Geoffrey W Greene, Wen Guo, Kendra K Kattelman, Melissa Olfert, Jennifer Walsh, et al. "Assessment of the dining environment on and near the campuses of fifteen post-secondary institutions". In: *Public health nutrition* 16.7 (2013).
 - [176] Abigail L Horn, Brooke M Bell, Bernardo Garcia Bulle Bueno, Mohsen Bahrami, Burcin Bozkaya, Yan Cui, John P Wilson, Alex Pentland, Esteban Moro Egido, and Kayla de la Haye. "Investigating mobility-based fast food outlet visits as indicators of dietary intake and diet-related disease". In: *medRxiv* (2021).
 - [177] Patrick D. Howell, Layla D. Martin, Hesamoddin Salehian, Chul Lee, Kyler M. Eastman, and Joohyun Kim. "Analyzing taste preferences from crowdsourced food entries". In: *Proc. of the 6th International Conference on Digital Health Conference (DH)*. 2016.
 - [178] Zhiting Hu and Li Erran Li. "A Causal Lens for Controllable Text Generation". In: *Advances in Neural Information Processing Systems* 34 (2021).
 - [179] Roger Hughes, Irene Serebryanikova, Katherine Donaldson, and Michael Leveritt. "Student food insecurity: The skeleton in the university closet". In: *Nutrition & dietetics* 68.1 (2011).
 - [180] Holly R Hull, Duncan Radley, Mary K Dinger, and David A Fields. "The effect of the Thanksgiving holiday on weight gain". In: *Nutrition journal* 5.1 (2006).
 - [181] Nick Huntington-Klein. *The effect: An introduction to research design and causality*. 2021.
 - [182] Marco Iacoboni, Roger P Woods, Marcel Brass, Harold Bekkering, John C Mazziotta, and Giacomo Rizzolatti. "Cortical mechanisms of human imitation". In: *Science* 286.5449 (1999), pp. 2526–2528.
 - [183] Mohammed Iddir, Alex Brito, Giulia Dingo, Sofia Sosa Fernandez Del Campo, Hanen Samouda, Michael R La Frano, and Torsten Bohn. "Strengthening the immune system and reducing inflammation and oxidative stress through diet and nutrition: considerations during the COVID-19 crisis". In: *Nutrients* 12.6 (2020).
 - [184] *Inetrnet Demographics & Use*. <https://www.statista.com/markets/424/topic/537/demographics-use>. Accessed: 2021-03-01.
 - [185] Joanne Ingram, Greg Maciejewski, and Christopher J Hand. "Changes in diet, sleep, and physical activity are associated with differences in negative mood during COVID-19 lockdown". In: *Frontiers in psychology* 11 (2020).

Bibliography

- [186] Leila Cheikh Ismail, Tareq M Osaili, Maysm N Mohamad, Amina Al Marzouqi, Amjad H Jarrar, Antonis Zampelas, Carla Habib-Mourad, Dima Omar Abu Jamous, Habiba I Ali, Haleama Al Sabbah, et al. "Assessment of eating habits and lifestyle during the coronavirus 2019 pandemic in the Middle East and North Africa region: a cross-sectional study". In: *British Journal of Nutrition* (2020).
- [187] Tomoharu Iwata, Shinji Watanabe, Takeshi Yamada, and Naonori Ueda. "Topic tracking model for analyzing consumer purchase behavior". In: *Proc. Twenty-First International Joint Conference on Artificial Intelligence (IJCAI)*. 2009.
- [188] Ranil Jayawardena and Anoop Misra. "Balanced diet is a major casualty in COVID-19". In: *Diabetes & Metabolic Syndrome* (2020).
- [189] Robert W Jeffery, Rena R Wing, Carolyn Thorson, Lisa R Burton, Cheryl Raether, Jean Harvey, and Monica Mullen. "Strengthening behavioral interventions for weight loss: a randomized trial of food provision and monetary incentives." In: *Journal of Consulting and Clinical Psychology* 61.6 (1993).
- [190] Peng Jia, Lei Zhang, Wanqi Yu, Bin Yu, Meijing Liu, Dong Zhang, and Shujuan Yang. "Impact of COVID-19 lockdown on activity patterns and weight status among youths in China: the COVID-19 Impact on Lifestyle Change Survey (COINLICS)". In: *International Journal of Obesity* (2020).
- [191] Eric S Johnson, Barbara A Bartman, Becky A Briesacher, Neil S Fleming, Tobias Gerhard, Cynthia J Kornegay, Parivash Nourjah, Brian Sauer, Glen T Schumock, Art Sedrakyan, et al. "The incident user design in comparative effectiveness research". In: *Pharmacoepidemiology and Drug Safety* 22.1 (2013).
- [192] Thomas F Johnson, Lisbeth A Hordley, Matthew P Greenwell, and Luke C Evans. "Effect of park use and landscape structure on COVID-19 transmission rates". In: *Science of The Total Environment* (2021).
- [193] James Jordon, Jinsung Yoon, and Mihaela Van Der Schaar. "PATE-GAN: Generating synthetic data with differential privacy guarantees". In: *International conference on learning representations*. 2018.
- [194] Anna Josephson, Talip Kilic, and Jeffrey D Michler. "Socioeconomic impacts of COVID-19 in low-income countries". In: *Nature Human Behaviour* (2021).
- [195] Dan Jurafsky. *The language of food: A linguist reads the menu*. 2014.
- [196] Dan Jurafsky, Victor Chahuneau, Bryan R Routledge, and Noah A Smith. "Narrative framing of consumer sentiment in online restaurant reviews". In: *First Monday* (2014).
- [197] O Kakaa, F Bert, C Botezatu, MR Gualano, and R Siliquini. "How we make choices about food? Analysis of factors influencing food expenditure in Northern Italy". In: *European Journal of Public Health* 28 (Nov. 2018).
- [198] Valentin Kassarnig, Enys Mones, Andreas Bjerre-Nielsen, Piotr Sapiezynski, David Dreyer Lassen, and Sune Lehmann. "Academic performance and behavioral patterns". In: *EPJ Data Science* 7 (2018).

-
- [199] Michail Katsoulis, Laura Pasea, Alvina G Lai, Richard JB Dobson, Spiros Denaxas, Harry Hemingway, and Amitava Banerjee. "Obesity during the COVID-19 pandemic: both cause of high risk and potential effect of lockdown? A population-based electronic health record study". In: *Public health* 191 (2021).
 - [200] Noriaki Kawamae. "Serendipitous recommendations via innovators". In: *Proc. of the 33rd International ACM Conference on Research and Development in Information Retrieval (SIGIR)*. 2010.
 - [201] Katherine A Keith, David Jensen, and Brendan O'Connor. "Text and causal inference: A review of using text to remove confounding from causal estimates". In: *arXiv preprint arXiv:2005.00649* (2020).
 - [202] Jiin Kim, Zara Ahmad, Yena Lee, Flora Nasri, Hartej Gill, Roger McIntyre, Lee Phan, and Leanna Lui. "Systematic Review of the Validity of Screening Depression through Facebook, Twitter, and Instagram". In: *Journal of Affective Disorders* (2021).
 - [203] Emre Kıcıman, Scott Counts, and Melissa Gasser. "Using longitudinal social media analysis to understand the effects of early college alcohol use". In: *Proc. of the 12th International AAAI Conference on Web and Social Media (ICWSM)* (2018).
 - [204] Gueorgi Kossinets and Duncan J Watts. "Empirical analysis of an evolving social network". In: *Science* 311.5757 (2006).
 - [205] Ramez Kouzy et al. "Coronavirus goes viral: quantifying the COVID-19 misinformation epidemic on Twitter". In: *Cureus* 12.3 (2020).
 - [206] Marina Krakovsky. "Solving for why". In: *Communications of the ACM* 65.2 (2022).
 - [207] Juhi Kulshrestha, Motahhare Eslami, Johnnatan Messias, Muhammad Bilal Zafar, Saptarshi Ghosh, Krishna P Gummadi, and Karrie Karahalios. "Search bias quantification: Investigating political bias in social media and web search". In: *Information Retrieval Journal* 22.1 (2019).
 - [208] David Laborde, Will Martin, Johan Swinnen, and Rob Vos. "COVID-19 risks to global food security". In: *Science* 369.6503 (2020).
 - [209] Carl K Lachat, Roosmarijn Verstraeten, Bruno De Meulenaer, Joris Menten, Lieven F Huybregts, John Van Camp, Dominique Roberfroid, and Patrick W Kolsteren. "Availability of free fruits and vegetables at canteen lunch improves lunch and daily nutritional profiles: a randomised controlled trial". In: *British journal of nutrition* 102.7 (2009).
 - [210] Peter A Lachenbruch and Cornelius J Lynch. "Assessing screening tests: extensions of McNemar's test". In: *Statistics in Medicine* 17.19 (1998).
 - [211] Laura Laguna, Susana Fiszman, Patricia Puerta, C Chaya, and Amparo Tárrega. "The impact of COVID-19 lockdown on food priorities. Results from a preliminary study using social media and an online survey with Spanish consumers". In: *Food quality and preference* 86 (2020).

Bibliography

- [212] Martin Lajous, Anne Bijon, Guy Fagherazzi, Emilie Rossignol, Marie-Christine Boutron-Ruault, and Françoise Clavel-Chapelon. “Processed and unprocessed red meat consumption and hypertension in women”. In: *The American journal of clinical nutrition* 100.3 (2014).
- [213] Camille Lassale, G David Batty, Amaria Baghdadli, Felice Jacka, Almudena Sánchez-Villegas, Mika Kivimäki, and Tasnime Akbaraly. “Healthy dietary indices and risk of depressive outcomes: a systematic review and meta-analysis of observational studies”. In: *Molecular psychiatry* 24.7 (2019).
- [214] David Lazer, Eszter Hargittai, Deen Freelon, Sandra Gonzalez-Bailon, Kevin Munger, Katherine Ognyanova, and Jason Radford. “Meaningful measures of human society in the twenty-first century”. In: *Nature* 595.7866 (2021).
- [215] David Lazer, Ryan Kennedy, Gary King, and Alessandro Vespignani. “The parable of Google Flu: traps in big data analysis”. In: *Science* 343.6176 (2014).
- [216] Christopher A Le Dantec, Erika Shehan Poole, and Susan P Wyche. “Values as lived experience: Evolving value sensitive design in support of value discovery”. In: *Proceedings of the SIGCHI conference on human factors in computing systems (CHI)*. 2009.
- [217] Michael Lechner. “The estimation of causal effects by difference-in-difference methods”. In: *Foundations and Trends in Econometrics* 4.3 (2011).
- [218] Douglas E Levy, Mark C Pachucki, A James O’Malley, Bianca Porneala, Awesta Yaqubi, and Anne N Thorndike. “Social connections and the healthfulness of food choices in an employee population”. In: *Nature Human Behaviour* 5.10 (2021), pp. 1349–1357.
- [219] Yu-Hsuan Lin, Chun-Hao Liu, and Yu-Chuan Chiu. “Google searches for the keywords of “wash hands” predict the speed of national spread of COVID-19 outbreak among 21 countries”. In: *Brain, Behavior, and Immunity* (2020).
- [220] Sandrine Lioret, Mathilde Touvier, Morgan Balin, Inge Huybrechts, Carine Dubuisson, Ariane Dufour, Mélanie Bertin, Bernard Maire, and Lionel Lafay. “Characteristics of energy under-reporting in children and adolescents”. In: *British journal of nutrition* 105.11 (2011).
- [221] Dimitra Liotsiou, Luc Moreau, and Susan Halford. “Social influence: From contagion to a richer causal understanding”. In: *International Conference on Social Informatics*. 2016, pp. 116–132.
- [222] Adam Liverpool et al. *Covid-19 news: Coronavirus restrictions to ease slightly in England*. <https://www.newscientist.com/article/2244516-uk-plans-to-further-ease-lockdown-as-new-case-rate-remains-high/>. Accessed: 2020-06-04. 2020.
- [223] Yonglong Lu, Nebojsa Nakicenovic, Martin Visbeck, and Anne-Sophie Stevance. “Policy: Five priorities for the UN sustainable development goals”. In: *Nature* 520.7548 (2015).
- [224] Jennifer L Lund, David B Richardson, and Til Stürmer. “The active comparator, new user study design in pharmacoepidemiology: historical foundations and contemporary application”. In: *Current Epidemiology Reports* 2.4 (2015).

- [225] Yuhan Luo, Peiyi Liu, and Eun Kyoung Choe. “Co-Designing food trackers with dietitians: Identifying design opportunities for food tracker customization”. In: *Proceedings of the 2019 Conference on Human Factors in Computing Systems (CHI)*. 2019.
- [226] JB Lynch and EY Hsiao. “Microbiomes as sources of emergent host phenotypes”. In: *Science* 365.6460 (2019).
- [227] Anmol Madan, Sai T Moturu, David Lazer, and Alex Sandy Pentland. “Social sensing: Obesity, unhealthy eating and exercise in face-to-face networks”. In: *Wireless Health 2010*. 2010.
- [228] Silvia Malaguzzi. *Food and feasting in art*. 2008.
- [229] Momin M Malik, Hemank Lamba, Constantine Nakos, and Jürgen Pfeffer. “Population bias in geotagged tweets”. In: *Proc. of the 9th International AAAI Conference on Web and Social Media (ICWSM)*. 2015.
- [230] Momin M Malik and Jürgen Pfeffer. “Identifying platform effects in social media data”. In: *Proc. of the 10th International AAAI Conference on Web and Social Media (ICWSM)*. 2016.
- [231] Kathryn E Mansfield, Rohini Mathur, John Tazare, Alasdair D Henderson, Amy R Mulick, Helena Carreira, Anthony A Matthews, Patrick Bidulka, Alicia Gayle, Harriet Forbes, et al. “Indirect acute effects of the COVID-19 pandemic on physical and mental health in the UK: a population-based study”. In: *The Lancet Digital Health* 3.4 (2021).
- [232] Olivia D Martinez, Christina A Roberto, Jane H Kim, Marlene B Schwartz, and Kelly D Brownell. “A survey of undergraduate student perceptions and use of nutrition information labels in a university dining hall”. In: *Health Education Journal* 72.3 (2013).
- [233] Naiara Martinez-Perez, Liv Elin Torheim, Nerea Castro-Diaz, and Marta Arroyo-Izaga. “On-campus food environment, purchase behaviours, preferences and opinions in a Norwegian university community”. In: *Public Health Nutrition* 25.6 (2022).
- [234] Anna V Mattioli, Susanna Sciomer, Camilla Cocchi, Silvia Maffei, and Sabina Gallina. “Quarantine during COVID-19 outbreak: Changes in diet and physical activity increase the risk of cardiovascular disease”. In: *Nutrition, Metabolism and Cardiovascular Diseases* 30.9 (2020).
- [235] Lucas Maystre and Matthias Grossglauser. “Fast and accurate inference of Plackett–Luce models”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. 2015.
- [236] Jim McCambridge, John Witton, and Diana R Elbourne. “Systematic review of the Hawthorne effect: New concepts are needed to study research participation effects”. In: *Journal of clinical epidemiology* 67.3 (2014), pp. 267–277.
- [237] Brent McFerran, Darren W. Dahl, Gavan J. Fitzsimons, and Andrea C. Morales. “I’ll have what she’s having: Effects of social influence and body type on the food choices of others”. In: *Journal of Consumer Research* 36.6 (2010).

Bibliography

- [238] David J. McIver and John S. Brownstein. “Wikipedia Usage Estimates Prevalence of Influenza-Like Illness in the United States in Near Real-Time”. In: *PLOS Computational Biology*. 2014.
- [239] William H McNeill. “How the potato changed the world’s history”. In: *Social Research* (1999).
- [240] Dave Mead, Karen Ransom, Stephen B Reed, and Scott Sager. “The impact of the COVID-19 pandemic on food price indexes and data collection”. In: *Monthly Lab. Rev.* 143 (2020).
- [241] Yelena Mejova, Sofiane Abbar, and Hamed Haddadi. “Fetishizing food in digital age: #foodporn around the world”. In: *Proc. of the 10th International AAAI Conference on Web and Social Media (ICWSM)*. 2016.
- [242] Yelena Mejova, Hamed Haddadi, Sofiane Abbar, Azadeh Ghahghaei, and Ingmar Weber. “Dietary habits of an expat nation: Case of Qatar”. In: *2015 International Conference on Healthcare Informatics*. 2015.
- [243] Yelena Mejova, Hamed Haddadi, Anastasios Noulas, and Ingmar Weber. “# foodporn: Obesity patterns in culinary interactions”. In: *Proc. of the 5th International Conference on Digital Health 2015*. 2015.
- [244] Yelena Mejova, Ingmar Weber, and Michael W Macy. *Twitter: A digital socioscope*. 2015.
- [245] Rebecca Mete, Alison Shield, Kristen Murray, Rachel Bacon, and Jane Kellett. “What is healthy eating? A qualitative exploration”. In: *Public Health Nutrition* 22.13 (2019).
- [246] S Meyre. *Agriculture et alimentation. Statistique de poche 2017*. 2017.
- [247] Katherine L Milkman, Dena Gromet, Hung Ho, Joseph S Kay, Timothy W Lee, Pepi Pandiloski, Yeji Park, Aneesh Rai, Max Bazerman, John Beshears, et al. “Megastudies improve the impact of applied behavioural science”. In: *Nature* 600.7889 (2021).
- [248] Sharada Prasanna Mohanty, Gaurav Singhal, Eric Antoine Scuccimarra, Djilani Kebaili, Harris Héritier, Victor Boulanger, and Marcel Salathé. “The Food Recognition Benchmark: Using Deep Learning to Recognize Food in Images”. In: *Frontiers in Nutrition* 9 (2022).
- [249] Esther Molina-Montes, Irina Uzhova, Vito Verardo, Reyes Artacho, Belén Garcia-Villanova, Eduardo Jesús Guerra-Hernández, Maria Kapsokafalou, Olga Malisova, Antonis Vlassopoulos, Alexandra Katidi, et al. “Impact of COVID-19 confinement on eating behaviours across 16 European countries: The COVIDiet cross-national study”. In: *Food Quality and Preference* 93 (2021).
- [250] Saar Mollen, Rajiv N Rimal, Robert AC Ruiter, and Gerjo Kok. “Healthy and unhealthy social norms and food selection. Findings from a field-experiment”. In: *Appetite* 65 (2013).
- [251] Anders Mollgaard, Sune Lehmann, and Joachim Mathiesen. “Correlations between human mobility and social interaction reveal general activity patterns”. In: *PloS one* 12.12 (2017).

- [252] Gil Morrot, Frédéric Brochet, and Denis Dubourdieu. "The color of odors". In: *Brain and language* 79.2 (2001).
- [253] Mehrab Bin Morshed, Samruddhi Shreeram Kulkarni, Koustuv Saha, Richard Li, Leah G Roper, Lama Nachman, Hong Lu, Lucia Mirabella, Sanjeev Srivastava, Kaya de Barbaro, et al. "Food, Mood, Context: Examining College Students' Eating Context and Mental Well-Being". In: *ACM Transactions on Computing for Healthcare* (2022).
- [254] Fred Morstatter, Jürgen Pfeffer, and Huan Liu. "When is it biased? Assessing the representativeness of Twitter's streaming API". In: *Proc. of the 23rd International Conference on World Wide Web (TheWebConf)*. 2014.
- [255] Andrew B Moynihan, Wijnand AP Van Tilburg, Eric R Igou, Arnaud Wisman, Alan E Donnelly, and Jessie B Mulcaire. "Eaten up by boredom: Consuming food to escape awareness of the bored self". In: *Frontiers in psychology* 6 (2015).
- [256] AE Munt, SR Partridge, and M Allman-Farinelli. "The barriers and enablers of healthy eating among young adults: A missing piece of the obesity puzzle: A scoping review". In: *Obesity Reviews* 18.1 (2017).
- [257] Shu Naritomi and Keiji Yanai. "CalorieCaptorGlass: Food calorie estimation based on actual size using hololens and deep learning". In: *IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW)*. 2020.
- [258] Aydin Nazmi, Suzanna Martinez, Ajani Byrd, Derrick Robinson, Stephanie Bianco, Jennifer Maguire, Rashida M Crutchfield, Kelly Condrón, and Lorrene Ritchie. "A systematic review of food insecurity among US students in higher education". In: *Journal of Hunger & Environmental Nutrition* 14.5 (2019).
- [259] Cliona Ni Mhurchu, Louise M Aston, and Susan A Jebb. "Effects of worksite health promotion interventions on employee diets: a systematic review". In: *BMC public health* 10.1 (2010).
- [260] Erik C Nook and Jamil Zaki. "Social norms shift behavioral and neural responses to foods". In: *Journal of Cognitive Neuroscience* 27.7 (2015).
- [261] Ferda Ofli, Yusuf Aytar, Ingmar Weber, Raggi Al Hammouri, and Antonio Torralba. "Is saki# delicious?: The food perception gap on Instagram and its relation to health". In: *Proc. of the 26th International Conference on World Wide Web (TheWebConf)*. 2017.
- [262] Kaimu Okamoto and Keiji Yanai. "UEC-FoodPIX Complete: A Large-scale Food Image Segmentation Dataset". In: *International Conference on Pattern Recognition*. 2021.
- [263] David Olganier and Trine H Mogensen. "The Covid-19 pandemic in Denmark: Big lessons from a small country". In: *Cytokine & growth factor reviews* 53 (2020).
- [264] Alexandra Olteanu, Carlos Castillo, Fernando Diaz, and Emre Kıcıman. "Social data: Biases, methodological pitfalls, and ethical boundaries". In: *Frontiers in Big Data* 2 (2019).
- [265] World Health Organization et al. *Healthy diet*. Tech. rep. World Health Organization. Regional Office for the Eastern Mediterranean, 2019.

Bibliography

- [266] Saskia Osendarp, Jonathan Kweku Akuoku, Robert E Black, Derek Headey, Marie Ruel, Nick Scott, Meera Shekar, Neff Walker, Augustin Flory, Lawrence Haddad, et al. "The COVID-19 crisis will exacerbate maternal and child undernutrition and child mortality in low-and middle-income countries". In: *Nature Food* (2021).
- [267] Mark A Pachucki, Paul F Jacques, and Nicholas A Christakis. "Social network concordance in food choice among spouses, friends, and siblings". In: *American Journal of Public Health* 101.11 (2011).
- [268] Jessica A. Pater, Oliver L. Haimson, Nazanin Andalibi, and Elizabeth D. Mynatt. "'Hunger hurts but starving works': Characterizing the presentation of eating disorders online". In: *Proc. of the 2016 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2016.
- [269] Heather Patrick and Theresa A Nicklas. "A review of family and social determinants of children's eating patterns and diet quality". In: *Journal of the American College of Nutrition* 24.2 (2005).
- [270] Devon C Payne-Sturges, Allison Tjaden, Kimberly M Caldeira, Kathryn B Vincent, and Amelia M Arria. "Student hunger on campus: Food insecurity among college students and implications for academic institutions". In: *American Journal of Health Promotion* 32.2 (2018).
- [271] Judea Pearl and Dana Mackenzie. *The book of why: The new science of cause and effect*. 2018.
- [272] Susanne Pedersen, Alice Grønhoj, and John Thøgersen. "Following family or friends. Social norms in adolescent healthy eating". In: *Appetite* 86 (2015).
- [273] Max Pellert, Hannah Metzler, Michael Matzenberger, and David Garcia. "Validating daily social media macroscopes of emotions". In: *Scientific Reports* 12.1 (2022).
- [274] Maria Perez-Ortiz and Rafal K Mantiuk. "A practical guide and software for analysing pairwise comparison experiments". In: *arXiv preprint arXiv:1712.03686* (2017).
- [275] Thanh-Trung Phan and Daniel Gatica-Perez. "Healthy# fondue# dinner: analysis and inference of food and drink consumption patterns on instagram". In: *Proc. of the 16th International Conference on Mobile and Ubiquitous Multimedia*. 2017.
- [276] Angelo Pietrobelli, Luca Pecoraro, Alessandro Ferruzzi, Moonseong Heo, Myles Faith, Thomas Zoller, Franco Antoniazzi, Giorgio Piacentini, S Nicole Fearnbach, and Steven B Heymsfield. "Effects of COVID-19 lockdown on lifestyle behaviors in children with obesity living in Verona, Italy: a longitudinal study". In: *Obesity* 28.8 (2020).
- [277] Haoyue Ping, Julia Stoyanovich, and Bill Howe. "Datasynthesizer: Privacy-preserving synthetic datasets". In: *Proceedings of the 29th International Conference on Scientific and Statistical Database Management*. 2017.
- [278] J. Poore and T. Nemecek. "Reducing food's environmental impacts through producers and consumers". In: *Science* (2018).

-
- [279] Joseph Poore and Thomas Nemecek. "Reducing food's environmental impacts through producers and consumers". In: *Science* 360.6392 (2018).
- [280] Jacob Poushter et al. "Smartphone ownership and internet usage continues to climb in emerging economies". In: *Pew research center* 22.1 (2016).
- [281] Tobias Preis, Helen Susannah Moat, and H Eugene Stanley. "Quantifying trading behavior in financial markets using Google Trends". In: *Scientific Reports* 3.1 (2013).
- [282] Reid Pryzant, Youngjoo Chung, and Dan Jurafsky. "Predicting sales from the language of product descriptions". In: *eCOM@ SIGIR*. 2017.
- [283] Giulia Pullano, Eugenio Valdano, Nicola Scarpa, Stefania Rubrichi, and Vittoria Colizza. "Evaluating the effect of demographic factors, socioeconomic factors, and risk aversion on mobility during the COVID-19 epidemic in France under lockdown: a population-based study". In: *The Lancet Digital Health* 2.12 (2020).
- [284] Rajagopal Raghunathan, Rebecca Walker Naylor, and Wayne D Hoyer. "The unhealthy=tasty intuition and its effects on taste inferences, enjoyment, and choice of food products". In: *Journal of Marketing* 70.4 (2006).
- [285] Rebecca K Ratner and Barbara E Kahn. "The impact of private versus public consumption on variety-seeking behavior". In: *Journal of Consumer Research* 29.2 (2002), pp. 246–257.
- [286] John J Reilly and Joanna Kelly. "Long-term impact of overweight and obesity in childhood and adolescence on morbidity and premature mortality in adulthood: systematic review". In: *International journal of obesity* 35.7 (2011).
- [287] Pengyu Patrick Ren, Zhenyu Cheryl Qian, and Jung Joo Sohn. "Learn to Cook for Yourself: Employing Gamification in a Recipe App Design to Promote a Healthy Living Experience to Young Generation". In: *International Conference on Human-Computer Interaction*. Springer. 2020.
- [288] *Restauration collective: En Suisse, près d'un million de personnes mangent chaque jour dans un établissement de restauration collective*. <https://www.blv.admin.ch/blv/fr/home/lebensmittel-und-ernaehrung/ernaehrung/gemeinschaftsgastronomie.html>. Accessed: 2021-03-01.
- [289] Peter R Reuter, Bridget L Forster, and Sierra R Brister. "The influence of eating habits on the academic performance of university students". In: *Journal of American College Health* 69.8 (2021).
- [290] Christian Reynolds, Liam Goucher, Tom Quested, Sarah Bromley, Sam Gillick, Victoria K Wells, David Evans, Lenny Koh, Annika Carlsson Kanyama, Cecilia Katzeff, et al. "Consumption-stage food waste reduction interventions—What works and how to design better interventions". In: *Food Policy* 83 (2019).

Bibliography

- [291] Manoel Horta Ribeiro, Kristina Gligorić, Maxime Peyrard, Florian Lemmerich, Markus Strohmaier, and Robert West. “Sudden attention shifts on Wikipedia during the COVID-19 crisis”. In: *Proc. of the 15th International AAAI Conference on Web and Social Media (ICWSM)*. 2021.
- [292] Hannah Ritchie and Max Roser. “Land Use”. In: *Our World in Data* (2013).
- [293] Eric Robinson, Jackie Blissett, and Suzanne Higgs. “Social influences on eating: implications for nutritional interventions”. In: *Nutrition Research Reviews* 26.2 (2013).
- [294] Eric Robinson and Suzanne Higgs. “Food choices in the presence of “healthy” and “unhealthy” eating partners”. In: *British Journal of Nutrition* 109.4 (2013).
- [295] Eric Robinson, Jason Thomas, Paul Aveyard, and Suzanne Higgs. “What everyone else is eating: A systematic review and meta-analysis of the effect of informational eating norms on eating behavior”. In: *Journal of the Academy of Nutrition and Dietetics* 114.3 (2014).
- [296] Rachel F Rodgers, Caterina Lombardo, Silvia Cerolini, Debra L Franko, Mika Omori, Matthew Fuller-Tyszkiewicz, Jake Linardon, Philippe Courtet, and Sebastien Guillaume. “The impact of the COVID-19 pandemic on eating disorder risk and symptoms”. In: *International Journal of Eating Disorders* 53.7 (2020).
- [297] Markus Rokicki, Christoph Trattner, and Eelco Herder. “The impact of recipe features, social cues and demographics on estimating the healthiness of online recipes”. In: *Proc. of the 12th International AAAI Conference on Web and Social Media (ICWSM)*. 2018.
- [298] E Romeo-Arroyo, M Mora, and L Vázquez-Araújo. “Consumer behavior in confinement times: Food choice and cooking attitudes in Spain”. In: *International Journal of Gastronomy and Food Science* 21 (2020).
- [299] Paul Rosenbaum. *Observation and experiment*. 2018.
- [300] Paul Rosenbaum and Donald B. Rubin. “The central role of the propensity score in observational studies for causal effects”. In: *Biometrika* 70.1 (Apr. 1983).
- [301] Paul Rosenbaum and Jeffrey H Silber. “Amplification of sensitivity analysis in matched observational studies”. In: *Journal of the American Statistical Association* 104.488 (2009).
- [302] Rajshri Roy, Bridget Kelly, Anna Rangan, and Margaret Allman-Farinelli. “Food environment interventions to improve the dietary behavior of young adults in tertiary education settings: a systematic literature review”. In: *Journal of the Academy of Nutrition and Dietetics* 115.10 (2015).
- [303] Rajshri Roy, Anna Rangan, Lana Hebden, Jimmy Chun Yu Louie, Lie Ming Tang, Judy Kay, and Margaret Allman-Farinelli. “Dietary contribution of foods and beverages sold within a university campus and its effect on diet quality of young adults”. In: *Nutrition* 34 (2017).

-
- [304] Rajshri Roy, Danielle Soo, Denise Conroy, Clare R Wall, and Boyd Swinburn. "Exploring university food environment and on-campus food purchasing behaviors, preferences, and opinions". In: *Journal of nutrition education and behavior* 51.7 (2019).
- [305] Donald B Rubin. "Causal inference using potential outcomes: Design, modeling, decisions". In: *Journal of the American Statistical Association* 100.469 (2005).
- [306] Derek Ruths and Jürgen Pfeffer. "Social media for large studies of behavior". In: *Science* 346.6213 (2014).
- [307] Benjamin J. Ryan, Damon Coppola, Deon V. Canyon, Mark Brickhouse, and Raymond Swienton. "COVID-19 Community Stabilization and Sustainability Framework". In: *Disaster Medicine and Public Health Preparedness* (2020).
- [308] Mohamed Sabt, Mohammed Achemlal, and Abdelmadjid Bouabdallah. "Trusted execution environment: what it is, and what it is not". In: *2015 IEEE Trustcom/BigDataSE/ISPA*. Vol. 1. IEEE. 2015.
- [309] Adam Sadilek, Stephanie Caty, Lauren DiPrete, Raed Mansour, Tom Schenk, Mark Bergtholdt, Ashish Jha, Prem Ramaswami, and Evgeniy Gabrilovich. "Machine-learned epidemiology: real-time detection of foodborne illness at scale". In: *npj Digital Medicine* 1.1 (2018).
- [310] Koustuv Saha, Jordyn Seybolt, Stephen M Mattingly, Talayeh Aledavood, Chaitanya Konjeti, Gonzalo J Martinez, Ted Grover, Gloria Mark, and Munmun De Choudhury. "What Life Events are Disclosed on Social Media, How, When, and By Whom?" In: *Proc. of the 2021 Conference on Human Factors in Computing Systems (CHI)*. 2021.
- [311] Koustuv Saha, John Torous, Eric D Caine, and Munmun De Choudhury. "Psychosocial Effects of the COVID-19 Pandemic: Large-scale Quasi-Experimental Study on Social Media". In: *J Med Internet Res* 22.11 (Nov. 2020).
- [312] Koustuv Saha, Asra Yousuf, Ryan L Boyd, James W Pennebaker, and Munmun De Choudhury. "Social media discussions predict mental health consultations on college campuses". In: *Scientific Reports* 12.1 (2022).
- [313] Doyen Sahoo, Wang Hao, Shu Ke, Wu Xiongwei, Hung Le, Palakorn Achananuparp, Ee-Peng Lim, and Steven CH Hoi. "FoodAI: food image recognition via deep learning for smart food logging". In: *Proc. of the 25th ACM International Conference on Knowledge Discovery & Data Mining (KDD)*. 2019.
- [314] Sina Sajadmanesh, Sina Jafarzadeh, Seyed Ali Ossia, Hamid R Rabiee, Hamed Haddadi, Yelena Mejova, Mirco Musolesi, Emiliano De Cristofaro, and Gianluca Stringhini. "Kissing cuisines: Exploring worldwide culinary habits on the web". In: *Proc. of the 26th International Conference on World Wide Web (TheWebConf)*. 2017.
- [315] Matthew J Salganik. *Bit by bit: Social research in the digital age*. 2019.

Bibliography

- [316] Amaia Salvador, Nicholas Hynes, Yusuf Aytar, Javier Marin, Ferda Ofli, Ingmar Weber, and Antonio Torralba. “Learning cross-modal embeddings for cooking recipes and food images”. In: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017.
- [317] Sarah-Jeanne Salvy, Kayla De La Haye, Julie C Bowker, and Roel CJ Hermans. “Influence of peers and friends on children’s and adolescents’ eating and activity behaviors”. In: *Physiology & Behavior* 106.3 (2012).
- [318] Sarah-Jeanne Salvy, Elizabeth Kieffer, and Leonard H Epstein. “Effects of social context on overweight and normal-weight children’s food selection”. In: *Eating Behaviors* 9.2 (2008).
- [319] Jose Ramon Saura, Ana Reyes-Menendez, and Stephen B Thomas. “Gaining a deeper understanding of nutrition using social networks and user-generated content”. In: *Internet Interventions* 20 (2020).
- [320] Federico Scarmozzino and Francesco Visioli. “Covid-19 and the subsequent lockdown modified dietary habits of almost half the population in an Italian sample”. In: *Foods* 9.5 (2020).
- [321] Claudia Schmidt, Stephan Goetz, Sarah Rocker, and Zheng Tian. “Google searches reveal changing consumer food sourcing in the COVID-19 pandemic”. In: *Journal of Agriculture, Food Systems, and Community Development* 9.3 (2020).
- [322] Jonathan D Schoenfeld and John PA Ioannidis. “Is everything we eat associated with cancer? A systematic cookbook review”. In: *The American journal of clinical nutrition* 97.1 (2013).
- [323] Jessica Schroeder, Jane Hoffswell, Chia-Fang Chung, James Fogarty, Sean Munson, and Jasmine Zia. “Supporting patient-provider collaboration to identify individual triggers using food and symptom journals”. In: *Proc. of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)*. 2017.
- [324] Yasaman S Sefidgar, Woosuk Seo, Kevin S Kuehn, Tim Althoff, Anne Browning, Eve Riskin, Paula S Nurius, Anind K Dey, and Jennifer Mankoff. “Passively-sensed behavioral correlates of discrimination events in college students”. In: *Proc. of the 2019 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW)* (2019).
- [325] Sara M. Seifert, Judith L. Schaechter, Eugene R. Hershorin, and Steven E. Lipshultz. “Health Effects of Energy Drinks on Children, Adolescents, and Young Adults”. In: *Pediatrics* 127.3 (2011).
- [326] Indira Sen, Fabian Flöck, and Claudia Wagner. “On the reliability and validity of detecting approval of political actors in tweets”. In: *Proc. of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 2020.

-
- [327] Indira Sen, Fabian Flöck, Katrin Weller, Bernd Weiß, and Claudia Wagner. “A total error framework for digital traces of human behavior on online platforms”. In: *Public Opinion Quarterly* (2021).
- [328] Laila Al-Shaar, Kelsey Vercammen, Chang Lu, Scott Richardson, Martha Tamez, and Josiemer Mattei. “Health effects and public health concerns of energy drink consumption in the United States: a mini-review”. In: *Frontiers in Public Health* (2017).
- [329] Cosma Rohilla Shalizi and Edward McFowland III. “Estimating causal peer influence in homophilous social networks by inferring latent locations”. In: (2016).
- [330] Cosma Rohilla Shalizi and Andrew C Thomas. “Homophily and contagion are generically confounded in observational social network studies”. In: *Sociological Methods & Research* 40.2 (2011).
- [331] Piyush Sharma, Bharadhwaj Sivakumaran, and Roger Marshall. “Impulse buying and variety seeking: A trait-correlates perspective”. In: *Journal of Business research* 63.3 (2010), pp. 276–283.
- [332] Sanket S Sharma and Munmun De Choudhury. “Measuring and characterizing nutritional information of food and ingestion content in Instagram”. In: *Proc. of the 24th International Conference on World Wide Web (TheWebConf)*. 2015.
- [333] Maxine Sharps, Suzanne Higgs, Jackie Blissett, Arie Nouwen, Magdalena Chechlac, Harriet A Allen, and Eric Robinson. “Examining evidence for behavioural mimicry of parental eating by adolescent females. An observational study”. In: *Appetite* 89 (2015), pp. 56–61.
- [334] Richard Shepherd. “Social determinants of food choice”. In: *Proceedings of the Nutrition Society* 58.4 (1999).
- [335] Jagdish Sheth. “Impact of Covid-19 on consumer behavior: Will the old habits return or die?”. In: *Journal of Business Research* 117 (2020).
- [336] Ashutosh Singla, Lin Yuan, and Touradj Ebrahimi. “Food/non-food image classification and food categorization using pre-trained googlenet model”. In: *Proceedings of the 2nd International Workshop on Multimedia Assisted Dietary Management*. 2016.
- [337] Meg Small, Lisa Bailey-Davis, Nicole Morgan, and Jennifer Maggs. “Changes in eating and physical activity behaviors across seven semesters of college: living on or off campus matters”. In: *Health Education & Behavior* 40.4 (2013).
- [338] Giovanni Sogari, Catalina Velez-Argumedo, Miguel I Gómez, and Cristina Mora. “College students and eating habits: A study using an ecological model for healthy behavior”. In: *Nutrients* 10.12 (2018).
- [339] S. Sohrabvandi, A.M. Mortazavian, and K. Rezaei. “Health-Related Aspects of Beer: A Review”. In: *International Journal of Food Properties* 15.2 (2012).
- [340] Barbara C Sorkin, Adam J Kuszak, John S Williamson, D Craig Hopp, and Joseph M Betz. “The challenge of reproducibility and accuracy in nutrition research: resources and pitfalls”. In: *Advances in Nutrition* 7.2 (2016).

Bibliography

- [341] *Sourdough baking sees rise in popularity during COVID-19 pandemic*. <https://www.cbc.ca/news/canada/north/sourdough-popular-covid-19-1.5529649>. Accessed: 2021-03-01.
- [342] Molly Spencer, Cesar Cienfuegos, and Jean-Xavier Guinard. “The Flexitarian Flip™ in university dining venues: Student and adult consumer acceptance of mixed dishes in which animal protein has been partially replaced with plant protein”. In: *Food Quality and Preference* 68 (2018).
- [343] David K Spierer, Nineequa Blanding, and Anthony Santella. “Energy drink consumption and associated health behaviors among university students in an urban setting”. In: *Journal of community health* 39.1 (2014).
- [344] Clay Spinuzzi. “The methodology of participatory design”. In: *Technical communication* 52.2 (2005).
- [345] Seth Stephens-Davidowitz. “The cost of racial animus on a black candidate: Evidence using Google search data”. In: *Journal of Public Economics* 118 (2014).
- [346] Clifford Stevenson, Glenda Doherty, Julie Barnett, Orla T Muldoon, and Karen Trew. “Adolescents’ views of food and eating: Identifying barriers to healthy eating”. In: *Journal of Adolescence* 30.3 (2007).
- [347] Stephanie Stockwell, Mike Trott, Mark Tully, Jae Shin, Yvonne Barnett, Laurie Butler, Daragh McDermott, Felipe Schuch, and Lee Smith. “Changes in physical activity and sedentary behaviours from before to during the COVID-19 pandemic lockdown: a systematic review”. In: *BMJ Open Sport & Exercise Medicine* 7.1 (2021).
- [348] Jina Suh, Eric Horvitz, Ryen W White, and Tim Althoff. “Population-scale study of human needs during the covid-19 pandemic: Analysis and implications”. In: *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM)*. 2021.
- [349] Nicolette J Sullivan and Scott A Huettel. “Healthful choices depend on the latency and rate of information accumulation”. In: *Nature Human Behaviour* 5.12 (2021), pp. 1698–1706.
- [350] V. Das Swain, H. Kwon, B. Saket, M. Bin Morshed, K. Tran, D. Patel, Y. Tian, J. Philipose, Y. Cui, T. Plötz, M. De Choudhury, and G. D. Abowd. *Leveraging WiFi network logs to infer social interactions: A case study of academic performance and student behavior*. 2020. arXiv: 2005.11228.
- [351] Ryan Tam, Barbara Yassa, Helen Parker, Helen O’Connor, and Margaret Allman-Farinelli. “University students’ on-campus food purchasing behaviors, preferences, and opinions on food availability”. In: *Nutrition* 37 (2017).
- [352] Takanao Tanaka and Shohei Okamoto. “Increase in suicide following an initial decline during the COVID-19 pandemic in Japan”. In: *Nature Human behaviour* 5.2 (2021).

- [353] Robin J Tanner, Rosellina Ferraro, Tanya L Chartrand, James R Bettman, and Rick Van Baaren. "Of chameleons and consumption: The impact of mimicry on choice and preferences". In: *Journal of Consumer Research* 34.6 (2008), pp. 754–766.
- [354] Sarah L Taubman, James M Robins, Murray A Mittleman, and Miguel A Hernán. "Intervening on risk factors for coronary heart disease: an application of the parametric g-formula". In: *International Journal of Epidemiology* 38.6 (2009).
- [355] Yla R. Tausczik, Kate Faasse, James W. Pennebaker, and Keith J. Petrie. "Public Anxiety and Information Seeking Following the H1N1 Outbreak: Blogs, Newspaper Articles, and Wikipedia Visits". In: *Health Communication*. 2012.
- [356] Edison Thomaz, Irfan Essa, and Gregory D. Abowd. "A practical approach for recognizing eating moments with wrist-mounted inertial sensing". In: *Proc. of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp)*. 2015.
- [357] Lukar E Thornton, Robert W Jeffery, and David A Crawford. "Barriers to avoiding fast-food consumption in an environment supportive of unhealthy eating". In: *Public Health Nutrition* 16.12 (2013).
- [358] Verena Tiefenbeck. "On the magnitude and persistence of the Hawthorne effect—evidence from four field studies". In: *4th European Conference on Behaviour and Energy Efficiency, Coimbra, Portugal*. 2016, pp. 8–9.
- [359] Michele Tizzoni, André Panisson, Daniela Paolotti, and Ciro Cattuto. "The impact of news exposure on collective attention in the United States during the 2016 Zika epidemic". In: *PLOS Computational Biology*. 2020.
- [360] Lauren N Tobey, Christine Mouzong, Joyce Senior Angulo, Sally Bowman, and Melinda M Manore. "How low-income mothers select and adapt recipes and implications for promoting healthy recipes online". In: *Nutrients* 11.2 (2019).
- [361] Linda Too and Bhishna Bajracharya. "Sustainable campus: engaging the community in sustainability". In: *International Journal of Sustainability in Higher Education* (2015).
- [362] Christoph Trattner, Dominik Moesslang, and David Elsweiler. "On the predictability of the popularity of online recipes". In: *EPJ Data Science* 7.1 (2018).
- [363] Annie Tubadji, Frederic Boy, and Don J Webber. "Narrative economics, public policy and mental health". In: *Covid Economics* (2020).
- [364] Bradley P Turnwald, Kathryn G Anderson, Dan Jurafsky, and Alia J Crum. "Five-star prices, appealing healthy item descriptions? Expensive restaurants' descriptive menu language." In: *Health Psychology* (2020).
- [365] Bradley P Turnwald, Jaclyn D Bertoldo, Margaret A Perry, Peggy Policastro, Maureen Timmons, Christopher Bosso, Priscilla Connors, Robert T Valgenti, Lindsey Pine, and Ghislaine Challamel. "Increasing vegetable intake by emphasizing tasty and enjoyable attributes: A randomized controlled multisite intervention for taste-focused labeling". In: *Psychological Science* 30.11 (2019).

Bibliography

- [366] Bradley P Turnwald, Danielle Z Boles, and Alia J Crum. “Association between indulgent descriptions and vegetable consumption: Twisted carrots and dynamite beets”. In: *JAMA internal medicine* 177.8 (2017).
- [367] Bradley P Turnwald and Alia J Crum. “Smart food policy for healthy food labeling: Leading with taste, not healthiness, to shift consumption and enjoyment of healthy foods”. In: *Preventive medicine* 119 (2019).
- [368] Bradley P Turnwald, Isaac J Handley-Miner, Natalie A Samuels, Hazel R Markus, and Alia J Crum. “Nutritional Analysis of Foods and Beverages Depicted in Top-Grossing US Movies, 1994-2018”. In: *JAMA Internal Medicine* 181.1 (2021).
- [369] Bradley P Turnwald, Dan Jurafsky, Alana Conner, and Alia J Crum. “Reading between the menu lines: Are restaurants’ descriptions of “healthy” foods unappealing?” In: *Health Psychology* 36.11 (2017).
- [370] Lenny R Vartanian, Samantha Spanos, C Peter Herman, and Janet Polivy. “Modeling of food intake: a meta-analytic review”. In: *Social Influence* 10.3 (2015).
- [371] Rachel A Vella-Zarb and Frank J Elgar. “The ‘freshman 5’: a meta-analysis of weight gain in the freshman year of college”. In: *Journal of American College Health* 58.2 (2009).
- [372] Kirill Veselkov, Guadalupe Gonzalez, Shahad Aljifri, Dieter Galea, Reza Mirnezami, Jozef Youssef, Michael Bronstein, and Ivan Laponogov. “HyperFoods: Machine intelligent mapping of cancer-beating molecules in foods”. In: *Scientific Reports* 9.1 (2019).
- [373] Joseph Viana, Stephanie A Leonard, Bob Kitay, Daryl Ansel, Peter Angelis, and Wendelin Slusser. “Healthier vending machines in a university setting: Effective and financially sustainable”. In: *Appetite* 121. February (2018).
- [374] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. “Data leverage: A framework for empowering the public in its relationship with technology companies”. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 2021.
- [375] Christina Vogel, Sarah Crozier, Daniel Penn-Newman, Kylie Ball, Graham Moon, Joanne Lord, Cyrus Cooper, and Janis Baird. “Altering product placement to create a healthier layout in supermarkets: Outcomes on store sales, customer purchasing, and diet in a prospective matched controlled cluster study”. In: *PLOS medicine* 18.9 (2021).
- [376] Simeon Vosen and Torsten Schmidt. “Forecasting private consumption: survey-based indicators vs. Google trends”. In: *Journal of Forecasting* 30.6 (2011).
- [377] Denny Vrandečić and Markus Krötzsch. “Wikidata: a free collaborative knowledgebase”. In: *Communications of the ACM* 57.10 (2014).
- [378] Claudia Wagner, Philipp Singer, and Markus Strohmaier. “Spatial and temporal patterns of online food preferences”. In: *Proc. of the 23rd International Conference on World Wide Web (TheWebConf)*. 2014.
- [379] Claudia Wagner, Philipp Singer, and Markus Strohmaier. “The nature and evolution of online food preferences”. In: *EPJ Data Science* 3 (2014).

-
- [380] Claudia Wagner, Markus Strohmaier, Alexandra Olteanu, Emre Kıcıman, Noshir Contractor, and Tina Eliassi-Rad. “Measuring algorithmically infused societies”. In: *Nature* 595.7866 (2021).
- [381] Paul Walter, Esther Infanger, and Pascale Mühlemann. “Food pyramid of the Swiss Society for Nutrition”. In: *Annals of Nutrition and Metabolism* 51.2 (2007).
- [382] Erpeng Wang, Ning An, Zhifeng Gao, Emmanuel Kiprop, and Xianhui Geng. “Consumer food stockpiling behavior and willingness to pay for food reserves in COVID-19”. In: *Food Security* 12.4 (2020).
- [383] Rui Wang, Fanglin Chen, Zhenyu Chen, Tianxing Li, Gabriella Harari, Stefanie Tignor, Xia Zhou, Dror Ben-Zeev, and Andrew T Campbell. “StudentLife: assessing mental health, academic performance and behavioral trends of college students using smartphones”. In: *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing*. 2014.
- [384] Rui Wang, Gabriella Harari, Peilin Hao, Xia Zhou, and Andrew T Campbell. “SmartGPA: how smartphones can assess and predict academic performance of college students”. In: *Proceedings of the 2015 ACM international joint conference on pervasive and ubiquitous computing*. 2015.
- [385] Wei Wang, David Rothschild, Sharad Goel, and Andrew Gelman. “Forecasting elections with non-representative polls”. In: *International Journal of Forecasting* 31.3 (2015).
- [386] Brian Wansink, Ying Cao, Prerna Saini, Mitsuru Shimizu, and David R Just. “College cafeteria snack food purchases become less healthy with each passing week of the semester”. In: *Public health nutrition* 16.7 (2013).
- [387] Brian Wansink, Anupama Mukund, and Andrew Weislogel. “Food art does not reflect reality: A quantitative content analysis of meals in popular paintings”. In: *SAGE Open* 6.3 (2016).
- [388] Kristi Weber, Mary Story, and Lisa Harnack. “Internet food marketing strategies aimed at children and adolescents: A content analysis of food and beverage brand web sites”. In: *Journal of the American Dietetic Association* 9 (2006).
- [389] Galen Weld, Peter West, Maria Glenski, David Arbour, Ryan A Rossi, and Tim Althoff. “Adjusting for confounders with text: Challenges and an empirical evaluation framework for causal inference”. In: *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)*. 2022.
- [390] André O Werneck, Danilo R Silva, Deborah C Malta, Crizian Saar Gomes, Paulo RB Souza-Júnior, Luiz O Azevedo, Marilisa BA Barros, and Célia L Szwarcwald. “Associations of sedentary behaviours and incidence of unhealthy diet during the COVID-19 quarantine in Brazil”. In: *Public health nutrition* 24.3 (2021).
- [391] Robert West. “Calibration of Google Trends Time Series”. In: *Proc. of the 29th ACM International Conference on Information & Knowledge Management (CIKM)*. 2020.

Bibliography

- [392] Robert West, Ryen W. White, and Eric Horvitz. "From Cookies to Cooks: Insights on Dietary Patterns via Analysis of Web Usage Logs". In: *Proc. of the 22nd International Conference on World Wide Web (TheWebConf)*. 2013.
- [393] Megan C Whatnall, Amanda J Patterson, and Melinda J Hutchesson. "Effectiveness of nutrition interventions in vending machines to encourage the purchase and consumption of healthier food and drinks in the university setting: a systematic review". In: *Nutrients* 12.3 (2020).
- [394] *Why Banana Bread Is the Official Comfort Food of the Coronavirus Quarantine*. <https://www.thekitchn.com/banana-bread-official-comfort-food-coronavirus-23021239>. Accessed: 2021-03-01.
- [395] Michael J Widener and Wenwen Li. "Using geolocated Twitter data to monitor the prevalence of healthy and unhealthy food references across the US". In: *Applied Geography* 54 (2014).
- [396] Jennifer L Wilkins, Elizabeth Bowdish, and Jeffery Sobal. "University student perceptions of seasonal and local foods". In: *Journal of nutrition education* 32.5 (2000).
- [397] Walter Willett, Johan Rockström, Brent Loken, Marco Springmann, Tim Lang, Sonja Vermeulen, Tara Garnett, David Tilman, Fabrice DeClerck, Amanda Wood, et al. "Food in the Anthropocene: the EAT-Lancet Commission on healthy diets from sustainable food systems". In: *The Lancet* 393.10170 (2019).
- [398] Courtney K Wilson, June I Matthews, Jamie A Seabrook, and Paula DN Dworatzek. "Self-reported food skills of university students". In: *Appetite* 108 (2017).
- [399] Eveline J Wouters, Junilla K Larsen, Stef P Kremers, Pieter C Dagnelie, and Rinie Geenen. "Peer influence on snacking behavior in adolescence". In: *Appetite* 55.1 (2010).
- [400] WWF. *Bending the Curve: The Restorative Power of Planet-Based Diets*. <https://www.worldwildlife.org/publications/bending-the-curve-the-restorative-power-of-planet-based-diets>. 2020. URL: <http://www.sciencedirect.com/science/article/pii/S0306437907000506>.
- [401] Najat Yahia, Carrie A Brown, Melyssa Rapley, and Mei Chung. "Level of nutrition knowledge and its association with fat consumption among college students". In: *BMC public health* 16.1 (2016).
- [402] Tetsuya Yamamoto, Chigusa Uchiumi, Naho Suzuki, Junichiro Yoshimoto, and Eric Murillo-Rodriguez. "The psychological impact of 'mild lockdown' in Japan during the COVID-19 pandemic: a nationwide survey under a declared state of emergency". In: *International journal of environmental research and public health* (2020).
- [403] Keiji Yanai and Yoshiyuki Kawano. "Twitter food photo mining and analysis for one hundred kinds of foods". In: *Pacific Rim Conference on Multimedia*. 2014.
- [404] Kai-Cheng Yang, Christopher Torres-Lugo, and Filippo Menczer. "Prevalence of low-credibility information on twitter during the covid-19 outbreak". In: *arXiv*. 2020.

-
- [405] Calvin Yip, Michelle Gates, Allison Gates, and Rhona M Hanning. "Peer-led nutrition education programs for school-aged youth: a systematic review of the literature". In: *Health Education Research* 31.1 (2016).
- [406] Kazuki Yoshida, Daniel H Solomon, and Seoyoung C Kim. "Active-comparator design and new-user design in observational studies". In: *Nature Reviews Rheumatology* 11.7 (2015).
- [407] Jessica G Young, Lauren E Cain, James M Robins, Eilis J O'Reilly, and Miguel A Hernán. "Comparative effectiveness of dynamic treatment regimes: an application of the parametric g-formula". In: *Statistics in Biosciences* 3.1 (2011).
- [408] Muhammad Bilal Zafar, Krishna P Gummadi, and Cristian Danescu-Niculescu-Mizil. "Message impartiality in social media discussions". In: *Proc. of the 10th International AAAI Conference on Web and Social Media (ICWSM)*. 2016.
- [409] Ark Fangzhou Zhang, Ruihan Wang, Eric Blohm, Ceren Budak, Lionel P Robert Jr, and Daniel M Romero. "Participation of New Editors after Times of Shock on Wikipedia". In: *Proc. International AAAI Conference on Web and Social Media (ICWSM)*. 2019.
- [410] Jian Zhang, Yumei Zhang, Shanshan Huo, Yidi Ma, Yalei Ke, Peiyu Wang, and Ai Zhao. "Emotional eating in pregnant women during the COVID-19 pandemic and its association with dietary intake and gestational weight gain". In: *Nutrients* 12.8 (2020).
- [411] Jun Zhang, Graham Cormode, Cecilia M Procopiuc, Divesh Srivastava, and Xiaokui Xiao. "Privbayes: Private data release via bayesian networks". In: *ACM Transactions on Database Systems (TODS)* 42.4 (2017).
- [412] George Kingsley Zipf. *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio Books, 2016.
- [413] Roberta Zupo, Fabio Castellana, Rodolfo Sardone, Annamaria Sila, Vito Angelo Giagulli, Vincenzo Triggiani, Raffaele Ivan Cincione, Gianluigi Giannelli, and Giovanni De Pergola. "Preliminary trajectories in dietary behaviors during the COVID-19 pandemic: a public health call to action to face obesity". In: *International Journal of Environmental Research and Public Health* 17.19 (2020).

EDUCATION

Ecole Polytechnique Federale de Lausanne

PhD in Computer and Communication Sciences

Mentor: Prof. Robert West

Sep. 2017 – present

Lausanne, Switzerland

University of Novi Sad

Bachelor with Honors in Electrical and Computer Engineering, GPA: 9.91/10.00

Mentor: Prof. Dejan Vukobratovic

Sep. 2013 – Jul. 2017

Novi Sad, Serbia

PROFESSIONAL EXPERIENCE

EPFL

Doctoral Assistant, mentor: Prof. Robert West

Sep. 2017 - present

Lausanne, Switzerland

Google Health

Research Intern, mentors: Shailesh Bavadekar and Dr. Evgeniy Gabrilovich

Jul. 2019 - Oct. 2019

Palo Alto CA, US

Max Planck Institute for Software Systems

Undergraduate Research Intern, mentor: Prof. Krishna P. Gummadi

Jul. 2016 - Oct. 2016

Saarbrücken, Germany

AWARDS, GRANTS, AND HONORS

Swiss National Science Foundation Postdoc Grant, 2022 (CHF 130,000)

CSCW 2021 Best Paper Honorable Mention Award

ICWSM 2021 Best Reviewer Award

EPFL IC School Dean's Recognition for Outstanding Performance, 2019, 2021

EPFL IC School Best Teaching Assistant Award, 2018

ICWSM Student Grant, 2018, 2021

EPFL IC Doctoral Fellowship for PhD Program in Computer Science, 2017

University of Novi Sad Award for Excellent Performance (awarded to top 2% most successful students), 2015, 2016, 2017, 2018

Dositeja Fellowship for Talented Students, Republic of Serbia Foundation, 2017

PEER-REVIEWED PUBLICATIONS

- [1] Biased Bytes: On the Validity of Estimating Food Consumption from Social Media Images (to appear). **Kristina Gligorić**, Irena Djordjevic and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing (**CSCW**), 2022.
- [2] Anticipated versus Actual Effects of Platform Design Change: A Case Study of Twitter's Character Limit (to appear). **Kristina Gligorić**, Justyna Czeszochowska, Ashton Anderson and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing (**CSCW**), 2022.

- [3] Population-scale dietary interests during the COVID-19 pandemic.
Kristina Gligorić, Arnaud Chiolero, Emre Kiciman, Ryen W White and Robert West. **Nature Communications**, 2022.
- [4] On the Context-Free Ambiguity of Emoji.
Justyna Czeszochowska*, **Kristina Gligorić***, Maxime Peyrard, Yann Mentha, Michał Bień, Andrea Grütter, Anita Auer, Aris Xanthos and Robert West (*equal contributions). International AAAI Conference on Web and Social Media (**ICWSM**), 2022.
- [5] Laughing Heads: Can Transformers Detect What Makes a Sentence Funny?
Maxime Peyrard, Beatriz Borges, **Kristina Gligorić** and Robert West. International Joint Conference on Artificial Intelligence (**IJCAI**), 2021.
- [6] Formation of Social Ties Influences Food Choice: A Campus-wide Longitudinal Study.
Kristina Gligorić, Ryen W White, Emre Kiciman, Eric Horvitz, Arnaud Chiolero and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing (**CSCW**), 2021. **Best Paper Honorable Mention Award**.
- [7] Linguistic effects on news headline success: Evidence from thousands of online field experiments (registered report protocol).
Kristina Gligorić, George Lifchits, Robert West and Ashton Anderson. **PLOS ONE**, 2021.
- [8] Sudden Attention Shifts on Wikipedia Following COVID-19 Mobility Restrictions.
Manoel Horta Ribeiro*, **Kristina Gligorić***, Maxime Peyrard*, Florian Lemmerich, Markus Strohmaier and Robert West (* equal contributions). International AAAI Conference on Web and Social Media (**ICWSM**), 2021.
- [9] Global maps of travel time to healthcare facilities.
Daniel Weiss, Andrew Nelson, Camilo Vargas-Ruiz, **Kristina Gligorić**, Shailesh Bavadekar, Evgeniy Gabrilovich, Amelia Bertozzi-Villa, Jennifer Rozier, Harry Gibson, Tomer Shekel, Chaitanya Kamath, Allison Lieber, Kevin Schulman, Yang Shao, Vesa Qarkaxhija, Anita Nandi, Suzanne Keddie, Susan Rumisha, Punam Amratia, Rohan Arambepola, Elisabeth Chestnutt, Justin Millar, Tasmin Symons, Ewan Cameron, Katherine Battle, Samir Bhatt and Peter Gething. **Nature Medicine**, 2020.
- [10] Causal Effects of Brevity on Style and Success in Social Media.
Kristina Gligorić, Ashton Anderson and Robert West. ACM Conference on Computer-Supported Cooperative Work and Social Computing (**CSCW**), 2019.
- [11] Comparing and Developing Tools to Measure the Readability of Domain-Specific Texts.
Elissa Redmiles, Lisa Maszkiewicz, Emily Hwang, Dhruv Kuchhal, Everest Liu, Miranda Morales, Denis Peskov, Sudha Rao, Rock Stevens, **Kristina Gligorić**, Sean Kross, Michelle Mazurek and Hal Daume III. The 2019 Conference on Empirical Methods in Natural Language Processing (**EMNLP**), 2019.
- [12] Message Distortion in Information Cascades.
Manoel Ribeiro, **Kristina Gligorić** and Robert West. The Web Conference (**WWW**), 2019.
- [13] How Constraints Affect Content: The Case of Twitter’s Switch from 140 to 280 Characters.
Kristina Gligorić, Ashton Anderson and Robert West. International AAAI Conference on Web and Social Media (**ICWSM**), 2018.
- [14] Visible Light Communication Based Indoor Positioning via Compressed Sensing.
Kristina Gligorić, Manisha Ajmani, Dejan Vukobratović and Sinan Sinanović. **IEEE Communication Letters**, 2018.

TEACHING ACTIVITIES

Applied Data Analysis (EPFL CS-401), Head Teaching Assistant
Fall 2018-2019, Fall 2019-2020, Fall 2020-2021, Fall 2021-2022

Data Visualization (EPFL COM-480), Teaching Assistant
Spring 2019-2020, Spring 2020-2021

Analysis II (EPFL MATH-106), Teaching Assistant
Spring 2018-2019

ACADEMIC SERVICE

Organization of conferences

ICWSM 2022 Sponsorship Co-chair
ICWSM 2021 Student Volunteer

Conference Program Committee Member

TheWebConf 2022, *ICWSM* 2021, 2022, *CIKM* 2020, 2021, *WebSci* 2022, *EMNLP* 2021, *ACL* 2021

Workshop Program Committee Member

Wiki Workshop @ TheWebConf 2020, 2021, 2022, *Euro CCS* 2018, 2019

Reviewer

TheWebConf 2020, *CSCW* 2021, 2022, *ICWSM* 2020, 2022, EPJ Data Science

Student Reviewer

CSCW 2019, *ACL* 2020

Institutional responsibilities

PhD admissions committee Jan. 2022 (EPFL School of Computer and Communication Sciences)