

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Towards theories of social interaction over networks

A dissertation submitted in partial satisfaction of the
requirements for the degree of Doctor of Philosophy

in

Electrical Engineering (Communication Theory and Systems)

by

Lorenzo Coviello

Committee in charge:

Professor Massimo Franceschetti, Chair
Professor James H. Fowler
Professor Young-Han Kim
Professor Alon Orlitsky
Professor Ramamohan Paturi

2015

Copyright

Lorenzo Coviello, 2015

All rights reserved.

The Dissertation of Lorenzo Coviello is approved and is acceptable in quality and form for publication on microfilm and electronically:

Chair

University of California, San Diego

2015

TABLE OF CONTENTS

Signature Page	iii
Table of Contents	iv
List of Figures	viii
List of Tables	xi
Acknowledgements	xiv
Vita	xvii
Abstract of the Dissertation	xix
Chapter 1 Introduction	1
1.1 Summary of contributions	5
Chapter 2 Human matching behavior in social networks: an algorithmic perspective	8
2.1 Introduction	8
2.1.1 Related literature	13
2.2 Methods	16
2.2.1 The Matching Games	16
2.2.2 The Algorithmic Model	18
2.3 Results	20
2.3.1 Mathematical Results	20
2.3.2 Validation	25
2.4 Discussion	26
2.5 Proof of Theorem 1	28
2.6 Proof of Theorem 2	31
2.7 Proof Theorem 3	33
2.7.1 Analysis	33
2.7.2 Properties of matchings in \mathcal{M}_1	33
2.7.3 The tree T_n^*	35
2.7.4 Completing the proof of of Theorem 3	38
2.8 Proof of Theorem 4	41
2.9 Acknowledgments	43
Chapter 3 An instance of distributed social computation: the multi-agent group membership problem	44
3.1 Introduction	44
3.2 Related work	47

3.3	The group membership task	50
3.4	The group membership algorithm	52
3.5	Complexity results	55
3.6	Experiments of human social computation	60
3.7	Discussion	63
3.8	Proof of Theorem 5	66
3.9	Proof of Lemma 6	68
3.10	Proof of Theorem 6	70
3.10.1	Properties of the matchings in \mathcal{M}_n	71
3.10.2	The tree T_m^*	73
3.10.3	The dynamics of the algorithm starting from $M \in \mathcal{M}_n$	76
3.10.4	The fraction of matchings $M \in \mathcal{M}_n$ such that $h(M) \geq \gamma n$	79
3.11	Acknowledgments	80
Chapter 4	Non-invasive detection of emotional contagion in on-line social networks	81
4.1	Introduction	81
4.1.1	Related work	85
4.2	Variables of the model	86
4.2.1	Quantifying the semantic of text-based expression	87
4.2.2	Exogenous control variable	89
4.2.3	Social network information	90
4.3	Data	91
4.4	A model of emotional contagion	93
4.5	Aggregating the model	95
4.6	Model estimation	97
4.6.1	Placebo test	101
4.6.2	Controlling for topic contagion	106
4.7	Quantifying the total effect of a user on her friends	111
4.8	How rain affects friends in other cities	114
4.9	Discussion	120
4.10	Acknowledgments	121
Chapter 5	Approximating physical encounter with friendship to predict epidemic outbreaks	122
5.1	Introduction	122
5.1.1	Outline	126
5.2	Dataset	126
5.3	Static and time-varying networks	130
5.4	Infection dynamics	132
5.4.1	Infection time	133
5.4.2	Seed selection	134
5.4.3	Detection time with sensors	134

5.4.4	Sensor selection	134
5.5	Infection detection – Time-varying networks	135
5.5.1	Single seed – Infection rate	136
5.5.2	Single seed – Sensor monitoring	137
5.6	The infected population – Time-varying networks	144
5.6.1	Certain infection ($\beta = 1$)	145
5.6.2	Stochastic infection ($\beta < 1$)	150
5.7	Infection detection – Static networks	158
5.7.1	Single seed – Infection rate	158
5.7.2	Single seed – Sensor monitoring	162
5.8	The infected population – Static networks	166
5.9	Acknowledgments	168
Chapter 6	Query incentive networks with split-contracts: robustness to individuals’ selfishness	171
6.1	Introduction	171
6.1.1	Query Incentive Networks	173
6.1.2	Our results	175
6.1.3	Additional related work	176
6.2	Preliminaries	178
6.2.1	Split contracts	179
6.2.2	Propagation of the payment	180
6.2.3	Difference with respect to previous work	181
6.2.4	Roadmap	182
6.3	Properties of Nash Equilibria	183
6.4	The Nash Equilibrium	185
6.5	Guaranteeing h -consistency	188
6.6	Efficiency	189
6.7	Proofs	191
6.7.1	Proof of Lemma 16	191
6.7.2	Proof of Lemma 17	192
6.7.3	Proof of Theorem 7	193
6.7.4	Proof of Theorem 8	196
6.7.5	Proof of Theorem 9	199
6.7.6	Proof of Lemma 19	201
6.7.7	Proof of Lemma 20	203
6.7.8	Non-uniqueness of the Nash equilibrium	205
6.8	Acknowledgments	207
Chapter 7	Matching markets with bundle discounts: computing efficient, stable and fair solutions	208
7.1	Introduction	208
7.2	The model	213

7.3	Existence of rational and stabilizing transfers	219
7.4	Computation of fair transfers	225
7.4.1	Step 1: rational and stabilizing group transfers	226
7.4.2	Step 2: transfer between buyers	228
7.5	Computation of social welfare maximizing matching	231
7.6	Discussion	234
7.7	Emptiness of the core	235
7.8	Maximizing the social welfare is not necessary for stability	235
7.9	Proof of Lemma 22	236
7.10	Proof of Proposition 1	239
7.11	Computational complexity for determining SWM matchings	240
7.12	Acknowledgments	241
	Bibliography	242

LIST OF FIGURES

Figure 2.1.	Computer interface. The subject is matched with the node on the right and is being requested by three unmatched nodes.	10
Figure 2.2.	Approximate and maximum matching. Left: an approximate maximum matching of size 5 on a network with 12 nodes. Right: a maximum matching of size 6 on the same network.	11
Figure 2.3.	Affinity between humans' and algorithm's performance, 16-node networks. Performance of the human subjects (red points) and of the algorithm (blue points) over eight bipartite 16-node networks (triangles) and eight non-bipartite 16-node networks (circles).	12
Figure 2.4.	Affinity between humans' and algorithm's performance. Performance of the human subjects (red points) and of the algorithm (blue points) over different 24-node networks: small-world (triangles), ring (diamonds), preferential attachment (circles).	13
Figure 2.5.	Algorithm's asymptotic performance. PRUDENCE algorithm's performance with respect to the network's size for the "bad" graph G_n (black diamonds), for preferential attachment model (green squares), small-world model (red triangles).	14
Figure 2.6.	Performance of the experimental subjects on networks of 24 nodes. The plot shows the time to reach a perfect matching of size 12 (red), an approximate matching of size 11 (a 0.92–approximate matching, in blue) and a matching of size 6 (a 1/2–approximate matching, in green).	15
Figure 2.7.	The bad graph. The "bad" graph G_n for $n = 3$. One of the "bad" matchings of Theorem 3 is highlighted in red.	24
Figure 2.8.	Tree T_n^* . Tree T_n^* with labels, for $n = 6$	36
Figure 3.1.	Example of a bipartite network between leaders and followers determined by physical constraints.	45
Figure 3.2.	Computer interface of a participant playing the role of a leader. . .	47
Figure 3.3.	Deficit-decreasing path. Top: a deficit-decreasing path of length 5. Bottom: the path is "solved" by turning each matched edge into an unmatched edge and vice versa.	56

Figure 3.4.	The network G_n for $n = 6$	58
Figure 3.5.	Algorithm's performance on the hard networks G_n	59
Figure 3.6.	Algorithm's performance to reach $(1 - \epsilon)$ -approximate best matchings on random bipartite networks.	60
Figure 3.7.	Algorithm performance versus human subjects performance.	61
Figure 3.8.	Human subjects performance: approximation versus stability.	62
Figure 3.9.	The leader ℓ in the proof of Lemma 3.9, matched edges are highlighted.	68
Figure 3.10.	An example of a matching M of G_6 with $d(M) = 1$. M is uniquely determined by the set $\mathcal{S}(M) = \{2, 4, 6\}$	72
Figure 3.11.	The three T_m^* for $m = 5$	73
Figure 4.1.	(a) Fraction of status updates containing positive words; (b) fraction of status updates containing negative words; (c) the 100 most populous cities in the U.S. and their average fraction of posts with positive words; (d) network of between-city ties.	93
Figure 4.2.	Models estimates. (a) Difference in emotional expression between days with and without rain, estimates derived from first stage regressions. (b) Estimates of emotional contagion between friends, from second stage regressions.	106
Figure 4.3.	Total number of negative posts generated by a day of rainfall within a city (direct) and in other cities via contagion (indirect).	117
Figure 5.1.	Inverse Cumulative Distribution Function of friends degree and encounter degree.	128
Figure 5.2.	Heat map of friend degree and encounter degree of all users with at least one friend and one encounter.	129
Figure 5.3.	Percentile plot of the Jaccard similarity of the set all user's friends and the set of all user's encounters.	130
Figure 5.4.	Susceptible-exposed-infected process on the time-varying friendship and encounter networks, with $\beta = 1$ (certain infection): Average of the last time of infection and average of the final infection.	138

Figure 5.5.	Fraction of infected nodes over time, for the time-varying friendship and encounter networks.	139
Figure 5.6.	Infection detection with random sensors and friend sensors on the friendship and encounter time-varying networks.	142
Figure 5.7.	Fraction of simulations that reached a target sensors' infection versus the infection starting time.....	143
Figure 5.8.	Metrics $J(m; s_i)$, $P_E(m; s_i)$, $P_F(m; s_i)$ for 5000 simulation pairs, each with a random choices of a single seeds, and different values of the target set size m	149
Figure 5.9.	Single seed infection on time-varying networks: Jaccard similarity.	156
Figure 5.10.	Single seed infection on time-varying networks: precision measures.	157
Figure 5.11.	Infection speed versus degree.	160
Figure 5.12.	Growth of the infection over time.	161
Figure 5.13.	Sensor infection monitoring versus seed degree.	164
Figure 5.14.	Sensor infection monitoring versus seed degree.	165
Figure 5.15.	Single seed infection on the static networks: Jaccard similarity. ..	169
Figure 5.16.	Single seed infection on the static networks: precision measures. .	170
Figure 7.1.	Example of cross-transfer graph.	223
Figure 7.2.	Scheme of the flow network \mathcal{G}	227
Figure 7.3.	Scheme for the flow network $\mathcal{G}(\pi)$	233

LIST OF TABLES

Table 4.1.	Summary statistics for each emotional and meteorological variable.	89
Table 4.2.	Number of rainy days and number of total observed days for each of the 100 cities in the dataset.	90
Table 4.3.	List of the 100 cities in the dataset.	92
Table 4.4.	Summary statistics of the dataset.	92
Table 4.5.	Instrumental variable regression estimates: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).....	102
Table 4.6.	Instrumental variable regression estimates: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).....	102
Table 4.7.	Instrumental variable regression estimates: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).....	103
Table 4.8.	Instrumental variable regression estimates: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).....	103
Table 4.9.	Instrumental variable regression estimates: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations).....	104
Table 4.10.	Instrumental variable regression estimates: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations).....	104
Table 4.11.	Instrumental variable regression estimates: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations).....	105
Table 4.12.	Instrumental variable regression estimates: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations).....	105

Table 4.13.	Placebo test: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations).	107
Table 4.14.	Placebo test: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations).	107
Table 4.15.	Placebo test: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations).	108
Table 4.16.	Placebo test: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations).	108
Table 4.17.	Terms used to identify status updates on the topic of weather.	110
Table 4.18.	Controlling for topic contagion: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).	110
Table 4.19.	Controlling for topic contagion: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).	111
Table 4.20.	Controlling for topic contagion: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).	112
Table 4.21.	Controlling for topic contagion: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations).	113
Table 4.22.	Average cumulative effect of a user on her friends.	114
Table 4.23.	Indirect and direct effect on negative emotion of rain in a city (the 50 most populous cities out of 100).	118
Table 4.24.	Indirect and direct effect on negative emotion of rain in a city (the 50 least populous cities out of 100).	119
Table 5.1.	Single seed infection on the time-varying networks. Jaccard similarity measures.	153

Table 5.2.	Single seed infection on the time-varying networks. Precision measures.	155
Table 5.3.	Single seed infection on the static networks. Jaccard similarity measures.	168
Table 5.4.	Single seed infection on the static networks. Precision measures ..	168

ACKNOWLEDGEMENTS

I thank everyone that in one way or another helped me getting to this point, in particular:

My advisor Massimo Franceschetti . He has been an inspiring mentor and has supported me since day one of my graduate studies. He taught me how to do rigorous research, to be the most severe critic of my own work, and that great content and great form must coexist in good research. I am definitely honored of being part of “Franceschetti’s group”.

Mohan Paturi, who introduced me to the study of social networks.

James Fowler, who introduced me to the incredible field of data science.

Alon Orlitsky and Young-Han Kim, for being part of my dissertation committee and for always giving me honest feedback.

Manuel Cebrian, who introduced me to the DARPA Network Challenge, to the data science team at Facebook, and to the Media Lab.

Andrea Vattani and Panos Voulgaris, coauthors and friends who made writing difficult papers fun.

Paolo Minero, for guiding me through my first steps in research.

Mathew McCubbins for teaching me how to run experiments. Daniel Enemark, Devin Barr and Matthew Jones for their help with the experiments.

Nicholas Christakis, Cameron Marlow and Adam Kramer for their help and feedback.

My coauthors and friends of the Human Nature Group, Jaime Settle, Jason Jones, Chris Fariss, Robert Bond, Yunkyu Sohn.

Everyone that stepped over the 4th floor of the Atkinson Hall in the past five years, for creating a fun working environment, and most of all for being patient about my (and Emanuele’s) loud Italian-speaking.

The professors at the University of Padova, Michele Zorzi and Andrea Zanella.

The great teachers I had the fortune to meet throughout the years and whose teaching has been important well beyond the classroom: Mary Jewel Bottaro, Maria Rotondo, Mario Cervelli and Angelo D’Onofrio.

Even before developing a method to detect and quantify their influence, I had no doubt about the importance of friends, especially some of them. Gabriele, Paolo, Zana, Elia, Nicola, Giacomo, Mark, Liz, Gioele, Ehsan (the order is strictly chronological).

I cannot express in writing how thankful I am to the most important people in my life. Silvia, the sugar, salt and pepper of my days. Emanuele, cloning myself would not have been comparable. Mamma, for her unconditional love and for introducing me to the literary production of Richard Scarry. Papà, for so proudly encouraging me to take this journey. Nonna, whose teaching is timeless and more valuable than a doctorate. His fluffiness Bimino.

Chapter 2, in full, is a reprint of the material as it appears in PLoS ONE, 7(8) e41900, 2012, L. Coviello, M. Franceschetti, M. McCubbins, R. Paturi, and A. Vattani. The dissertation author was the primary investigator and co-author of this paper.

Chapter 3, in full, has been submitted for publication of the material as it might appear in the IEEE Transaction on Network Control Systems. L. Coviello, and M. Franceschetti. The dissertation author was primary investigator and co-author of this paper.

Chapter 4, in full, is a reprint of the material as it appears in the Proceedings of the IEEE, Vol. 102, Issue 12, pp. 1911–1921, 2014, L. Coviello, J. Fowler, and M. Franceschetti, and in PLoS ONE, 9(3), e90315, 2014, L. Coviello, Y. Sohn, A. Kramer, C. Marlow, M. Franceschetti, N. Christakis, and J. Fowler. The dissertation author was the primary investigator and co-author of these papers.

Chapter 5, in part, is currently being prepared for submission for publication of

the material. L. Coviello, and M. Franceschetti. The dissertation author was the primary investigator and co-author of this paper.

Chapter 6, in full, is a reprint of the material as it appears in the Proceedings of the forty-fourth annual ACM symposium on Theory of computing, STOC 2012, pp. 775–788, M. Cebrian, L. Coviello, A. Vattani, and P. Voulgaris. The dissertation author was co-primary investigator and co-author of this paper.

Chapter 7, in full, has been submitted for publication of the material as it might appear in the ACM Transaction on Economics and Computation, under submission. L. Coviello, and M. Franceschetti.

VITA

- 2006 Bachelor of Science in Information Engineering, Università degli Studi di Padova
- 2008 Master of Science in Telecommunication Engineering, Università degli Studi di Padova
- 2013 Master of Science in Electrical Engineering (Communication Theory and Systems), University of California, San Diego
- 2015 Doctor of Philosophy in Electrical Engineering (Communication Theory and Systems), University of California, San Diego

PUBLICATIONS

L. Coviello, J. Fowler, and M. Franceschetti, “Words on the web: non-invasive detection of the spread of emotion in on-line social networks.” *Proceedings of the IEEE*, Vol. 102, Issue 12, pp. 1911–1921, 2014.

L. Coviello, Y. Sohn, A. Kramer, C. Marlow, M. Franceschetti, N. Christakis, and J. Fowler. “Detecting emotional contagion in massive social networks.” *PLoS ONE*, 9(3), e90315, 2014.

L. Coviello, M. Franceschetti, M. McCubbins, R. Paturi, and A. Vattani. “Human matching behavior in social networks: an algorithmic perspective.” *PLoS ONE*, 7(8) e41900, 2012.

L. Coviello, and M. Franceschetti. “Distributed social computation: the multi-agent group membership problem.” *IEEE Transaction on Network Control Systems*, accepted for publication.

L. Coviello, M. Franceschetti, “Distributed team formation in multi-agent systems: stability and approximation.” In *Proceedings of the fifty-first IEEE Conference on Decision and Control, CDC 2012*, pp. 3776-3782.

M. Cebrian, L. Coviello, A. Vattani, and P. Voulgaris. “Finding red balloons with split contracts: robustness to individual selfishness.” In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing, STOC 2012*, pp. 775–788.

L. Coviello, and M. Franceschetti. “Matching markets with bundle discounts: computing efficient, stable and fair solutions.” Submitted for publication.

P. Minero, L. Coviello, and M. Franceschetti. “Stabilization over Markov feedback channels: the general case.” *IEEE Transaction on Automatic Control*, Vol. 58, Issue 2, pp. 349–362, 2013.

L. Coviello, P. Minero, M. Franceschetti, “Stabilization over Markov feedback channels.” In *Proceedings of the fiftieth IEEE Conference on Decision and Control and European Control Conference, CDC-ECC 2011*, pp. 3776-3782.

ABSTRACT OF THE DISSERTATION

Towards theories of social interaction over networks

by

Lorenzo Coviello

Doctor of Philosophy in Electrical Engineering (Communication Theory and Systems)

University of California, San Diego, 2015

Professor Massimo Franceschetti, Chair

The study of human interaction has taken an unprecedented quantitative and technological turn, driven by the advent of online social networks, the consequent availability of massive tracks of human behavior, and advancements in computing and analytical tools. To provide a rigorous understanding of social networks, this dissertation proposes theories of human networked interaction, able to explain and predict individual and global outcomes observable in a population. Populations of heterogeneous, complex individuals are modeled using homogeneous, simple agents who act according to simple incentives and local rules. The resulting models are prone to analysis using algorithmic

complexity, game theory and statistics, enabling the test of predictions about the original population. Rigorous models and methods are proposed to study several aspects of human interaction: scenarios of social computation, in which interconnected individuals cooperate to solve a problem in a distributed fashion; the analysis of online social networks, to study influence spreading over networks and to test hypothesis of behavioral interaction; the theoretical analysis of human populations, to understand fundamental capabilities and limits of social coordination in complex scenarios.

Chapter 1

Introduction

To what extent can complex interactions between multiple people be predicted? Can the rigor of mathematics and the insight of engineering bring clarity to the complex mechanisms underlying social interaction?

The discourse on social behavior is probably as old as societies themselves. Written records of social reasoning date back to the ancient Greek philosopher Plato in the 4th Century BC [59], and are common to most cultures and philosophies, from Confucianism to medieval Islam. The study of society began to gain his modern form during the 19th Century, when French philosopher August Comte proposed sociological positivism [56], the application of the scientific method to society, to address and solve social issues, and first used the expression “social physics”. Since then, scholars have investigated all facets of social interaction, incorporating methods and techniques ranging from theory and statistics to surveys and controlled experiments. In the last decades, this discipline has taken an unprecedented quantitative and technological turn, driven by the advent of the Internet, the consequent availability of massive tracks of human behavior and interaction, and advancements in computing and analytical tools [138]. Computational social science and social network analysis flourished at the intersection of computer science and the social sciences, gaining increasing relevance in academia, industry and policymaking. The most prestigious universities have founded graduate

programs focused on data science. Social network and data-centered technologies are disrupting every possible market. In February 2015, President Barak Obama appointed Doctor DJ Patil as the first U.S. Chief Data Scientist at the White House Office of Science and Technology Policy.

This dissertation builds knowledge about the networked behavior of individuals and groups, proposing *theories of human networked interaction*, capable of explaining and predicting individual and global outcomes observable in a population. The investigation herein contained is motivated by questions at the foundation of disciplines such as economics, social science, health and computation. In order to answer these questions, given a heterogeneous population of complex individuals who exchange information, my modeling approach employs a synthetic population of individuals, or *agents*, who make decisions according to simple incentives and rules of local interaction. Despite their simple rules, agents have the descriptive power of the original population, and their actions are prone to analysis using algorithmic complexity, game theory and statistics. Such descriptive models allow to make and test predictions about the original population, including how behaviors may be triggered by events and conditions that would not be apparent in controlled experiments.

This dissertation covers three areas that I explored through my research: social network computation, to describe the behavior of groups of individuals cooperating to perform a task in a distributed fashion (Chapter 2 and Chapter 3); the analysis of online social networks, to study information and influence spreading over networks, and to test hypotheses of behavioral interaction (Chapter 4 and Chapter 5); the theoretical analysis of human populations, to understand fundamental capabilities and limits of social coordination in complex scenarios (Chapter 6 and Chapter 7).

Regarding the first area, Chapter 2 and Chapter 3 consider scenarios in which a group of interconnected individuals have to compute the solution to a problem. A single

individual does not have knowledge of the entire network structure and state, and is incapable of computing the solution alone. Therefore, the solution must be reached in a distributed fashion, through local interaction and information exchange. A recent line of research based on controlled laboratory experiments initiated by Michael Kearns [124] and followed by others [122, 114, 67, 151] shows that humans can successfully coordinate in a wide range of tasks and that the network structure and incentives affect this capability to coordinate. Moving beyond the purely observational approach, and inspired by Nobel laureate Herbert Simon's claim that algorithmic processes are at the basis of human decision [194], my work attempts to discover the fundamental mechanisms enabling the distributed solutions observed in such networked tasks. In particular, starting from controlled laboratory experiments on human networks, I propose algorithmic models of local interaction that match the empirical data and are prone to mathematical analysis, and I state and test predictions in the form of performance guarantees on any network – rather than the limited set initially observed in the laboratory. My work shows the feasibility of this approach, and the possibility of identifying simple rules of interaction (among the multitude reported in the post-experimental surveys) that can describe and predict global outcomes in a population, addressing an open question about the possibility to model human behavior with algorithms [121].

Regarding the second area, Chapter 4 considers the problem of causal inference from observational data. A fundamental difficulty of observational studies is to distinguish between causation and correlation [150, 152] – the researcher's ultimate goal is to perform causal inference, but the lack of controlled experimental treatments usually hinders its feasibility. On the other hand, methods based on observational rather than experimental approaches are desirable in that they avoid the manipulation of user experience, a research practice which recently became object of criticism and concern. In particular, Chapter 4 proposes a method to detect and quantify the spread of semantic expression in on-line

social networks using observational data. Successfully applied to the emotional content of a corpus of billions of Facebook status updates spanning a period of three years, the method responds to the need of developing rigorous, noninvasive methods for the study of human interaction. The limitations typical to observational studies are overcome through mathematical modeling of the influence process and estimation based on instrumental variables regression [9], a technique introduced in econometrics. In particular, the method allows to measure the effect of an exogenous variable (for example rainfall) on the expression of an individual and, consequently, its effect on the expression of others to whom that individual is socially connected.

Chapter 5 studies the diffusion of infectious processes on real-world complex networks. The theme is of interest to diverse disciplines as similar models have been proposed to characterize the spread of information, behaviors, cultural norms, innovation, as well as the diffusion of computer viruses [180, 170], and is of current importance in light of events such as the recent ebola epidemics [86, 93]. Considering a data set of human behavior and social relationships from the online review service Yelp, and assuming an infection that spreads through physical contact between individuals, Chapter 5 characterizes how accurately and within which limits said process can be predicted and approximated if the researcher has only access to explicit relationship ties between individuals (i.e., friendship ties). At a macroscopic level, friendship is shown to provide a valid approximation with respect to physical encounter, confirming the observation that different real-world complex networks present similar structural properties and are governed by similar rules [5]. At a microscopic level, instead, it is shown that friendship does not provide accurate prediction of the individuals at risk if the infection is driven by physical encounter, highlighting the importance of local connectivity in network dynamics.

Regarding the third area, my work considers complex coordination scenarios that

can be formalized in terms of cooperation games and network dynamics. In a wide variety of settings, how to design incentives that guarantee the performance and the outcomes of a networked population is not known. Yet the potential benefits of such design capabilities could be far reaching in domains as diverse as crowd sourcing, social mobilization and habit formation. On the one hand, if the limitations to successful coordination are of a strategic nature (social behavior), then the designer's effort should be directed towards the definition of more efficient strategies. In the context of crowdsourcing, motivated by the worldwide attention received by the DARPA Network Challenge¹, Chapter 6 considers the problem of large-scale, time-critical social mobilization for information acquisition. Inspired by the winning strategy of the MIT Media Lab [177], it shows that social mobilization is always efficient if suitable strategies are available, positively answering to a question posed by Jon Kleinberg and Prabhakar Raghavan [129]. On the other hand, if there exist provable theoretical limits to coordination (as graph structural constraints, or complexity constraints) that hold for a variety of strategies, then only proper market or policy design can overcome such limitations. Such inefficiencies often occur in games between selfish agents, where the Price of Anarchy [134] encodes the gap between the equilibria reached by the population and the socially optimal outcome. Motivated by the observation that stability and social efficiency do not always coexist in a market of selfish agents [98], Chapter 7 proposes cooperation between agents as a solution to such misalignment, able to drive the population towards socially desirable outcomes.

1.1 Summary of contributions

This dissertation is divided into three parts. The first part (Chapter 2 and Chapters 3) proposes the use of simple algorithms to model the behavior of interconnected

¹<https://networkchallenge.darpa.mil/>

individuals. The second part (Chapter 4) proposes a non-invasive, statistical methods to perform causal inference in social networks, based on observational data. The third part (Chapter 6 and Chapters 7) proposes game theoretical models to investigate stability and efficiency in coordination problems.

Chapter 2 considers an experimental scenario in which individuals on a virtual network can form pairs by means of simple interaction. Individuals receive a monetary reward if they find a maximum matching of the network (i.e., the maximum possible number of mutually disjoint pairs) in a timely fashion. Algorithmic modeling is shown to be a powerful approach to understand and predict the collective dynamics of human behavior.

Chapter 3 extends the framework of algorithmic behavioral modeling of Chapter 2 to a more complex coordination scenario, in which individuals of one type, the leaders, have to recruit individuals of a different type, the followers, in order to form groups in a timely and stable fashion.

Chapter 4 proposes a method to detect and quantify the spread of semantic expression online, able to perform causal inference from observational data. In particular, it describes a contagion model and an estimation approach based on instrumental variable regression, applied to the emotional content of a corpus of billions of Facebook status updates.

Chapter 5 studies the diffusion of infectious processes on complex networks, considering a data set of human behavior and social relationships from the online review service Yelp. In particular, assuming an epidemics that spreads via physical encounter, it characterizes how accurately said process can be approximated knowing only the friendship ties between individuals.

Chapter 6 considers a problem of information acquisition in a strategic networked environment, in which agents receive incentives to recursively recruit each other and to

take part to the search. A recursive recruitment and payment strategy is defined and its stability and efficiency are proven.

Chapter 7 considers a market scenario in which selfish agents are interested in triggering discounts in order to maximize their utilities. Despite stability and social efficiency do not coexist in general, if cooperation between agents in the form of utility transfers is allowed, then this misalignment is always resolved, and it is computationally feasible to determine the amount of cooperation required.

Chapter 2

Human matching behavior in social networks: an algorithmic perspective

2.1 Introduction

The modeling and prediction of collective human behavior has been one of the key challenges of social sciences for several decades. As early as 1947, Herbert Simon argued that *information processing* constitutes the core of human decision-making [194]. A corollary of his argument is that human decision-making processes can be modeled *algorithmically*. However, such algorithmic modeling and prediction is challenging, considering that collective decision-making processes are driven by both individual attitudes and collective dynamics, and often involve social interchange and mutual agreement.

This paper argues that despite the inherent complexity of human social interactions, it is possible to isolate basic behavioral principles, formulate mathematical models, and predict collective dynamics, using an algorithmic approach. As a simple example of this approach, in the context of a distributed coordination game on networks (i.e., the *maximum matching game*), we present an algorithmic model of human behavior that is based on *simple* principles of local interaction and that is able to capture *complex* collective coordination.

Our approach is similar in spirit to the one in physics where particle systems and cellular automata described by simple rules are known to generate complex behaviors, such as phase transitions and universal computability [141, 36, 211, 57]. However, our algorithmic modeling approach embeds individual interaction behavior as part of a distributed computing system and leads to computational complexity analysis.

Our work is influenced by the work of Kearns et al. [124] who studied the effect of network topology on subjects' ability to color a graph, and by subsequent work in the context of distributed coloring and consensus games [122, 151, 114, 67]. However, our focus is on algorithmic modeling and analysis, rather than on observing the effect of network topology on performance.

We have conducted over 250 experiments with human subjects on a pool of over 80 networks with up to 24 nodes each, ranging from simple networks to more complex stochastic models including preferential attachment [26, 33] and small-world networks [215]. Our experimental set-up is simple. Subjects are represented by nodes of a network with edges representing potential matches. In our experiments, human subjects are connected over a virtual network and interact with their neighbors through a computer interface, see Figure 2.1. Subjects can form and destroy pairs with their neighbors, and each subject can be part of a single pair at a time. Subjects are given only local information about their immediate neighbors, and can only interact with them. They are able to propose to match with a neighbor and accept a proposal from a neighbor. While matched, a subject can also make a proposal to or accept a proposal from another neighbor; in both cases, the existing match would automatically be broken. Moreover, a subject can only have a single outstanding proposal at a time. Therefore, at any time, a subject can either be part of a matched pair, or not be matched and have at most a single outstanding proposal. Subjects are equally incentivized to achieve a *maximum matching*, namely to form the maximum number of disjoint mutual pairs, without regard to whom

is matched with whom. Specifically, they are given an equal monetary reward for each game where a maximum matching is found within the allotted time.




Figure 2.1. Computer interface. The subject is matched with the node on the right and is being requested by three unmatched nodes.

To better understand this setup, consider the following metaphor: imagine that incoming graduate students are pairing up with faculty members. Further imagine that every member of the department prefers every graduate student to have one adviser and every adviser to have one graduate student, and only certain faculty and graduate students share interests. Communication is limited so that individuals can only tell if members with whom they share an interest are already matched. Each member of the department is now a node, the edges represent shared interest, and individuals can then propose to work with members with whom they share an edge.

Our algorithmic model is based on a simple property that we call “prudence” and that emerges from the analysis of a first set of experimental data. This property states that *individuals do not break existing matched pairs unless they receive an alternative proposal by an unmatched neighbor*. Based on this property, we propose a simple distributed algorithm, analyze its performance, validate the model with additional experimental results, and predict outcomes. The prudence property is reminiscent of the notion of risk aversion, a relevant topic in the economics literature [179, 116].

We now briefly summarize our findings. Throughout the paper we use the graph-

theoretic terminology, according to which a matching is a set of edges without common nodes. The size of a matching is the number of edges in it. A maximum matching is a matching with the largest size. For $0 < c \leq 1$, a matching is a c -approximate maximum matching if its size is within a factor of c from that of a maximum matching. A matching M is maximal if it is not a proper subset of any other matching, i.e., for any new edge added to it, it is no longer a matching. Figure 2.2 depicts an approximate and a maximum matching of a network.




Figure 2.2. Approximate and maximum matching. Left: an approximate maximum matching of size 5 on a network with 12 nodes (matching edges are represented in bold red, matched nodes are colored, unmatched nodes are white). Right: a maximum matching of size 6 on the same network (note that the maximum matching is also a perfect matching, as all nodes are matched).

We show that the convergence time to the maximum matching in computer simulations of the prudence algorithm fits well the experimental data (after scaling by a constant factor), see Figures 3.7 and 2.4. By computer simulations we also predict that convergence to a maximum matching is slower on preferential attachment networks than on small-world networks, see Figure 2.5. This prediction is validated by our experiments with human subjects. It is also in agreement with the experimental results by Kearns et al. [124] regarding the coloring problem, and with the theoretical results by Montanari and Saberi [156] regarding the spread of innovation in networks.

On the theoretical side, we analyze the dynamics of the prudence algorithm and show that for all graphs of bounded degree a $1/2$ -approximate maximum matching is reached quickly, on average in $O(\log n)$ rounds, where n refers to the number of nodes in




Figure 2.3. Affinity between humans’ and algorithm’s performance, 16-node networks. The performance of the human subjects (red points joined by continuous line) and of the algorithm (blue points) over eight bipartite 16-node networks (triangles) and eight non-bipartite 16-node networks (circles) are plotted. The experiment was run multiple times on each network and the average behavior is reported. The x -axis shows the indexes of the networks sorted by increasing average time required to reach a maximum matching. Bipartite networks are labeled from 1 to 8, while non-bipartite networks are labeled from 9 to 16. The y -axis shows the average time (in seconds) required to reach a maximum matching for humans, while the average number of rounds of the algorithm is scaled by a constant factor.

the network (Theorem 1); and for all graphs a $(1 - \epsilon)$ -approximate maximum matching is reached in polynomially many rounds with high probability (Theorem 2). We also show that there are instances (called “bad” graphs) for which reaching a *maximum* matching requires exponential time with high probability when starting from a set of configurations (called “bad” matchings) which constitute almost all possible configurations (Theorems 3 and 4). These results show that in the worst case there is an exponential gap between reaching a good matching (i.e., an approximate maximum matching whose cardinality is close to a maximum matching) versus the best (i.e., perfect) matching. The experimental data shows (consistently with the theoretical analysis) that human subjects always find a “good” matching quickly, while they can take much longer to improve the solution to a maximum matching, see Figure 2.6. In particular, on the bad graph, human subjects




Figure 2.4. Affinity between humans' and algorithm's performance, 24-node networks. The performance of the human subjects (red points joined by continuous line) and of the algorithm (blue points) over different 24-node networks are plotted. In particular, small-world networks (triangles), a ring network (diamonds), and preferential attachment networks (circles) were tested. The experiment was run multiple times on each network and the average behavior is reported. The x -axis shows the indexes of the networks sorted by increasing average time required to reach a maximum matching. The y -axis shows the average time (in seconds) required to reach a maximum matching for humans, while the average number of rounds of the algorithm is scaled by a constant factor.

could not converge to a maximum matching in the allotted time.

2.1.1 Related literature

The experimental study of human strategic behavior over networks is a topic of great current interest in the literature. The work by Kearns and others on network coloring and consensus games [124, 122, 151, 114, 67] has been particularly influential.

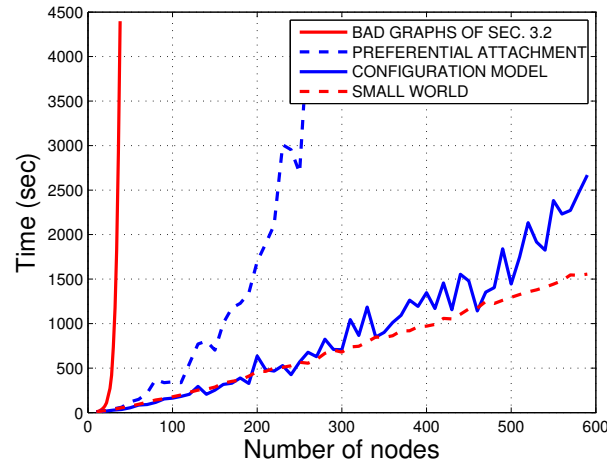


Figure 2.5. Algorithm’s asymptotic performance. PRUDENCE algorithm’s performance with respect to the network’s size for the “bad” graph G_n (black diamonds), for preferential attachment model (green squares), small-world model (red triangles). For each generative model and network size we generated 100 networks and run the algorithm 1000 times on each. The average behavior is reported. The x -axis shows the network size, and the y -axis shows the average number of rounds required by the algorithm to converge to a maximum matching.

Judd et al. [115] investigated how subjects choose between playing either a dominant or a submissive role in a network game, documenting the importance of fairness. Kearns et al. [123] performed experiments on network formation games when there is a cost for creating links. Suri and Watts [204] conducted experiments in which individuals connected over networks play local public good games. Wang et al. [212] studied multi-player prisoner’s dilemma games in which subjects can propose and delete links to other players, showing that partner selection increases cooperation. Brautbar and Kearns [38] introduced a network formation game in which players need to maximize their clustering coefficients. Compared to these previous works, we focus on isolating behavioral principles of human interaction (in the context of maximum matching games) and using these principles to formulate algorithmic predictions of outcomes.

As social interaction naturally induces strategic behavior, our work is also closely related to game theory. Indeed, several authors proposed game theoretical models of

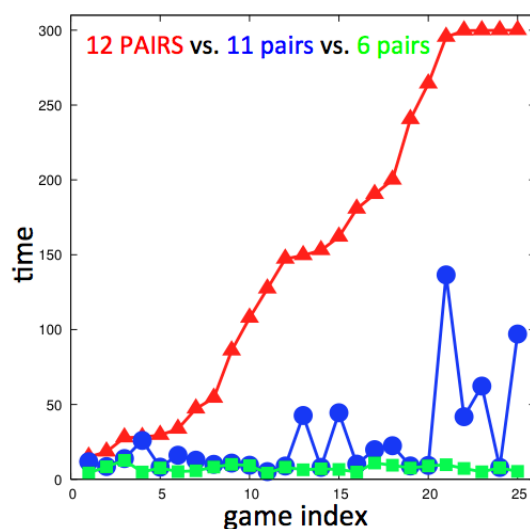


Figure 2.6. Performance of the experimental subjects on networks of 24 nodes. The plot shows the time to reach a perfect matching of size 12 (red), an approximate matching of size 11 (a 0.92–approximate matching, in blue) and a matching of size 6 (a 1/2–approximate matching, in green). Results for single games are reported. The x -axis shows the indexes of the games sorted by increasing solving time, while the y -axis shows the time in seconds. The right-most four games on the red plot did not converge to a maximum matching and correspond to three instances of the “bad” graph G_n and to one instance of the preferential attachment network.

human interaction over social networks. Topics vary from diffusion and contagion over networks [156, 88, 2] to strategic information retrieval [129] (see also Chapter 6), models of segregation [37] and bargaining over networks [117], to mention a few. The main element that distinguishes our work from the game theory literature is that we focus on the algorithmic processes involved in strategic thinking and the ensuing collective dynamics rather than on equilibria. Moreover, our algorithmic model is motivated and supported by experimental data.

Finally, matching theory has received notable attention throughout the decades, both in the context of game theory and economics [186, 77, 96, 17], and in the development of algorithms for the maximum matching problem [103, 118, 106, 148, 175].

We point out that our simplified setup constitutes a simplification of the richness and heterogeneity of the ties in real social networks, as the subjects have no preference over each other, all the ties are equivalent, and interaction has no cost. However, such a simplified model leads to a tractable analysis and to the formulation of a general principle of collective behavior.

2.2 Methods

The experiments included the interaction of the participants through a computer interface, and were conducted in accordance with the ethical standards specified in the 1964 declaration of Helsinki. Written consent was granted before participation in the experiments. Our institutional review boards approved this study (UCSD IRB approval 111213SX, US Army Human Research Protection Office ARO-HRPO Log A-17038).

2.2.1 The Matching Games

Before formulating our algorithmic model, we conducted four sessions of experiments, each with a different pool of sixteen undergraduate students connected over a virtual network. Subsequently, to validate our model, we ran an additional session of experiments with a pool of twenty four subjects on a set of networks that included small world and preferential attachment networks. In each of the first four sessions the subjects were asked to solve the matching game on a pool of over 70 networks. All networks admitted a *perfect matching*, namely a maximum matching with no unmatched nodes. We considered networks classified into four groups: bipartite networks admitting a unique perfect matching, bipartite networks admitting multiple perfect matchings, non-bipartite networks admitting a unique perfect matching, non-bipartite networks admitting multiple perfect matchings. Within each group, networks were randomly generated. As a remark, a bipartite network is a network whose nodes can be divided into two disjoint sets V_1

and V_2 such that every edge connects a vertex in V_1 to one in V_2 . If this property does not hold, we say that the network is non-bipartite. Subjects sat in front of workstations for the entire two-hour duration of the session and had no eye-contact with each other. For each matching game, a network was chosen, subjects were randomly assigned to its nodes, and each subject interacted with its neighbors by making or accepting proposals to form matched pairs using the interface shown in Figure 2.1. Each subject could control the node in the center of the screen and could only see its neighbors and, among those, distinguish which of them were currently matched (marked in dark green). A subject could make proposals or accept proposals by selecting a neighbor with a mouse click, and could only have one outstanding proposal at a time to form a matched pair (circled in yellow). While subjects knew whether a neighbor is matched or unmatched, they did not have direct knowledge of any outstanding requests to their neighbors other than their own. If two neighbors selected each other, a pair was formed (a bright green link appeared between them) which could be broken when one of the partners selected another neighbor. As a remark, since a pair was formed when two subjects selected each other and each subject could make a single selection at a time, each subject could be part of a single pair at a time (with one of its neighbors).

If a perfect matching was found within the time limit of five minutes, the game was declared solved and each participant was rewarded by \$.50 or \$1 depending on the session, otherwise the game ended with no reward. The number of games in an experimental session was not fixed, but games were run for the two-hour duration of the session. Therefore, the number of games and the cumulative reward in a session depended on the performance of the participants, providing an additional incentive to coordinate.

The networks used in this first set of experiments can be divided into four classes: bipartite, non bipartite, unique perfect matching, multiple perfect matchings. Two one-

tailed Welch’s t-tests confirmed the hypotheses that it is harder for humans to complete the matching game on non-bipartite than on bipartite networks (p -value < 0.001); and that non-bipartite networks with unique perfect matching are more difficult to solve than non-bipartite networks with multiple perfect matchings (p -value < 0.001). No statistically significant difference was found between the completion time of bipartite networks with unique and with multiple perfect matchings. We believe that this depended on the small network size of sixteen nodes and we did not explore larger bipartite networks further.

2.2.2 The Algorithmic Model

One of the main behavioral properties that emerged from the experimental data is that matched players *may* break their current matching *only if* they have been requested by an unmatched neighbor. In particular, in 30% of the games no player ever violated this rule at any time during the game. In the remaining games, over 93% percent of the moves were in agreement with this rule. Therefore, this property led to the following modeling assumption:

Assumption 1 (Prudence) *A matched node does not break its current matched pair if it does not receive any request from other neighbors.*

Two remarks are in order. First, note that this behavioral rule is peculiar to the matching problem since each player needs to choose a partner but also needs to be chosen. Second, notice that a matched subject with unmatched neighbors has some incentive to behave non-prudently and break the current match, because the subject can infer from the status of its neighbors that the perfect matching is not reached yet. However, experimental data shows that this rarely happens.

For each node u , let $f(u)$ indicate u ’s current preference. In other words, $f(u)$ is the unique node to which u has currently proposed to. $f(u)$ will be null if u does not have a current proposal. If two neighbors u and v currently prefer each other (i.e., $u = f(v)$ and

$v = f(u)$), then consider them matched and the edge $e = \{u, v\}$ as part of the matching. Assume that each node knows if a neighbor is matched or unmatched.

Given the prudence property, we model the algorithmic behavior of humans using the PRUDENCE algorithm shown in the algorithm box 1. The algorithm is specified by the implementation of two functions, called $\text{MATCHEDCHOOSE}(u)$ and $\text{UNMATCHEDCHOOSE}(u)$, which are placeholders for the behavior that node u would follow depending on whether u is matched or unmatched. We consider a synchronous setting, in which time is divided into rounds, and at the beginning of each round each node observes its status and the status of its neighborhood and then decides on an action to take.

ALGORITHM 1: The PRUDENCE algorithm.

```

if unmatched then
  Set  $f(u) \leftarrow \text{UNMATCHEDCHOOSE}(u)$ 
else
  if matched and  $\exists$  neighbor  $v$  s.t.  $f(v) = u$  then
    Set  $f(u) \leftarrow \text{MATCHEDCHOOSE}(u)$ 
  end
end

```

In the following we provide a canonical implementation of the two functions $\text{UNMATCHEDCHOOSE}(u)$ and $\text{MATCHEDCHOOSE}(u)$ which are consistent with the prudence property. $\text{UNMATCHEDCHOOSE}(u)$ does not change the current value of $f(u)$ with probability p , while with probability $1 - p$ accepts the proposal from a neighbor uniformly at random from among the neighbors v with $f(v) = u$ if any; if there is no neighbor v with $f(v) = u$, then it proposes to a node uniformly at random from among the unmatched neighbors if any; otherwise it proposes to a node uniformly at random from among all the matched neighbors. In other words, unmatched nodes prefer neighbors who requested them over other unmatched neighbors, and unmatched neighbors over matched neighbors. As for matched nodes, $\text{MATCHEDCHOOSE}(u)$ accepts a proposal

from a neighbor uniformly at random from among the neighbors v with $f(v) = u$ (note that u 's current partner is one of them). We remark that the simulations' performance and the fit with the experimental data was practically insensitive to the value of p chosen in the run of the algorithm.

2.3 Results

2.3.1 Mathematical Results

In this section we present our analytical results regarding the convergence behavior of the PRUDENCE algorithm. In particular, our results describe how well the algorithm performs in finding a large matching and the time it takes in terms of the number of rounds required. Due to space constraints, we only present proof sketches here. Complete details of the proofs are deferred to the SI.

We define a *matching* at round t as the set of matched edges at the beginning of round t of the algorithm. We first claim that the prudence property implies that the size of the matching does not decrease with time. The proof is immediate and it is omitted.

Claim 1 *The size of the matching at round t is non-decreasing as t increases.*

We then observe that the behavior of the PRUDENCE algorithm can be described by a Markov chain over matchings. A transition from a matching M to a matching M' is made by selecting an edge $e = \{u, v\}$ such that at least one among u and v is unmatched, and setting $M' = M + e$ if u, v are both unmatched, and $M' = M + e - e'$ if exactly one of u and v is matched in M and e' is the matching edge. This Markov chain is reversible when restricted to matchings of the same size. Since the Markov chain is memory-less and has positive probability of reaching a maximum matching, we conclude that the PRUDENCE algorithm enjoys self-stabilization.

Claim 2 *The PRUDENCE algorithm is a self-stabilizing algorithm.*

Our first theorem says that a $1/2$ -approximate matching will be reached quickly in networks with bounded degree.

Theorem 1 *In any bounded-degree graph on n nodes, the expected number of rounds for the PRUDENCE algorithm to reach a $1/2$ -approximate matching is $O(\log n)$.*

The proof of Theorem 1 is in Section 2.5. The key idea of the proof is to show that, in expectation, the “distance” in terms of number of matched pairs to the smallest *maximal* matching shrinks by a constant factor in each round of the PRUDENCE algorithm. Since it is well known that any maximal matching is a $1/2$ -approximation of the maximum matching, the result then follows.

We remark that the assumption of having bounded degrees is necessary as there are unbounded degree graphs in which a polynomial number of rounds is required with high probability to achieve a $1/2$ -approximation. However, in this case, a polynomial number of rounds is also enough to achieve *any* constant approximation: indeed, as the next theorem states, the PRUDENCE algorithm provides a PTAS (polynomial time approximation scheme) for the maximum matching problem. Given a graph G , Δ denotes its maximum degree.

Theorem 2 *For any graph G of n nodes, $\varepsilon > 0$ and $c \geq 1/2$, the PRUDENCE algorithm reaches a $(1 - \varepsilon)$ -approximate matching in $\frac{c}{\varepsilon} n \Delta^{2/\varepsilon}$ rounds with probability at least $1 - \exp(-c\varepsilon^2 n/2)$.*

The theorem implies that, for any constant $\varepsilon > 0$, a matching whose size is within a $(1 - \varepsilon)$ fraction of the size of the maximum matching is reached in polynomial time. For bounded-degree graphs, this result also holds for $\varepsilon = \Omega(1/\log n)$, implying that in this case a maximum matching can be reached in polynomial time.

The proof of Theorem 2 is in Section 2.6. To prove the theorem, we track the progress of the algorithm towards an approximate maximum matching, using the concept

of an *augmenting path*. An augmenting path is a path of odd length which alternates between matched and unmatched edges and whose extreme edges are unmatched. It turns out that there is a close connection between the size of a shortest augmenting path in a matching and how close the matching size is to the size of a maximum matching. More specifically, we use the following lemma due to Hopcroft and Karp [103].

Lemma 1 *Consider any matching M that does not admit augmenting paths of odd length k or smaller. Then, the size of M is at least a fraction $\frac{k+1}{k+3}$ of the size of a maximum matching.*

Hence, to prove Theorem 2, we need to show that short augmenting paths (for a suitably chosen k) are solved in a short amount of time. It is useful to consider a particle analogy to understand the process that eliminates short augmenting paths. We consider each unmatched node as a particle. Particles move around the graph from node to node as nodes change their status between matched and unmatched states dictated by the random choices in the algorithm. There are exactly two particles along an augmenting path, situated at the extreme nodes. To understand how an augmenting path gets shorter and eventually vanishes, we consider how the two particles move closer to each other along the path.

Let $u_0, u_1, u_2, \dots, u_\ell$ denote a *shortest* augmenting path. If the extreme unmatched node u_0 proposes to u_1 and u_1 accepts the proposal breaking the current match with u_2 , then the particle moves from u_0 to u_2 . A similar argument applies to the other end of the path. Also, the minimality of the path guarantees that the internal nodes do not change their current matching as they have no unmatched neighbor. It follows that the particles become closer to each other and the augmenting path gets shorter. Using this approach, we can prove that with suitable probability the length of the *shortest* augmenting path shrinks after each round. When an augmenting path becomes an edge (that is, a path

of length one), and if the extreme unmatched nodes select each other as partners, the particles and the path vanish, yielding an increment to the size of the matching. Hence, a key step of our proof is to lower bound the probability that an augmenting path of length k vanishes, and then to apply Lemma 1 to relate the existing augmenting paths and the matching size.

We remark that the random process governing the movement of the particles in the network is not a classical random walk over the nodes of the graph. Indeed, if that were the case, a *maximum* matching would always be reached in polynomial time by a simple cat-and-mouse argument. Instead, a random move of a particle depends on the current matching, which in turn changes when the particle moves. This modest difference can lead to an exponential time gap between convergence to an approximate matching and convergence to a maximum matching. Indeed, exploiting the dependence of the particles' movements on the current matching, we show that there is a family of graphs for which the PRUDENCE algorithm takes exponentially many rounds with high probability to reach a maximum matching starting from a set of configurations that cover almost all possible cases. This family of “bad” graphs is defined as follows (see also Figure 2.7).

Definition 1 (Bad graph G_n) *The bipartite graph $G_n = (A \cup B, E)$ has $4n$ nodes $A = \{a_1, \dots, a_{2n}\}$ and $B = \{b_1, \dots, b_{2n}\}$, and its edges are (a_{n+1}, b_n) , (a_i, b_j) for all $1 \leq i \leq n$ and $1 \leq j \leq i$, and (a_i, b_j) for all $n+1 \leq i \leq 2n$ and $n+1 \leq j \leq i$.*

Note that the set of “horizontal” edges (a_i, b_i) , for $1 \leq i \leq 2n$ is the unique perfect matching for G_n .

Theorem 3 *The PRUDENCE algorithm requires $2^{\Omega(n/\log^2 n)}$ many rounds with high probability to reach the perfect matching when starting from any $(2n - 1)$ -matching in which the two unmatched nodes are in opposite sides of G_n .*




Figure 2.7. The bad graph. The “bad” graph G_n for $n = 3$. One of the “bad” matchings of Theorem 3 is highlighted in red.

The proof of Theorem 3 is in Section 2.7. The main idea of the proof is to track the positions of the unmatched nodes throughout the course of the algorithm and to lower bound the number of rounds needed before they meet as an adjacent pair.

We first prove a one-to-one correspondence between the Markov process of the state evolution between matchings and a classical random walk on a tree (represented in Figure 2.8) whose size is exponential in n . We show that this classical random walk takes exponential time to reach the root of the tree starting at any one of its nodes, thus providing a lower bound on the convergence time of the PRUDENCE algorithm.

We say that a matching M of G_n of size $2n - 1$ is *bad* if the PRUDENCE algorithm requires exponentially many rounds with high probability to converge to the perfect matching when starting from M . Observe that all matchings considered by Theorem 3 are bad. The following theorem states that *almost all* matchings of size $2n - 1$ are bad.

Theorem 4 *The ratio between the number of “bad” matchings and the number of all $(2n - 1)$ -matchings of G_n is $1 - O(2^{-n})$.*

Theorems 3 and 4 show that the PRUDENCE algorithm requires exponentially many rounds to converge to the perfect matching of G_n when starting from a set of

configurations (the bad matchings) constituting almost all possible cases (the matchings of size $2n - 1$).

2.3.2 Validation

Figure 3.7 compares the performance of the human subjects (red) with that of simulations (blue) on a set of 16-node networks (8 bipartite networks and 8 non-bipartite networks) with unique perfect matchings. The networks are sorted by increasing average completion time, and as a result bipartite networks are labeled from 1 to 8, while non-bipartite networks are labeled from 9 to 16. Each of these networks was tested at least 6 times over all sessions. The vertical axis represents the time (in seconds), and the numerical values of the convergence time of the algorithm are scaled by a constant factor to best match the experimental data.

In an additional experimental session, we tested twenty four subjects connected over small-world, preferential attachment and ring networks as well as over the “bad” graph G_n . The games on the bad graph were never solved, consistent with the prediction of exponentially slow convergence. Furthermore, we found that preferential attachment networks were more difficult to solve than small-world networks (one-tailed Welch’s t -test, p -value < 0.01). Figure 2.4 shows the affinity between humans’ (red) and algorithm’s (blue) performance, on this set of 24-node networks: small-world networks (triangles), ring network (diamonds), preferential attachment networks (circles). The x -axis shows the indices of the networks sorted by increasing average time to find the perfect matching, and the y -axis shows the average time.

Figure 2.5 shows, by simulation, that the algorithm scales linearly in the size of the network in the case of small-world networks [215], while it scales polynomially for preferential attachment networks [26, 33], and exponentially on the “bad” graph G_n . These results closely resemble the experimental data of the coloring games performed

by Kearns et al. [124], where preferential attachment networks resulted in the worst performance among all tested networks, while small-worlds networks appeared to be much easier to solve.

Figure 2.6 shows the performance of the experimental subjects on networks of 24 nodes, each admitting a perfect matching. In particular, it reports results for single games, and it compares the time to reach a perfect matching of size 12 (red), an approximate matching of size 11 (a 0.92-approximate matching, in blue) and a matching of size 6 (a 1/2-approximate matching, in green) in each game. The x -axis shows the indexes of the games sorted by increasing solution time, while the y -axis shows time in seconds. The plot shows (consistent with the theoretical analysis) that a 1/2-approximate matching is reached almost immediately in all games, an almost maximum matching is reached quickly, while reaching a perfect matching can take a large amount of time.

2.4 Discussion

While it is challenging to characterize the strategies used by humans in performing even simple social tasks, as they may depend on diverse individual cognitive and psychological attitudes, we argue that it is possible to isolate simple behavioral invariants of individual behavior, which are useful for algorithmic modeling, analysis and prediction of collective dynamics of coordination.

To illustrate our approach, we have focused on a simple matching game over networks and presented a combination of theoretical, experimental, and simulation results. From the experiments, we identified the prudence property as a common behavioral invariant of human subjects when they coordinate to find a maximum matching. We proposed an algorithm as model of human behavior and showed that it can successfully predict dynamics of coordination.

We have shown that our approach is able to uncover basic behavioral properties

that may not be apparent from off-line surveys. Indeed, when subjects were asked to report on their strategies in post-experimental surveys, we obtained a list of diverse strategies, including: choose a partner and never disengage from it, always accept proposals from neighbors, try to change partner if the game is not solved for a while. Moreover, our results demonstrate that algorithmic modeling and the mathematical analysis of algorithms can be useful in systematically predicting the aggregate behavior and in deriving results that hold for any graph, or for a large family of graphs. This general conclusions cannot be derived rigorously from experimental observations and computer simulations.

Our work suggests further research in several directions. A natural question is whether non-prudent behavior by a subset of the nodes can help. In a preliminary investigation, we have evaluated the performance of a variant of our algorithm where a subset of nodes behave non-prudently with a positive probability. In our simulations, these populations do not offer significant improvement in terms of finding a maximum matching. Furthermore, populations entirely composed of non-prudent nodes seem to perform poorly. In other words, a group of aggressive and risk-taking individuals might not achieve coordination easily.

Our PRUDENCE algorithm is memoryless. It is an interesting question as to what extent human subjects use memory in distributed games, and how memory could be incorporated in modeling human strategies. In an initial attempt to study this, we implemented a variant of the PRUDENCE algorithm in which a node remembers its recent history and gives less preference to neighbors who recently rejected it. In simulations on preferential attachment and small world networks, memory did not result in significant improvement over the memoryless case. Furthermore, simulations show that making decisions based on events in a distant past (that is, tracking events that happened in a distant past) might hurt performance. A careful investigation of the role of memory in

human strategies in distributed games is of fundamental interest.

Regarding the incentives, in our matching games each subject obtains the same reward when a maximum matching is reached, regardless of the chosen partner. How does the introduction of preferences affect the overall coordination? Preferences could be “enforced” for example by rewarding subjects based on the partners they match with. There is likely to be a trade-off between the collective task of finding a maximum matching and the individual profit maximization.

As a final remark, the proposed PRUDENCE algorithm constitutes a *possible* reasonable explanation of human coordination behavior in the distributed matching game. Apart from the simple variations mentioned above, we did not test how well other alternative algorithmic models could fit the experimental data.

2.5 Proof of Theorem 1

For ease of presentation, we assume $p = 0$, and remark that this result holds for all choices of $0 \leq p < 1$. Let G be a graph of n nodes and maximum degree Δ . Let m be the number of matched nodes in the smallest maximal matching of G . For $t \geq 0$, denote by W_t the set of nodes of G which are unmatched and have at least an unmatched neighbor at the beginning of round t , and let $|W_t|$ be the cardinality of W_t . Also, let M_t be the matching of G obtained by the PRUDENCE algorithm at the beginning of round t and N_t be the number of nodes matched by M_t . For $t \geq 0$, define the random variable

$$D_t = m - N_t.$$

We devote the rest of the proof to showing that

$$E[D_t] \leq (1 - (\Delta + 1)^{-3})^t E[D_0] \tag{2.1}$$

The theorem then follows by the observations that $E[D_0] \leq n$ and that any maximal matching is at least a $1/2$ -approximation of the maximum matching.

To prove (2.1), we will first show that $E[B_t(W_t)|W_t] \geq (\Delta + 1)^{-3}|W_t|$, where $B_t(W_t)$ is the number of nodes in W_t that match with nodes in W_t during round t (here the expectation is taken over the randomness of the algorithm during round t). For $u \in W_t$, let $Z_t(u)$ be the indicator random variable that takes value 1 if and only if u gets matched with a node in W_t during round t . By linearity of expectation, we have that

$$E[B_t(W_t)|W_t] = \sum_{u \in W_t} E[Z_t(u)] = \sum_{u \in W_t} \Pr(Z_t(u) = 1).$$

Let A_t be the set of nodes $u \in W_t$ such that (i) u has no incoming or outgoing request to nodes in W_t , and (ii) all neighbors $v \in W_t$ of u have an incoming request. Let $\bar{A}_t = W_t \setminus A_t$. For $u \in A_t$, we have that $\Pr(Z_t(u) = 1) = 0$, as unmatched nodes prefer neighbors who requested them over other unmatched neighbors. On the other hand, for $u \in \bar{A}_t$, we have $\Pr(Z_t(u) = 1) \geq \Delta^{-2}$. To see this, note that a pending request involving u (if any) will be honored with probability at least Δ^{-2} ; if no such request exists, the co-occurrence of the event of u requesting a neighbor with no incoming request and of that neighbor requesting u happens with probability at least Δ^{-2} . By definition of A_t , no two nodes in A_t can be neighbors. Also, by definition of W_t , every node $u \in W_t$ has at least one neighbor in W_t . These two facts imply that $|\bar{A}_t| \geq (\Delta + 1)^{-1}|W_t|$. We can conclude that $E[B_t(W_t)|W_t] \geq (\Delta + 1)^{-3}|W_t|$.

We now relate D_{t+1} to $B_t(W_t)$. First, note that $D_{t+1} \leq D_t - B_t(W_t)$. By itself, this bound is not strong as W_t can be small. However, when W_t is small, the current matching must be close to a maximal matching. Indeed, by considering the union of M_t and any maximal matching of W_t , we have that $m \leq N_t + |W_t|$. This implies that

$D_t = m - N_t \leq |W_t|$ and hence $D_{t+1} \leq D_t - B_t(W_t) \leq |W_t| - B_t(W_t)$. Therefore, we have

$$D_{t+1} \leq D_t - B_t(W_t),$$

$$D_{t+1} \leq |W_t| - B_t(W_t).$$

By taking the expectations with respect to the randomness of the algorithm during round t , we get

$$E[D_{t+1}|W_t, D_t] \leq D_t - E[B_t(W_t)|W_t] \leq D_t - (\Delta + 1)^{-3}|W_t|,$$

$$E[D_{t+1}|W_t, D_t] \leq |W_t| - E[B_t(W_t)|W_t] \leq |W_t| - (\Delta + 1)^{-3}|W_t| = (1 - (\Delta + 1)^{-3})|W_t|.$$

Now, by taking the expectation with respect to the randomness of the algorithm during rounds up to t , we obtain

$$E[D_{t+1}] \leq E[D_t] - (\Delta + 1)^{-3}E[|W_t|],$$

$$E[D_{t+1}] \leq (1 - (\Delta + 1)^{-3})E[|W_t|].$$

Letting $d_t = E[D_t]$, $w_t = E[|W_t|]$, and $\alpha = (\Delta + 1)^{-3}$, the bounds above can be rewritten as

$$d_{t+1} \leq \min \{d_t - \alpha w_t, (1 - \alpha)w_t\}.$$

To conclude the proof of (2.1), we show by induction that $d_t \leq d_0(1 - \alpha)^t$. For $t = 0$, as $d_0 \leq w_0$, we have $d_1 \leq d_0 - \alpha w_0 \leq (1 - \alpha)w_0$. Now, let us consider any $t \geq 1$ and distinguish between the cases of $w_t \leq d_0(1 - \alpha)^t$ and $w_t > d_0(1 - \alpha)^t$. If $w_t \leq d_0(1 - \alpha)^t$, we have $d_{t+1} \leq (1 - \alpha)w_t \leq d_0(1 - \alpha)^{t+1}$. Otherwise, if $w_t > d_0(1 - \alpha)^t$, using the

induction hypothesis, we have that

$$d_{t+1} \leq d_t - \alpha w_t \leq d_0(1 - \alpha)^t - \alpha w_t \leq d_0(1 - \alpha)^t - d_0\alpha(1 - \alpha)^t = d_0(1 - \alpha)^{t+1},$$

which completes the proof.

2.6 Proof of Theorem 2

For ease of presentation, we assume $p = 0$, and remark that this result holds for all choices of $0 \leq p < 1$. Let G be a graph of n nodes, maximum degree Δ , and maximum matching of size OPT . We will consider the unmatched nodes as *particles* randomly moving on the nodes of the network as per the algorithm choices. To see how the particle move, consider the particle positioned at any unmatched node u . If u requests a matched neighbor v and v accepts the requests, then the particle will move to v 's old partner (which is left unmatched). If u requests an unmatched neighbor z and z accepts the request, then both the particles at u and z will dissolve. Note that when two particles dissolve the size of the matching increases by one.

An augmenting path is a path of odd length which alternates matched and unmatched edges and whose extreme edges are unmatched. Observe that by switching each unmatched edge of an augmenting path into a matched edge, and viceversa, the size of the matching increases by one.

We split the rounds into epochs of $\lfloor 1/\varepsilon \rfloor$ rounds each. We claim that if at the beginning of any epoch the size of the matching is less than a $(1 - \varepsilon)\text{OPT}$, then the size of the matching increases by at least one by the end of that epoch with probability at least $\Delta^{-2/\varepsilon}$. To prove the claim, consider the first round of any epoch and let u_0, u_1, \dots, u_ℓ be any *shortest* augmenting path at the beginning of that round. It must be that $\ell < 2(\varepsilon^{-1} - 1)$, otherwise Lemma 1 would imply that the size of the matching is at least a

$\frac{\ell+1}{\ell+3} \geq 1 - \varepsilon$ fraction of OPT. For $\ell = 1$, u_0 and u_1 will match with each other during the first round with probability at least Δ^{-2} , hence the claim is true. For $\ell = 3$, u_0 and u_3 will request respectively u_1 and u_2 with probability at least Δ^{-2} during the first round of the epoch, and these requests will be accepted in the second round with probability at least Δ^{-2} — hence, the size of the matching increases by one within 2 rounds with probability at least Δ^{-4} . Now consider $5 \leq \ell < 2(\varepsilon^{-1} - 1)$. We have that two particles occupy the nodes u_0 and u_ℓ at the extremes of the augmenting path. With probability at least Δ^{-2} , u_0 requests to match with u_1 during the first round and u_1 accepts in the second round, making the corresponding particle move from u_0 to u_2 . A similar argument yields that the particle at u_ℓ moves to $u_{\ell-2}$ within two rounds with probability at least Δ^{-2} . Moreover, as the augmenting path under consideration is a shortest augmenting path, nodes $u_2, \dots, u_{\ell-2}$ have no unmatched neighbors at the beginning of the first round and hence do not receive any matching request during that round. Therefore, with probability at least Δ^{-4} , at the end of the second round the nodes u_2 and $u_{\ell-2}$ are unmatched whereas nodes $u_3, \dots, u_{\ell-3}$ did not change their partner. That is, the length of the shortest path at the beginning of the third round of the epoch is at most $\ell - 4$ with probability at least Δ^{-4} . By means of the same argument, we can conclude that with probability at least $(\Delta^{-4})^{\ell/4} > \Delta^{-2/\varepsilon}$, all nodes in an augmenting path are matched within $\ell/2 \leq \lceil 1/\varepsilon \rceil$ rounds, which proves the claim.

For any epoch i , we now associate a binary random variable X_i which takes on value 1 with probability $p = \Delta^{-2/\varepsilon}$. The claim guarantees that the size of the matching after T epochs is at least $\min\{(1 - \varepsilon)\text{OPT}, \sum_{i=1}^T X_i\}$. Also, as successive rounds of the algorithm are independent, the X_i 's are independent random variables. For any $0 < \delta \leq 1$, the Chernoff bound states that

$$\Pr \left[\sum_{i=1}^T X_i < (1 - \delta)Tp \right] < \exp(-Tp\delta^2/2).$$

For any $c \geq 1/2$, by setting $T := cn\Delta^{-2/\varepsilon}$ and $\delta := \varepsilon$, the above yields that the size of the matching after T epochs (i.e., after $T \lfloor 1/\varepsilon \rfloor \leq \frac{c}{\varepsilon} n\Delta^{2/\varepsilon}$ rounds) is at least $\min\{(1 - \varepsilon)\text{OPT}, (1 - \varepsilon)cn\} = (1 - \varepsilon)\text{OPT}$ with probability at least $1 - \exp(-c\varepsilon^2 n/2)$.

2.7 Proof Theorem 3

2.7.1 Analysis

For ease of presentation, we assume $p = 0$, and remark that this result holds for all choices of $0 \leq p < 1$. We say that the nodes $\{a_i : 1 \leq i \leq n\} \cup \{b_i : 1 \leq i \leq n\}$ constitute the upper half of G_n , and the remaining ones constitute the lower half of G_n . Let $\mathcal{M} = \mathcal{M}_1 \cup \mathcal{M}_2$ be the set of all matchings of G_n of size $2n - 1$, where \mathcal{M}_1 is the set of matchings of size $2n - 1$ in which the two unmatched nodes are in opposite halves of G_n , and $\mathcal{M}_2 = \mathcal{M} \setminus \mathcal{M}_1$ are the remaining ones.

Our goal is to show that the PRUDENCE algorithm requires $2^{\Omega(n/\log^2 n)}$ rounds with high probability to reach the perfect matching of G_n when starting from any matching in \mathcal{M}_1 . We first prove certain properties for the matchings in \mathcal{M}_1 . We then establish a correspondence between the Markov chain over matchings induced by the PRUDENCE algorithm and a classical random walk on the tree T_n^* . In particular, we show that the hitting time of the root of T_n^* is a lower bound on the number of rounds to reach the perfect matching of G_n .

2.7.2 Properties of matchings in \mathcal{M}_1

We begin by characterizing the matchings in \mathcal{M}_1 .

Lemma 2 *Consider any matching $M \in \mathcal{M}_1$, and let a_k, b_ℓ be the unmatched nodes in the upper and lower half of G_n , respectively. Then, the following properties hold:*

1. *For all $i < k$ and $i > \ell$, the matching M contains the edges (a_i, b_i) .*

2. If $k < n$, M contains the edge (a_n, b_j) for some $1 \leq j < n$. Similarly, if $\ell > n + 1$, M contains the edge (a_i, b_{n+1}) for some $n + 1 < i \leq 2n$. That is, the nodes a_n and b_{n+1} can be matched only through non-horizontal edges.
3. If in its upper half M contains a pair of edges $(a_{i_1}, b_{j_1}), (a_{i_2}, b_{j_2})$ with $i_1 \neq j_1$, $i_2 \neq j_2$, and $1 \leq i_1 < i_2 \leq n$, then $1 \leq k \leq j_1 < i_1 \leq j_2 < i_2 \leq n$. Similarly, if in its lower half M contains a pair of edges $(a_{i_1}, b_{j_1}), (a_{i_2}, b_{j_2})$ with $i_1 \neq j_1$, $i_2 \neq j_2$, and $n + 1 \leq j_1 < j_2 \leq 2n$, then $n + 1 \leq j_1 < i_1 \leq j_2 < i_2 \leq \ell \leq 2n$. That is, non-horizontal matching edges do not cross.

Proof. To prove the first property, we show that $(a_i, b_i) \in M$ for all $i < k$ (the claim for $i > \ell$ can be proven in the same way). We show by induction on $1 \leq j \leq k - 1$ that $(a_i, b_i) \in M$ for all $i \leq j$. For $j = 1$, we have that a_1 must be matched to b_1 (its only neighbor), and therefore the claim holds true. Suppose the claim holds true for some $j < k - 1$. By the inductive assumption we have that $(a_i, b_i) \in M$ for all $i \leq j$. As $(a_{j+1}, b_i) \in E$ if and only if $i \leq j + 1$, a_{j+1} must be matched to b_{j+1} , and therefore the claim holds for $j + 1$.

The second property follows by observing that $M \in \mathcal{M}_1$ implies that the bridge edge (a_{n+1}, b_n) is in M , and therefore a_n cannot be matched to b_n , and a_{n+1} cannot be matched to b_{n+1} in M . To see this, suppose by contradiction that $(a_{n+1}, b_n) \notin M$. Then, b_n must be matched to a_n (its only neighbor besides a_{n+1}), and a node in $\{a_1, \dots, a_{n-1}\}$ is unmatched. Then, each of the $n - 1$ nodes in $\{b_1, \dots, b_{n+1}\}$ must be matched with one of the $n - 2$ matched nodes in $\{a_1, \dots, a_{n-1}\}$, generating a contradiction. This implies that $(a_{n+1}, b_n) \in M$.

To prove the third property, assume that, in its upper half, M contains edges $(a_{i_1}, b_{j_1}), (a_{i_2}, b_{j_2})$ with $i_1 \neq j_1$, $i_2 \neq j_2$, and $1 \leq i_1 < i_2 \leq n$. Then, it must be that $j_1 < i_1$ and $j_2 < i_2$. Moreover, Property 1 implies that $k \leq j_1$. Therefore, it only remains to

show that $i_1 \leq j_2$. Suppose by contradiction that $i_1 > j_2$. As $i_1 > j_1 \geq k$ and $j_1 \neq j_2$, it must be that $i_1 \geq k + 2$. As b_{j_2} is matched to a_{i_2} and $i_2 > i_1$, we have that each of the $i_1 - k \geq 2$ nodes in $\{a_{k+1}, \dots, a_{i_1}\}$ must be matched to one of the $i_1 - k - 1$ nodes in $\{b_k, \dots, b_{i_1-1}\} \setminus \{b_{j_2}\}$, generating a contradiction. This implies that $i_1 \leq j_2$ and therefore $1 \leq k \leq j_1 < i_1 \leq j_2 < i_2 \leq n$. The claim in Property 3 regarding the lower half of M is similarly proved. \square

It follows from Lemma 2 that a matching $M \in \mathcal{M}_1$ can be uniquely reconstructed by specifying the two unmatched nodes and the nodes in $\{a_1, \dots, a_n\} \cup \{b_{n+1}, \dots, b_{2n}\}$ whose matching edges are *non-horizontal*. To see this, consider the upper half of G_n : assume $a_{j_0} \neq a_n$ is the unmatched node and $S = \{j_1, \dots, j_m\}$, with $1 \leq j_0 < j_1 < j_2 < \dots < j_m = n$, is the set of indexes of the left nodes whose matching edges are non-horizontal. (Note that $n \in S$ by Lemma 2.) Then, $j_0 < j_1$ and $(a_i, b_i) \in M$ for all i such that $i \notin S \cup \{j_0\}$ and $1 \leq i \leq n$. Hence, it necessarily holds that $(a_{j_i}, b_{j_{i-1}}) \in M$ for all $1 \leq i \leq m$. This completes the construction of the matching in the upper half of G_n . A similar argument can be applied to the lower half. These two arguments imply the following lemma.

Lemma 3 *There exists a bijection ψ between matchings in \mathcal{M}_1 and elements of $\mathcal{P} \times \mathcal{P}'$, where*

$$\mathcal{P} = \{(x, S) : x \in \{1, \dots, n-1\}, \{n\} \subseteq S \subseteq \{x+1, \dots, n\}\} \cup \{(n, \emptyset)\},$$

$$\mathcal{P}' = \{(y, S') : y \in \{n+1, \dots, 2n\}, \{n+1\} \subseteq S' \subseteq \{n+1, \dots, y-1\}\} \cup \{(n+1, \emptyset)\}.$$

2.7.3 The tree T_n^*

Definition 2 *Let T_1 be a labelled rooted tree with a singleton node with label 1. Inductively, for $2 \leq i \leq n-1$, let T_i be the labelled rooted tree whose root is labelled with i and its children are T_1, \dots, T_{i-1} . We define T_n^* to be the tree with an unlabelled root whose*




Figure 2.8. Tree T_n^* . Tree T_n^* with labels, for $n = 6$.

only child is T_n (also see Figure 2.8). Let r^* denote the root of T_n^* .

We show that the hitting time of r^* when starting at any node $u \neq r^*$ is exponential with high probability. For a node $u \neq r^*$, we call the edge that connects u to its parent u 's *exit edge*. For any subtree $T_i \subset T_n^*$, let Z_i be the random variable denoting the number of steps that it takes for a walk starting at the root of T_i to “exit” T_i . The following lemma provides an exponential lower bound on Z_i

Lemma 4 *There exist positive constants $\alpha, \gamma > 0$ such that, for all $i \geq 2$,*

$$\Pr[Z_i \geq \gamma \cdot 2^{i/(\alpha \log^2 i)}] \geq 1 - \frac{1}{\log i}.$$

Proof. We proceed by induction on i . For convenience, define $g(i) = \alpha \log^2 i$ and $f(i) = \gamma \cdot 2^{i/g(i)}$ for some $\alpha, \gamma > 0$. For any constant $\alpha > 0$, there exists a small enough constant $\gamma > 0$ such that $f(i) \leq 1$; therefore, as $Z_i \geq 1$ with probability 1, the claim holds trivially for any $i \leq i^*$, where i^* is a suitable large constant.

Now consider any $i \geq i^*$ and suppose the claim holds up to $i - 1$. Every time the walk is on the root of T_i , it exits T_i with probability $1/i$. Therefore, letting E_t be

the event that the first t times the walk is on the root of T_i it does *not* exit T_i , we have $\Pr[E_t] \geq 1 - t/i$. Let $t = i/(2\log i)$, and let D_j , $1 \leq j \leq t$, be the event that, when it is on the root of T_i for the j -th time, the walk moves to the root of one of the subtrees $T_{i-g(i)}, \dots, T_{i-1}$ and takes at least $f(i-g(i))$ steps to exit that subtree. For $1 \leq j \leq t$, we have

$$\begin{aligned} \Pr[D_j | E_t] &\geq \frac{g(i)}{i} \cdot \Pr[Z_{i-g(i)} \geq f(i-g(i))] \\ &\geq \frac{g(i)}{i} \cdot \left(1 - \frac{1}{\log(i-g(i))}\right), \end{aligned}$$

by the induction hypothesis on $Z_{i-g(i)}$. Letting χ_j be the indicator function of the event D_j for $1 \leq j \leq t$, the probability that at least two of the events D_j happen, given E_t , is lower bounded by:

$$\Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] \geq \Pr \left[\sum_{j=1}^{t/2} \chi_j \geq 1, \sum_{j=t/2+1}^t \chi_j \geq 1 \mid E_t \right] = \Pr \left[\sum_{j=1}^{t/2} \chi_j \geq 1 \mid E_t \right]^2.$$

By union bound, we can write

$$\begin{aligned} \Pr \left[\sum_{j=1}^{t/2} \chi_j \geq 1 \mid E_t \right] &\geq 1 - \prod_{i=1}^{t/2} (1 - \Pr[D_j | E_t]) \\ &\geq 1 - \left(1 - \frac{g(i)}{i} \left(1 - \frac{1}{\log(i-g(i))}\right)\right)^{t/2} \\ &\geq 1 - \exp \left[-\frac{\alpha \log i}{4} \left(1 - \frac{1}{\log(i-g(i))}\right) \right] \\ &\geq 1 - \frac{1}{i^{\alpha/8}}, \end{aligned}$$

where the last step holds for i sufficiently large so that $\log(i-g(i)) \geq 2$. This implies that

$$\Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] \geq \left(1 - \frac{1}{i^{\alpha/8}}\right)^2 \geq 1 - \frac{2}{i^{\alpha/8}}.$$

Therefore, we conclude that

$$\begin{aligned} \Pr[Z_i \geq 2 \cdot f(i - g(i))] &\geq \Pr \left[\sum_{j=1}^t \chi_j \geq 2 \right] \geq \Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] \Pr[E_t] \\ &\geq \left(1 - \frac{2}{i^{\alpha/8}} \right) \left(1 - \frac{t}{i} \right) \geq 1 - \frac{1}{\log i}, \end{aligned}$$

where the last step holds by choosing α sufficiently large. The claim now follows since $2 \cdot f(i - g(i)) \geq f(i)$. \square

Note that any random walk starting at any node $u \neq r^*$ has to exit T_n before hitting r^* . Therefore, an application of Lemma 4 to T_n yields a lower bound to the hitting time of r^* when starting at any node $u \neq r^*$.

Corollary 1 *The hitting time of r^* of a random walk starting at any node $u \neq r^*$ is $2^{\Omega(n/\log^2 n)}$ with high probability.*

2.7.4 Completing the proof of of Theorem 3

For $t \geq 0$, let $\mathcal{M}(t)$ be the matching at the beginning of round t and assume $\mathcal{M}(0) \in \mathcal{M}_1$. To analyze the convergence to a perfect matching, we will consider on the event that $\mathcal{M}(t) \notin \mathcal{M}_1$. Note that in order for this event to happen, the bridge edge (a_{n+1}, b_n) of G_n will have to swap out of the matching. Let $E(t)$ be the event that a_n requests b_n during round t . Similarly, let $E'(t)$ be the event that b_{n+1} requests a_{n+1} during round t . Define the random variables

$$\begin{aligned} \tau_n &= \min\{t : E(t) \text{ happens}\}, \\ \tau'_n &= \min\{t : E'(t) \text{ happens}\}, \\ \tau_n^* &= \min\{\tau_n, \tau'_n\}. \end{aligned}$$

Then τ_n^* is a lower bound on the number of rounds to reach the perfect matching. Lemma 5 below states that, for some $c > 0$,

$$\Pr \left[\tau_n \leq 2^{cn/\log^2 n} \mid \tau_n \leq \tau'_n \right] = o(1)$$

and

$$\Pr \left[\tau'_n \leq 2^{cn/\log^2 n} \mid \tau'_n \leq \tau_n \right] = o(1).$$

Then the main theorem follows as

$$\begin{aligned} \Pr \left[\tau_n^* \leq 2^{cn/\log^2 n} \right] &= \Pr \left[\tau_n^* \leq 2^{cn/\log^2 n} \mid \tau_n \leq \tau'_n \right] \Pr \left[\tau_n \leq \tau'_n \right] \\ &\quad + \Pr \left[\tau_n^* \leq 2^{cn/\log^2 n} \mid \tau'_n < \tau_n \right] \Pr \left[\tau'_n < \tau_n \right] \\ &= \Pr \left[\tau_n \leq 2^{cn/\log^2 n} \mid \tau_n \leq \tau'_n \right] \Pr \left[\tau_n \leq \tau'_n \right] \\ &\quad + \Pr \left[\tau'_n \leq 2^{cn/\log^2 n} \mid \tau'_n < \tau_n \right] \Pr \left[\tau'_n < \tau_n \right] \\ &= o(1). \end{aligned}$$

Lemma 5

$$\Pr \left[\tau_n \leq 2^{cn/\log^2 n} \mid \tau_n \leq \tau'_n \right] = o(1)$$

and

$$\Pr \left[\tau'_n \leq 2^{cn/\log^2 n} \mid \tau'_n \leq \tau_n \right] = o(1).$$

Proof. We will prove the first bound. The second one follows by symmetry. Conditioning on the event that $\tau_n \leq \tau'_n$, we will analyze the matching in the upper half of G_n induced by $\mathcal{M}(t)$. Since $\tau_n \leq \tau'_n$, $\mathcal{M}(t) \in \mathcal{M}_1$ as long as $E(t)$ does not happen. By Lemma 3, it is equivalent to study the Markov process $\{(X(t), \mathcal{S}(t)), t \geq 0\}$ over $\mathcal{P} \cup \{(\perp, \emptyset)\}$, where $(X(t), \mathcal{S}(t))$ is defined as the first marginal of $\psi(\mathcal{M}(t))$, and the additional state (\perp, \emptyset) is reached when the event $E(t)$ happens. That is, conditioning on

the event $\tau_n \leq \tau'_n$, it follows that

$$\tau_n = \min\{t : (X(t), \mathcal{S}(t)) = (\perp, \emptyset)\}. \quad (2.2)$$

If $\tau_n \leq \tau'_n$ and $(X(t), \mathcal{S}(t)) \neq (\perp, \emptyset)$, all the neighbors of the unmatched node in the upper half of G_n are matched at the beginning of round t , and hence are *equally likely* to be requested during round t . Therefore, the Markov process $(X(t), \mathcal{S}(t))$ has the following transition probabilities.

$$\Pr \left[(X(t+1), \mathcal{S}(t+1)) = (x', S') \mid (X(t), \mathcal{S}(t)) = (x, S) \neq (\perp, \emptyset), \tau_n \leq \tau'_n \right] = \frac{1}{x},$$

for any

$$(x', S') \in \begin{cases} \{(x'', S \cup x) : x'' < x\} \cup \{(\min(S), S \setminus \min(S))\}, & \text{if } x < n \text{ (and } S \neq \emptyset) \\ \{(x'', S \cup x) : x'' < x\} \cup \{(\perp, \emptyset)\}, & \text{if } x = n \text{ (and } S = \emptyset) \end{cases}$$

The case $(x', S') \in \{(x'', S \cup x) : x'' < x\}$ represents the scenario in which the unmatched node a_x requests a node through a non-horizontal edge: in this case, no progress is made as the unmatched node in the next round will be further away from a_n . If the unmatched node a_x requests the node on its horizontal edge, the next unmatched node will be closer to a_n . In the special case $(x, S) = (n, \emptyset)$, if the unmatched node requests the neighbor on its horizontal edge, then the bridge edge is swapped out of the matching and $\mathcal{M}(t+1) \notin \mathcal{M}_1$.

We will now show that the Markov chain $\{(X(t), \mathcal{S}(t)), t \geq 0\}$ is equivalent to the random walk on T_n^* . For a node v of T_n^* , let x_v be its label and S_v be the set of labels

of its ancestors. Define the function ϕ from nodes of T_n^* to states of the chain as follows:

$$\phi(v) = \begin{cases} (\perp, \emptyset), & v = r^* \\ (x_v, S_v), & v \neq r^* \end{cases}$$

It is easy to verify that ϕ is a bijection. Two nodes u and v are adjacent in T_n^* if and only if there is a nonzero transition probability between the states $\phi(u)$ and $\phi(v)$. To see this, suppose there is a nonzero transition probability from (x_u, S_u) to (x_v, S_v) in the Markov chain. Let $u = \phi^{-1}(x_u, S_u)$ and $v = \phi^{-1}(x_v, S_v)$ be the corresponding nodes in T_n^* . There are two cases: (a) if $x_v < x_u$ then $S_v = S_u \cup x_u$, and v is a child of u ; (b) if $x_v > x_u$ then $x_v = \min(S_u)$, $S_v = S_u \setminus \min(S_u)$, and v is the parent of u . The other direction is analogous. Therefore, conditioning on $\tau_n \leq \tau'_n$ and $(X(0), \mathcal{S}(0)) \neq (\perp, \emptyset)$, we can conclude that $\min\{t : (X(t), \mathcal{S}(t)) = (\perp, \emptyset)\}$ equals the hitting time of r^* for a random walk on T_n^* starting at the node $\phi^{-1}(X(0), \mathcal{S}(0)) \neq r^*$. The lemma follows by equation (2.2) and Corollary 3.

2.8 Proof of Theorem 4

As in the proof of Theorem 3, we let $\mathcal{M} = \mathcal{M}_1 \cup \mathcal{M}_2$ be the set of all matchings of G_n of size $2n - 1$, where \mathcal{M}_1 and \mathcal{M}_2 contain all matchings in which the two unmatched nodes are in opposite sides of G_n and in the same side of G_n , respectively. By Theorem 3 we know that starting from any matching in \mathcal{M}_1 requires exponentially many steps to reach the perfect matching of G_n with high probability. We will show that these matchings substantially make up for the whole \mathcal{M} . Indeed, we prove that

$$|\mathcal{M}_1| = 2^{2n-2}, \quad \text{and} \quad |\mathcal{M}_2| = 2^{n+1} - 2.$$

To compute the size of \mathcal{M}_1 , using Corollary 3 we have that

$$|\mathcal{M}_1| = \left(1 + \sum_{i=1}^{n-1} 2^{n-i-1}\right)^2 = \left(1 + \sum_{j=0}^{n-2} 2^j\right)^2 = (1 + (2^{n-1} - 1))^2 = 2^{2n-2}.$$

To compute the size of \mathcal{M}_2 , let \mathcal{M}'_2 contain the matchings of \mathcal{M}_2 in which the two unmatched nodes are in the upper half of G_n . Observe that by symmetry $|\mathcal{M}_2| = 2 \cdot |\mathcal{M}'_2|$. To determine the size of \mathcal{M}'_2 , note first that every matching in \mathcal{M}'_2 is such that the nodes in the lower half of G_n are matched through parallel edges, i.e. a_j is matched with b_j for every $n+1 \leq j \leq 2n$. Now consider all matchings in \mathcal{M}'_2 where a_k, b_ℓ are the two unmatched nodes, and observe that it must be that $1 \leq k \leq \ell \leq n$ (if not, we would have at least another unmatched node a_t with $t < \ell$). Also, note that for every $1 \leq j \leq k-1$ and every $\ell+1 \leq j \leq n$, it must be that a_j is matched with b_j . Hence, for $k = \ell$, there is a single matching. For $k < \ell$, we show that the remaining nodes can be matched in $2^{\ell-k-1}$ ways. To prove this, first observe that a_{k+1} can be matched to either b_{k+1} or b_k . Then, given the choice for a_{k+1} , a_{k+2} can be matched to either b_{k+2} or the node in $\{b_{k+1}, b_k\}$ which is not matched to a_{k+1} . Similarly, for $i+1 \leq j \leq \ell-1$, there are two possible choices for a_j given the choice for $\{a_{k+1}, \dots, a_{j-1}\}$. Finally, given the choices for $\{a_{k+1}, \dots, a_{\ell-1}\}$ there is only one possible match for a_ℓ , thus obtaining $2^{\ell-k-1}$ matchings with a_k, b_ℓ unmatched, $1 \leq k < \ell \leq n$. We can conclude that

$$\begin{aligned} |\mathcal{M}'_2| &= \sum_{k=1}^n \left(1 + \sum_{\ell=k+1}^n 2^{\ell-k-1}\right) = n + \sum_{k=1}^{n-1} 2^{-k} \sum_{\ell=k+1}^n 2^{\ell-1} \\ &= n + \sum_{k=1}^{n-1} 2^{-k} (2^n - 2^k) = 1 + \sum_{k=1}^{n-1} 2^{n-k} = 1 + (2^n - 2) = 2^n - 1. \end{aligned}$$

2.9 Acknowledgments

Chapter 2, in full, is a reprint of the material as it appears in PLoS ONE, 7(8) e41900, 2012, L. Coviello, M. Franceschetti, M. McCubbins, R. Paturi, and A. Vattani. The dissertation author was the primary investigator and co-author of this paper.

I would like to thank Devin Barr and Daniel Enemark for helping us with the experiments. I would also like to thank Joel Sobel for his feedback. I am indebted to Stephen Judd and Michael Kearns for allowing us to modify their software for running our experiments.

Chapter 3

An instance of distributed social computation: the multi-agent group membership problem

3.1 Introduction

We consider a distributed computation scenario in which there are agents of two types, *leaders* and *followers*. Each leader is equipped with the task to form a group of followers of a certain cardinality, by sending them requests. Followers can either accept or reject incoming leaders' requests. Each follower can be part of a single group at any time but can change group over time. Multiple followers can be part of a leader's group, but each leader can only recruit followers with whom it shares a communication link. These communication links are described by an arbitrary bipartite network, and we assume that each agent has knowledge of, and can interact with its neighbors over the network. In practice, the structure of the network can be dictated by physical or social constraints, see Figure 3.1. Leaders and followers share the *common goal* of reaching a state in which each leader formed a group of the right size, and we call *stable* such a state of "social welfare." We refer to this scenario as the *group membership problem*.

The contribution of the present work is twofold. First, we show that simple local rules of interaction lead to stable, or close to stable, group membership in reasonable




Figure 3.1. Example of a bipartite network between leaders and followers determined by physical constraints. Left: each leader can recruit the followers in its visibility range (dotted circle), arrows represent group membership, the set of arrows defines a partition of the followers into groups. Right: the resulting bipartite network. An edge between leader ℓ and follower f exists if and only if f is in ℓ 's visibility range. Matching edges are highlighted.

time, where by “close to stable” we mean that the total number of additional followers required to satisfy all group size constraints is an arbitrary small fraction of the entire population. Then, we show that such rules predict the performance of a group of human subjects solving the same group membership task in a laboratory setting, suggesting that multi-agent systems are useful to describe complex heterogeneous systems as human populations.

Regarding the first contribution, we propose a simple, distributed, memoryless algorithm in which leaders only pursue local stability, and we show that, in any network of size n , any constant approximation of a globally stable outcome (or of a suitably defined *best* outcome if a stable one does not exist) is reached in time polynomial in n with high probability. In other words, within an acceptable approximation, our algorithm is able to find a solution in feasible time on any instance of the problem. In contrast, we show that there exist networks requiring an exponential gap between the time needed to reach stability and that needed to reach approximate stability, that is, to find the *best*

solution compared to a *good* solution. In order to characterize what types of solution can be reached in feasible (i.e., polynomial) time, we do not restrict our attention to a subclass of problem instances to provide guarantees on the convergence to the optimal solution. Instead, we consider approximate solutions and derive a result that holds for any instance of the problem.

Regarding the second contribution, we created an artificial environment in which human subjects have to solve a group membership task on virtual networks of leaders and followers. We conducted 36 experiments of group membership on a pool of 10 different networks with 16 nodes each. In each experiment, participants controlled the nodes of a virtual network and interacted with their neighbors via the point-and-click interface shown in Fig. 3.2. In order to elicit the common goal of reaching stability, they received a monetary reward if they reached a stable state within a maximum time of 5 minutes. We observe a good fit between experimental data and the algorithm's predictions. On the one hand, the algorithm was able to predict which networks were the most difficult to solve by the human subjects. On the other hand, the human subjects always found good solutions quickly and spent most of the time attempting to improve to the optimum. These results suggest that, at least in the specific context of the group membership problem considered here, simple local rules of interaction are able to simulate complex global dynamics, and therefore tools from traditional computation theory can be used to study distributed social computation.

We point out that the idea of using simple interactions to predict global outcomes resulting from possibly complex and diverse microscopic effects is not new. The theme is recurrent in statistical physics and cellular automata [57, 141, 36, 211], but has yet to gain popularity in the context of social computation. Our reduction of social interaction to algorithmic modeling is also reminiscent of the work of Herbert Simon, who claimed that information processing is at the basis of human decision-making [194]. Finally, we refer

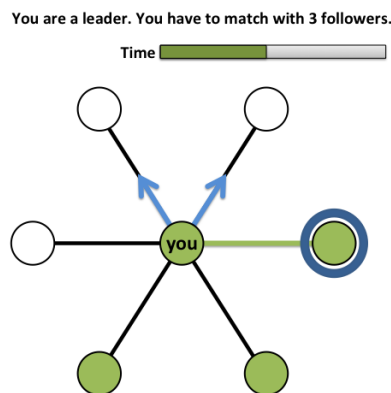


Figure 3.2. Computer interface of a participant playing the role of a leader. The participant controls the node in the middle (with the label “you”) and has to match with 3 followers (neighboring nodes). The green line indicates a matched pair with the follower on the right. Followers marked in green (on the bottom) are currently matched with other leaders. Blue arrows represent outgoing matching requests (to the followers on top). The bar on top shows the remaining time (out of the 5 minutes allowed).

to the influential work of Kearns and his collaborators [124, 113, 122, 114, 123, 121], who pointed out the need to study the principal mechanisms of social computation and strategic interaction over networks. Our experimental approach, based on a highly constrained laboratory setting, where language and other natural forms of communication are eliminated in favor of enforcing simple actions, follows closely this line of work. In addition, we advocate for simple computational models of individual behavior for predictive and explanatory purposes.

3.2 Related work

The study of distributed multi-agent coordination has received a great deal of attention by the control and computer science communities, particularly in settings where the agents perform simple local updates, do not have complete knowledge of the entire population, and communication between them is limited. On the other hand, in computational social science similar coordination problems have been considered with

the goal of providing a model of distributed dynamics of human networks.

Within the first group of studies, one of the main issues is whether distributed multi-agent coordination dynamics converge in finite time to a set of desirable configurations defined by notions of stability and optimality. For example, Roth and Vande Vate [187] considered two-sided marriage markets and showed that better- and best-response dynamics always reach stability in finite time. Bertsekas and Castanon [29, 30, 31] studied distributed dynamics for the assignment problem and proved their converge with probability one to optimal assignments. In our work, convergence to the set of optimal solutions always occurs in any instance of the problem. In particular, we define a potential function for our algorithm (that we call *deficit*) and a structural result guarantees that this decreases in a finite number of iterations (see Lemma 6).

Beside convergence in finite time, another important issue is convergence in *feasible* time. In this case, results are based on computer simulation [31, 22], or require specific modeling assumptions. For example, in the context of two-sided matching, best- or better-response dynamics are shown to converge to stability in polynomial time in the cases of global rankings [1], correlated markets [3] and geometric preferences [16], but might require exponential time in the general case (see for example [3]). In the context of distributed network coloring, Vattani et al. [107] proposed simple local dynamics that converge in polynomial time on any bipartite network. These works provide relevant insights about certain classes of problems, but their analysis is restricted by the specific assumptions they make to provide convergence guarantees.

An alternative approach, that allows to obtain results in more general scenarios, is to introduce a notion of *approximate* solution and to quantify the tradeoff between the quality of a solution and the time needed to reach it. The objective is to provide provable performance guarantees that hold for any instance of the problem, and to understand what types of configuration can be reached in practical time. In this context, Nedic

and Ozdaglar [160] studied the distributed optimization of sums of convex functions by agents who optimize their local objectives and exchange information locally, focusing on the tradeoff between solution accuracy and time. In the same spirit, our work considers a setup in which a globally stable solution is approximated by agents concerned by their local stability. In the context of bargaining over social exchange networks, Kanoria et al. [117] showed that ε -approximate Nash bargaining solutions are reached by a simple distributed algorithm in time polynomial in ε^{-1} and in the network size. In Chapter 2, we considered the problem of maximum matching and showed that a ε -approximate maximum matching is reached by simple local interaction in time increasing in ε^{-1} and polynomial in the network size. In the context of distributed consensus and averaging, Olshevsky and Tsitsiklis [165] showed that the number of iterations to convergence is polynomial in the number of agents and increases only logarithmically in the target accuracy ε^{-1} (where ε is the maximum allowed distance of an agent's opinion from consensus). Nax et al. [159] considered a formulation of the assignment problem in which limited information is available and can be exchanged between the agents, and proposed a simple distributed scheme for the agents' local updates. While the proposed algorithm is shown to always converge to optimal and stable allocations, the authors do not study its rate of convergence.

A different aspect of our work is the proposal of a simple distributed dynamical model to describe and predict the outcomes of groups of humans who have to coordinate over networks.

In computer science, human coordination has been studied under the premise that coordination constitutes the basis for *social computing* [205]. Following this approach, distributed collections of humans are tried to collectively solve traditional algorithmic tasks, such as coloring, consensus, and various forms of matching. For example, in the work of Kearns et al. [124], human subjects positioned at the vertices of a virtual network

were shown to be able to collectively reach a coloring of the network, given only local information about their neighbors. Other works further investigated human coordination in the case of coloring [67, 114, 151], consensus [114, 122], matching (see Chapter 2), bargaining and trade [117, 46, 113], and network formation [123]. Quoting Kearns [121], the main findings of research on experimental social computation up to date are the ability of humans to solve a wide range of tasks in a distributed fashion, the effect of the network structure on performance, with opposite effects for different tasks [114], and the emergence of behavioral characteristics of individuals [122]. However, the effectiveness of mathematical models of social computation to predict performance still needs to be assessed.

In this work, we address the question posed in [121] regarding the possibility of using simple models of social computation to predict the performance of humans on specific computation tasks over networks. Within an extremely wide design space, we focus on the distributed task of group membership, and use computational complexity and equilibrium concepts as the rigorous language to express these predictions.

3.3 The group membership task

We consider a bipartite network $G = (L \cup F, E)$ whose nodes are the disjoint sets L of leaders and F of followers, and where there exists an edge $(f, \ell) \in E$ between follower f and leader ℓ if and only if f and ℓ can communicate between each other (see Figure 3.1). Let $N_\ell = \{f \in F : (f, \ell) \in E\}$ be the neighborhood of $\ell \in L$. For each $\ell \in L$, leader ℓ has to form a group of c_ℓ followers from N_ℓ , where $c_\ell \geq 1$.

Definition 3 (Matching) *A subset $M \subseteq E$ is a matching of G if for each $f \in F$ there exists at most a single $\ell \in L$ such that $(\ell, f) \in M$.*

The definition of matching permits multiple followers to be part of a leader's group. There is a one-to-one correspondence between matchings M of G and groups $\{T_\ell(M) : \ell \in L\}$, where $T_\ell(M)$ denotes the group of leader ℓ under the matching M . We have that $T_\ell(M) = \{f \in F : (\ell, f) \in M\} \subseteq N_\ell$ for every matching M . Agents in $L \cup F$ are rewarded if each leader ℓ controls a team of c_ℓ followers, therefore we consider the following notion of stability.

Definition 4 (Stable matching) *Given constraints c_ℓ for each $\ell \in L$, a matching M of G is stable if and only if $|T_\ell(M)| = c_\ell$ for each $\ell \in L$.*

Having a local view of the network, each leader ℓ can only assess if “local stability” holds (i.e., if it is matched with c_ℓ followers), in contrast with the notion of “global stability” defined above.

Given the constraints c_ℓ , a network G might not admit a stable matching. Nonetheless, given a matching of G , we are interested in assessing its *quality*. Our main result builds on the following definitions of *deficit* of a leader and a matching.

Definition 5 (Deficit) *Let ℓ be a leader with constraint c_ℓ , and M be a matching of G . The deficit of ℓ under M is*

$$d_\ell(M) = c_\ell - |T_\ell(M)|.$$

The deficit of M is

$$d(M) = \sum_{\ell \in L} d_\ell(M) = \sum_{\ell \in L} (c_\ell - |T_\ell(M)|).$$

In words, $d_\ell(M)$ is the number of additional followers leader ℓ needs to satisfy its size constraint. Similarly, $d(M)$ sums the numbers of additional followers all leaders need to satisfy their size constraints. Given a matching M , we say that a leader ℓ is *poor* if $d_\ell(M) > 0$ (that is, $|T_\ell(M)| < c_\ell$) and *stable* if $|T_\ell(M)| = c_\ell$ (we exclude $|T_\ell(M)| > c_\ell$ assuming that matching with additional followers is costly). Observe that only poor

leaders contribute to $d(M)$, and that M is stable if and only if $d(M) = 0$. Given G , two matchings can be compared with respect to their deficit, and the best matching of G can be defined as one minimizing the deficit.

Definition 6 (Best matching) *A matching M of G is a best matching if $d(M) \leq d(M')$ for every matching M' of G .*

Observe that a stable matching is also a best matching, and that a best matching always exists for any network G and constraints c_ℓ . Moreover, if G admits a stable matching, $d(M)$ quantifies how much M differs from a stable matching of G . In general, if M^* is a best matching of G with $d(M^*) = d^*$, then, $d(M) - d^*$ tells how much M differs from a best matching of G . Given a matching M of G , the following definitions provide a measure of how well M approximates a best matching of G , or a stable matching (if one exists).

Definition 7 (Approximate best matching) *Fix $\varepsilon \in [0, 1]$, and let m be the number of followers in G . Let M^* be a best matching of G . Then, a matching M is a $(1 - \varepsilon)$ -approximate best matching of G if $d(M) - d(M^*) < \varepsilon m$.*

Definition 8 (Approximate stable matching) *Let G admit a stable matching. Fix $\varepsilon \in [0, 1]$, and let m be the number of followers in G . Then, a matching M is a $(1 - \varepsilon)$ -approximate stable matching of G if $d(M) < \varepsilon m$.*

3.4 The group membership algorithm

For ease of presentation, we assume that agents are synchronized. However, our results continue to hold also in the case of asynchronous agents (see discussion in Section 3.7). We assume that time is divided into rounds and each round is composed of two stages. In the first stage, each leader acts according to the algorithm in algorithm

box 2, and in the second stage each follower acts according to the algorithm in algorithm box 3.

First consider a leader ℓ , and let M be the matching at the beginning of a given round. If ℓ is poor (that is, $|T_\ell(M)| < c_\ell$) and $|T_\ell(M)| < |N_\ell|$ (that is, ℓ is not already matched with all followers in N_ℓ) then, with probability p (where $p \in (0, 1]$ is a fixed constant), ℓ attempts to match with an additional follower. We assume that leaders always prefer followers that are currently unmatched over matched ones. Note that a leader first checks if *local stability* holds (i.e., its group size is c_ℓ).

Consider now a follower f . During each round, if f has incoming requests then each is rejected independently of the others with probability $1 - q$ (where $q \in (0, 1]$ is a fixed constant). If all incoming requests are rejected, then f does not change group (if currently matched) or it remains unmatched (if currently unmatched). Otherwise, one among the active requests is chosen uniformly at random, f matches with the corresponding leader, and all the other requests are discarded. For ease of presentation, we assume that a follower is equally likely to accept a request when unmatched or matched, and that p and q are the same for all agents. Our results hold for more general choices of the parameters, that can vary between agents, as long as they remain bounded away from zero¹.

The proposed algorithm has the following desirable features aimed at modeling distributed social computation: agents have no memory of the past, decision are based only on local information, it is *self-stabilizing* (i.e. it stops when a stable matching is reached), the exchanged messages can be represented by a single bit², and each leader

¹ If each leader ℓ has a parameter p_ℓ and each follower f has a parameter q_f , our main result in Theorem 5 holds with $p = \min_\ell p_\ell$ and $q = \min_f q_f$ as long as these lower bounds are bounded away from zero.

² For example, when a follower accepts a matching request by leader ℓ , she might communicate it by sending a bit ‘1’ to leader ℓ and a bit ‘0’ to all other neighboring leaders. Our main result (Theorem 5) does not consider the total number of bits exchanged to reach a given approximation of the optimal solution.

only pursues local stability. The single invariant of the algorithm is that leaders prefer unmatched followers and pursue local stability. Followers, on the other hand, act in a randomized fashion and ensure exploration of the state space. Despite their simplicity, these simple rules allow to reach a good approximate solution and capture the collective behavior of the real human network. In practice, preferring unmatched followers appears to be a natural strategy, pursuing local stability is an inherent characteristic of human behavior –although subjects might not admit it explicitly in a survey– while randomization captures the diversity of the actions of the population as evidenced in the exit pools. As a remark, we consider an algorithm in which agents have no memory of the past for ease of analysis. We believe that allowing agents’ decisions to depend on the past actions (made by them and their neighbors) would not change our results.

ALGORITHM 2: Algorithm for leader $\ell \in L$.

```

if  $|T_\ell(M)| < \min\{c_\ell, |N_\ell|\}$  then
  with probability  $p$  do the following;
  if  $\exists$  unmatched  $f \in N_\ell$  then
    choose an unmatched follower  $f' \in N_\ell$  u.a.r.
  else
    choose a follower  $f' \in N_\ell \setminus T_\ell(M)$  u.a.r.
  end
end
send a matching request to  $f'$ ;

```

ALGORITHM 3: Algorithm for follower $f \in F$.

```

if  $f$  has incoming requests then
  for each leader  $\ell$  requesting  $f$  do
    with probability  $1 - q$  reject  $\ell$ 's request
  end
  if there are active requests then
    select one u.a.r. and join corresponding team;
    reject all other requests;
  end
end

```

3.5 Complexity results

For ease of presentation, we only consider networks admitting stable matchings and show that, given any network and any constant $\varepsilon \in (0, 1)$, a $(1 - \varepsilon)$ -approximate stable matching is reached in a number of rounds that is polynomial in the network size with high probability (Theorem 5). Then, we show through a counterexample that improving from approximate stability to stability might require time exponentially large in the network size (Theorem 6). Our results hold in general for reaching approximate best matchings.

Theorem 5 *Let G be a network with m followers and which admits a stable matching. Let $\Delta = \max_{\ell \in L} |N_\ell|$ be the maximum degree of the leaders. Fix $0 < \varepsilon < 1$, and let $c \geq 1 + \frac{1}{m(1-\varepsilon)}$. Then, a $(1 - \varepsilon)$ -approximate stable matching of G is reached within $c \lceil 1/\varepsilon \rceil (\Delta/pq)^{\lceil 1/\varepsilon \rceil} m$ rounds of the algorithm with probability at least $1 - e^{-cm\varepsilon^2/2}$.*

As an example, if Δ is constant in the network size, then one can choose $\varepsilon = 1/\log m$, and Theorem 5 implies that a $(1 - 1/\log m)$ -approximate stable matching is reached in at most $\mathcal{O}(m^2 \log m)$ rounds with probability that goes to one as $m \rightarrow \infty$.

In order to prove Theorem 5 (see Section 3.8) we introduce the notion of *deficit-decreasing* path, that in our setup plays the role of the augmenting path in the context of one-to-one matching. Since we consider bipartite networks, a path alternates leaders and followers.

Definition 9 (Deficit-decreasing path) *Given a matching M of G , a cycle-free path $P = \ell_0, f_1, \ell_1, \dots, f_{k-1}, \ell_{k-1}, f_k$ (of odd length $2k-1$) is a deficit-decreasing path relative to M if $(\ell_i, f_i) \in M$ for all $1 \leq i \leq k-1$, ℓ_0 is a poor leader, and f_k is an unmatched follower.*




Figure 3.3. Deficit-decreasing path. Top: a deficit-decreasing path of length 5: ℓ_0 is a poor leader, f_3 is an unmatched follower, and matching edges are highlighted. Bottom: the path is “solved” by turning each matched edge into an unmatched edge and vice versa; ℓ_0 obtains an additional follower (and its deficit decreases by one), ℓ_1 and ℓ_2 do not change their numbers of followers.

In words, a deficit-decreasing path starts at a poor leader with an edge not in M , ends at a follower that is not matched, and alternates edges in M and edges not in M . Observe that a new matching M' such that $d(M') = d(M) - 1$ is obtained by flipping each unmatched edge of a deficit-decreasing path P into a matched edge, and vice versa. This is shown in Figure 3.3. The proof of Theorem 5 is based on a technical lemma (see Section 3.9) that extends a previous combinatorial result by Hopcroft and Karp [103, Theorem 1]. Given a matching M with $d(M) \geq \varepsilon m$, we guarantee the existence of a deficit-decreasing path of length at most $2\lceil 1/\varepsilon \rceil$. Such a “short” path allows us to bound the number of rounds needed for a one-unit reduction of the deficit. The symmetric difference of two sets A and B is $A \oplus B = (A \setminus B) \cup (B \setminus A)$. Two paths are *follower-disjoint* if they do not share any follower (even though they might share some leader).

Lemma 6 *Let G admit a stable matching N . Let M be a matching of G with deficit $d(M) > 0$. Then, in $M \oplus N$ there are at least $d(M)$ follower-disjoint deficit-decreasing paths relative to M .*

We make use of Lemma 6 through the following corollary, which holds as the lengths of a set of follower-disjoint paths sum to at most $2m$.

Corollary 2 *Let G be a network with m followers, admitting a stable matching N . Let M be a matching of G with deficit $d(M) \geq \varepsilon m$, for some $\varepsilon > 0$. Then, in $M \oplus N$ there exists*

a deficit-decreasing path relative to M of length at most $2\lfloor 1/\varepsilon \rfloor - 1$.

As a remark, Corollary 2 and the observation that the deficit is non-increasing guarantee that our algorithm always converges to the set of optimal solutions in finite time in any instance of the problem³.

Theorem 5 gives a polynomial bound for reaching a $(1 - \varepsilon)$ -approximate stable matching for any constant $0 < \varepsilon < 1$ and any network. However, a polynomial guarantee cannot be derived for the case of a stable matching (that is, for $\varepsilon = 1/m$). To show this, we define a sequence of networks in which the number of rounds required to converge from an approximate matching M with $d(M) = 1$ to the stable matching is exponentially large in the network's size with high probability from an overwhelming fraction of the approximate matchings M such that $d(M) = 1$.

For $n \geq 1$, let $G_n = (L_n \cup F_n, E_n)$ be the network with leaders $L_n = \{\ell_1, \dots, \ell_n\}$, followers $F_n = \{f_1, \dots, f_n\}$, edges $E_n = \{(\ell_i, f_j) : 1 \leq i \leq n, j \leq i\}$, and group size constraints $c_\ell = 1$ for all $\ell \in L_n$, see Figure 3.4. G_n has a unique stable matching given by $M_n^* = \{(\ell_i, f_i) : 1 \leq i \leq n\}$.

Theorem 6 *For any matching M of G_n , let $\tau^*(M)$ denote the number of rounds to converge to the perfect matching when starting from M . Then, for any fixed constant $0 < \gamma < 1$, $\tau^*(M)$ is exponentially large in γn with high probability for a $1 - O(n2^{-(1-\gamma)n})$ fraction of all the matchings M such that $d(M) = 1$.*

³The deficit never increases over time as a leader never voluntarily disengages from a follower (see algorithm box 2), and a follower disengages from a leader only when she accepts a new matching request (see algorithm box 3). Convergence follows by considering the Markov chain whose state space is the set of all matchings. For each matching M , only matchings M' with $d(M') \leq d(M)$ can be reached from M , and Corollary 2 guarantees the existence of a finite sequence of transitions that lead from M to M' such that $d(M') < d(M)$ with finite probability.




Figure 3.4. The network G_n for $n = 6$.
The matching M'_n is highlighted.

Here we only provide the idea of the proof, whose details are presented in Section 3.10. To get an understanding of the algorithm's dynamics, consider the matching

$$M'_n = \{(l_i, f_{i-1}) : 2 \leq i \leq n\},$$

highlighted in Figure 3.4 for the case of $n = 6$. Observe that $d(M'_n) = 1$ and l_1 is poor. According to the algorithm, l_1 tries to match with f_1 . If f_1 accepts, then l_2 becomes poor (and tries to match with f_1 or f_2). After each round, there exists a unique poor leader until the stable matching is reached. The stable matching is reached when l_{n-1} (l_5 in Figure 3.4) becomes poor and matches with f_{n-1} (f_5 in Figure 3.4), and finally l_n matches with f_n . The stochastic process tracking the position of the poor leader is not a classical random walk and its transition probabilities at each time depend on the current matching. We show that convergence to stability requires a number of rounds that is exponential in n with high probability, and this holds for an overwhelming fraction of all matchings with $d(M) = 1$.

Fig. 3.5 shows the algorithm's average convergence time on the sequence of networks G_n (in logarithmic scale). The average number of rounds to reach a 0.9-approximate stable matching is upper bounded by a polynomial of small degree (bottom,




Figure 3.5. Algorithm’s performance on the hard networks G_n . The x -axis reports n , the y -axis reports the average time over 1000 simulations (in base-10 logarithm). The red line (top) shows the time to reach the stable matching; the blue line (bottom) shows the time to reach a $(1 - \varepsilon)$ -approximate matching for $\varepsilon = 0.1$.

blue line), consistently with Theorem 5, while convergence to the stable matching requires an average number of rounds that grows exponentially in n (top, red line), as predicted by Theorem 6.

Fig. 3.6 shows the algorithm’s performance in reaching successively finer approximations of the best matching on random networks $G(n, m, \rho)$. Here, $G(n, m, \rho)$ refers to a random bipartite network with n leaders and m followers, in which each edge exists independently of the others with probability ρ (we fixed $\rho = 0.04$), and with constraint $c_\ell = \min\{m/n, |N_\ell|\}$ for each leader ℓ . For each choice of n and m that we considered, 100 random $G(n, m, \rho)$ were generated, and the algorithm was run 100 times on each. We observe that, consistently with Theorem 5, $\tau(\varepsilon)$ increases both when ε decreases (i.e., a finer approximation is desired) and when the number m of followers increases. The plot shows that a good solution is reached quickly, while most of the time is spent in the attempt of improving it to the best solution.




Figure 3.6. Algorithm’s performance to reach a $(1 - \epsilon)$ -approximate best matching on random bipartite networks $G(n, m, \rho)$, for $\rho = 0.04$ and different choices of n and m . For each choice of n and m , 100 random networks were generated, and each was simulated 100 times.

3.6 Experiments of human social computation

We conducted 36 experiments on a pool of 10 networks of 16 nodes each (each network was tested 3 or 4 times). Each of sixteen participants controls a node in the network via a computer interface which shows only its immediate neighbors (see Figure 3.2). During each experiment, a network is chosen and subjects are randomly assigned to nodes and informed whether they are playing the role of followers or leaders (in the latter case, the target number of followers is also specified). In order to elicit the common goal of reaching stability, each subject is paid a reward of \$1 if stability is reached within the maximum time of 5 minutes. Subjects can only interact via the computer interface: leaders can send matching requests to followers and break them with clicks of the mouse (for each leader, the number of concurrent outgoing requests plus matched followers can be at most equal to its group target size); followers can accept or reject leaders’ requests and break their own existing matched pairs, with clicks of the mouse.




Figure 3.7. Algorithm performance versus human subjects performance. The x -axis shows the NetworkID for the 10 different networks tested in the experiments, sorted by increasing average solving time in the experiments. The red line shows the average solving time (in seconds) for each network in the experiments (each tested three of four times); the blue triangles show the average number of rounds needed by the algorithm to solve the same networks (over 1000 simulations). The correlation between experiments' average time and algorithm's average number of rounds is 0.64 (p-value=0.04).

The networks range from simple random topologies to topologies similar to the network G_n defined above (and are not shown due to space constraints). After the experiments, each network was assigned a networkID such that higher IDs correspond to higher average solving time (if an experiment is not solved within the 5 minutes maximum time, a time of 5 minutes is considered). Figure 3.7 compares the performance of the human subjects (average number of seconds for each network, sorted by increasing solving time) and of the algorithm (average number of rounds over 10000 simulations on each network). Networks that required more time to be solved during the experiments also required more rounds of the algorithm (correlation 0.64 between number of seconds in the experiments and average number of rounds for the algorithm, p-value=0.04). Moreover, the networks with NetworkID from 8 to 10 (the most difficult to solve for the human subjects) are the topologies similar to the network G_n and were solved 6 times out of 11 (all other experiments were solved). Our results do not seem to be determined by




Figure 3.8. Human subjects performance: approximation versus stability. The x -axis shows the experimentID (sorted by increasing solving time). The red line shows the time spend by the human subjects to solve each experiment (300 seconds if the experiment was not solved), the green line shows the corresponding time spent to reach an approximate matching in which only a single leader needs an additional follower.

participants learning or getting tired (i.e., improving or worsening over time).⁴

Figure 3.8 compares the time needed by the human subjects to reach a stable matching versus an approximate solution in which only a single leader needs an additional follower (that is, with deficit equal to 1, according to the definition given above). Experiments are sorted by increasing solving time and we observe that a good solution is reached quickly while most of the time is spent improving it to the optimum, in agreement with the algorithm predictions. The time to reach an approximate solution with deficit equal to 1 and a stable solution are correlated (0.55, p -value 0.0004), and on average reaching the approximate solution requires about 7% of the total time (least squares regression, 0.065 p -value 0.0005 without controlling for NetworkID, 0.073 p -value 0.005 controlling for NetworkID). At the end of the experimental session, subjects were asked to complete an exit survey about their strategies. A wide range of strategies was reported. As for the leaders, participants reported to favor unmatched followers

⁴Difference between the solution time of subsequent experiments not significant (least squares regression, average additional 1.44sec. for each subsequent experiment, p -value 0.35), also controlling for NetworkID (0.51sec. less for each subsequent experiment on the same network, p -value 0.70).

(7 surveys, notice how this criterion agrees with our algorithm), blink (that is, quickly sending and canceling requests) in order to capture the attention of a follower (3 surveys), only request followers who did not break the matching earlier (2 surveys), try to match with new followers if the game is not solved for a while (3 surveys), and many other criteria. As for the followers, participants reported to always accept new requests (5 surveys), match with the leaders who are more persistent (4 surveys), match with leaders that are blinking (2 surveys), and so on. Clearly, trying to take all reported strategies into account (that might not correspond to the real strategies employed) would result in a complex and mathematically intractable model, and prevent us from deriving the clean trade-off between time and quality of the solution as stated in Theorem 5.

In practice, preferring unmatched followers appears to be a natural strategy, pursuing local stability is an inherent characteristic of human behavior – although subjects might not admit it explicitly – while randomization captures the diversity of the actions of the population as evidenced in the exit pools.

3.7 Discussion

The algorithmic model we proposed presents a set of desirable features aimed at modeling distributed social computation, and is simple enough to be prone to rigorous mathematical analysis. Despite its simplicity, it is able to predict human performance and fits the experimental data, showing that the global dynamics of complex agents with possibly diverse strategies can be well described by simple synthetic agents with uniform strategies. We advocate the usage of similarly simple algorithmic models to capture the essence of social interaction and to investigate a wider variety of social computation tasks.

In order to evaluate the proposed algorithmic model as a possible description of human behavior, we created an artificial environment in which human subjects solve

the group membership task on virtual networks. In these experiments, participants have the possibility to send, accept or decline matching requests as long as a solution is not reached and they are given a monetary reward upon successful and timely completion of each task. Two important features of the experiments are absent from the algorithmic model – the monetary reward and the time threshold. The monetary reward provides an incentive to solve each experiment and, as such, it appears necessary in the experimental design. On the other hand, in the algorithmic model, agents act until a stable solution is reached. The five-minute time threshold on each experiment guarantees that the entire experimental session has a constrained duration even in the presence of hard-to-solve networks. For such networks, we can compare the time to reach approximate solutions and a lower bound for the time to reach a stable solution. Despite these differences and the fact that participants can follow arbitrary strategies, the proposed algorithmic model is able to qualitatively capture the dynamics of the human subjects.

For ease of presentation, we assumed that agents are synchronized. We can consider an asynchronous setting where each of the m leaders has a clock that activates at random times. When a leader’s clock rings, the leader acts according to the algorithm in Table 1. Followers are activated by incoming requests from leaders and act according to the algorithm in Table 2. Our results would continue to hold substantially unchanged. For example, if we consider independent Poisson processes with inter-point times that are exponentially distributed with parameter $\lambda = 1$, a total of m events occur on average in a unit of time. This is comparable to the synchronous scenario in which all m leaders have the possibility to act during each time interval. The argument in the proof Theorem 5 would follow similarly by considering time intervals of fixed duration δ and replacing p by $p(1 - e^{-\delta})$, where the term between parentheses is a lower bound for the probability that a given clock rings within δ . The upper bound for the time to reach a $(1 - \epsilon)$ -approximate stable matching would present a multiplicative factor that depends on δ .

The best choice of δ depends on ε and for constant ε the multiplicative factor is constant in m .

In practical scenarios, the constraints c_ℓ can be lower bounds rather than exact targets. In this case, we assume that leader ℓ sends matching requests as long as her team's size is smaller than c_ℓ or there are unmatched followers in her neighborhood, and our main result in Theorem 5 continues to hold. In addition, if $\sum_\ell c_\ell < n$, where n is the number of followers, some followers will be unmatched in any configuration, and particularly in the first stable matching that is reached (if a stable matching exists), unless the constraints c_ℓ are lower bounds.

In the present paper, a fixed network topology is assumed. Multi-agent systems are often characterized by time-changing topologies, which result from either agents mobility or unreliable communication and whose effect on the system dynamics depends on the particular scenario. For example, Sarwate and Dimakis [190] showed that evolving topologies might help the diffusion of information in the context of averaging, Xiao and Wang [218] showed that they can harm convergence to consensus unless certain connectivity conditions hold. If the network topology is allowed to vary arbitrarily, our main result (Theorem 1) gives a performance guarantee in terms of the time *since* the most recent change in the network topology. We observe that a change in network topology that affects only few edges might in general result in a significant change in the optimal solution, and we leave the rigorous analysis of time-varying topologies to future investigation.

Finally, in the present work, we defined the quality of a solution in terms of the additional number of followers needed by all leaders in order to satisfy all size constraints (and we called this quantity the *deficit* of a matching). We proved that our algorithm constitutes a Polynomial Time Approximation Scheme (PTAS) for the minimization of this quantity, that is, *any* constant approximation of the optimal solution is reached in

polynomial time in *any* instance of the problem. Different quantities might be better suited to express the quality of an approximate solution in different applications. However, provable performance guarantees might in general be derived only for a subclass of the problem instances, and the analysis depends on the particular definition of approximation that is considered.

3.8 Proof of Theorem 5

Fix $0 < \varepsilon < 1$. Observe that $d(M(t))$ is non-increasing in t , as leaders do not voluntarily disengage from the followers in their groups (and therefore the deficit of a leader increases of one unit only if the deficit of another leader decreases by one unit). Moreover, since $c_\ell \geq 1$ for every leader ℓ , and G admits a stable matching, $d(M(t)) \leq m$ for every t .

For every $0 < x \leq 1$, let

$$\tau(x) = \min \left\{ t \geq 0 : d(M(t)) < xm \right\}$$

be the first round at whose beginning the deficit is strictly smaller than xm . We want to find an upper bound for $\tau(\varepsilon)$.

Consider any round $t \geq 0$. Since $d(M(t)) \leq m$, there exists $0 < \varepsilon' \leq 1$ such that $d(M(t)) = \varepsilon' m$ (we assume $\varepsilon' > 0$, as the case of $\varepsilon' = 0$ is trivial). The following lemma bounds the number of rounds $\tau(\varepsilon') - t$ needed for a one-unit reduction of the deficit. Let $\Delta = \max_{\ell \in L} |N_\ell|$.

Lemma 7 *Let $d(M(t)) = \varepsilon' m$ for some $0 < \varepsilon' \leq 1$. Then*

$$\Pr \left(\tau(\varepsilon') - t \leq \lfloor 1/\varepsilon' \rfloor \right) \geq \left(\frac{pq}{\Delta} \right)^{\lfloor 1/\varepsilon' \rfloor}.$$

Proof: Let $h(t) \geq 1$ be the odd length of each shortest deficit-decreasing path relative to $M(t)$. By Corollary 1, $h(t) \leq 2\lfloor 1/\varepsilon' \rfloor - 1$. We distinguish the cases of $h(t) = 1$ and $h(t) \geq 3$. First consider $h(t) = 1$. With probability at least pq/Δ the deficit decreases by at least one unit during the next round of the algorithm. To see this, consider a deficit-decreasing path ℓ, f . With probability at least p/Δ , ℓ tries to match with f and, conditional on this, f considers ℓ 's proposal with probability q , resulting in the lower bound pq/Δ .

Now consider $h(t) \geq 3$, and let P be a shortest deficit-decreasing path of length $h(t)$ ending at an unmatched follower f . The length of P decreases by one in the next round with probability at least pq/Δ (f is unmatched after round t as P is a shortest deficit decreasing path and $h(t) > 1$).

By independence of successive rounds of the algorithm and the bound $h(t) \leq 2\lfloor 1/\varepsilon' \rfloor - 1$, with probability at least $(pq/\Delta)^{\lfloor 1/\varepsilon' \rfloor}$, a sequence of $\lfloor 1/\varepsilon' \rfloor - 1$ rounds reduces the length of P to 1 and then in one additional round P is ‘‘solved’’ and the deficit decreases by one unit. \square

Consider consecutive phases of $\lfloor 1/\varepsilon \rfloor$ rounds each. For phases $i = 0, 1, 2, \dots$, let X_i be *iid* Bernoulli random variables with $\Pr(X_i = 1) = (pq/\Delta)^{\lfloor 1/\varepsilon \rfloor}$. By Lemma 7, after T phases (i.e., at the beginning of round $t^* = T\lfloor 1/\varepsilon \rfloor$), the deficit of the matching is upper bounded by

$$d(M(t^*)) < \max\{\varepsilon m, m + 1 - \sum_{i=1}^T X_i\},$$

as the matching at the beginning of round 0 has deficit $d(M(0)) \leq m$. By independence of the phases, a Chernoff bound implies that for any $0 < \delta \leq 1$

$$\Pr\left(\sum_{i=1}^T X_i < (1 - \delta)T(pq/\Delta)^{\lfloor 1/\varepsilon \rfloor}\right) < e^{-T(pq/\Delta)^{\lfloor 1/\varepsilon \rfloor} \delta^2/2}.$$

Setting $\delta = \varepsilon$ and $T = cm(\Delta/pq)^{\lfloor 1/\varepsilon \rfloor}$ (where c is a constant to be specified later),




Figure 3.9. The leader ℓ in the proof of Lemma 3.9, matched edges are highlighted.

the deficit of the matching at the beginning of round $t^* = \lfloor 1/\varepsilon \rfloor cm(\Delta/pq)^{\lfloor 1/\varepsilon \rfloor}$ is

$$d(M(t^*)) < \max \{ \varepsilon m, m + 1 - (1 - \varepsilon)cm \}$$

with probability at least $1 - e^{-cm\varepsilon^2/2}$. To conclude the proof of the theorem we need that $\varepsilon m \geq m + 1 - (1 - \varepsilon)cm$, which is true for any $c \geq 1 + \frac{1}{m(1-\varepsilon)}$.

3.9 Proof of Lemma 6

Given the matching M and the stable matching N , for brevity we write deficit-decreasing path instead of deficit-decreasing path in $M \oplus N$ relative to M .

The proof is divided in two parts. First, we show that for each leader ℓ with deficit $d_\ell(M) > 0$ there are at least $d_\ell(M)$ follower-disjoint deficit-decreasing paths starting at ℓ . Then, we argue that $d(M)$ follower-disjoint deficit-decreasing paths can be chosen, $d_\ell(M)$ of which start at each leader ℓ with deficit $d_\ell(M) > 0$, and the claim of the lemma follows.

Consider a leader ℓ with $d_\ell(M) > 0$. Assume by contradiction that there are strictly less than $d_\ell(M)$ follower-disjoint deficit-decreasing paths starting at ℓ (see Figure 3.9). Since ℓ has a team size constraint $c_\ell > 0$, there are exactly $c_\ell - d_\ell(M)$ followers that are matched to ℓ . Observe that no follower matched to ℓ can be the first follower of a deficit-decreasing path starting at ℓ (because the first edge must be in $N \setminus M$).

Since G admits a stable matching, the neighborhood N_ℓ of ℓ has size $|N_\ell| \geq c_\ell$.

Therefore, there are $k \geq d_\ell(M)$ followers in N_ℓ that are not matched to ℓ . Assume that $h < d_\ell(M)$ of the followers in N_ℓ are the first followers of h follower-disjoint deficit-decreasing paths starting at ℓ (P_1, \dots, P_h in Figure 3.9). Denote the remaining $k - h > 0$ followers by f_1, \dots, f_{k-h} , and assume by contradiction that none among them is the first follower of a deficit-decreasing path starting at ℓ (i.e., there are strictly less than $d_\ell(M)$ follower-disjoint deficit-decreasing paths starting at ℓ).

Observe that, in order to become stable, ℓ needs to match with at least one additional follower among $\{f_1, \dots, f_{k-h}\}$. We show that, under the assumption above, a one-unit reduction in the deficit of ℓ would eventually result in a one-unit increase of the deficit of another leader, implying that G does not admit a stable matching, generating a contradiction.

Consider any follower $f' \in \{f_1, \dots, f_{k-h}\}$, and observe that f' is matched in M since otherwise $\ell f'$ would be a deficit-decreasing path starting at ℓ . Let ℓ' be the leader such that $(\ell', f') \in M$, and observe that if ℓ' is matched to all followers in $N_{\ell'}$ then ℓ cannot match to f' without causing a one-unit increase of the deficit of ℓ' . Therefore assume that in $N_{\ell'}$ there is a follower f'' such that $(\ell'', f'') \in M$ for some leader $\ell'' \neq \ell'$ (f'' is matched in M since otherwise ℓ, f', ℓ', f'' is a deficit-decreasing path). In the following two cases ℓ cannot match to f' without eventually increasing the deficit of another leader: (i) $\ell'' = \ell$, and $\ell, f', \ell', f'', \ell$ is a cycle; (ii) $\ell'' \neq \ell$ and ℓ'' is matched to all followers in $N_{\ell''}$ other than f' .

Therefore assume that in $N_{\ell''}$ there is a follower f''' such that $(\ell''', f''') \in M$ for some leader $\ell''' \neq \ell''$ (f''' is matched in M). Again, ℓ cannot match to f' without eventually increasing the deficit of another leader if either $\ell''' = \ell$ or $\ell''' = \ell'$ (each similar to the case (i) above), or if ℓ''' is matched to all followers in $N_{\ell''}$ other than f', f'' (similar to the case (ii) above). By iteration, it follows that ℓ cannot match to any follower $f' \in \{f_1, \dots, f_{k-h}\}$ without eventually increasing the deficit of another leader,

contradicting with the existence of the stable matching N . Hence, there are at least $d_\ell(M)$ follower-disjoint deficit-decreasing paths starting at ℓ .

To complete the proof of the lemma, we show that we can choose $d(M)$ follower-disjoint deficit-decreasing paths, $d_\ell(M)$ of which start at each leader ℓ with $d_\ell(M) > 0$. We proceed by contradiction, and make the following assumption. For any set \mathcal{P} of $d(M)$ deficit-decreasing paths, $d_\ell(M)$ of which start at each leader ℓ with $d_\ell(M) > 0$ (denote by \mathcal{P}_ℓ the elements of \mathcal{P} starting at ℓ), there are leaders ℓ, ℓ' and paths $P \in \mathcal{P}_\ell, P' \in \mathcal{P}_{\ell'}$ that are not follower-disjoint. In order to reach the stable matching N starting from M , a set of $d(M)$ deficit-decreasing paths must be solved. However, if P is solved (by “flipping” matched edges into unmatched edges, and vice versa) then P' is not solved, and if P' is solved then P is not solved. It follows that N cannot be reached from M by solving the $d(M)$ deficit-decreasing paths in \mathcal{P} .

By the assumption above, the last argument holds for any choice of \mathcal{P} , and this generates a contradiction on the reachability of N starting from M (observe that N can be reached from M in finite time, e.g. by a cat-and-mouse argument on the space of all the matchings of G). The lemma is proven.

3.10 Proof of Theorem 6

Let \mathcal{M}_n be the set of all the matchings of G_n such that $d(M) = 1$. We proceed as follows. First, we show that each $M \in \mathcal{M}_n$ is uniquely identified by the set of the leaders that are not matched with “horizontal” edges (that is, leaders ℓ_i such that $(\ell_i, f_i) \notin M$). Second, we define trees T_m^* , $m \geq 1$ such that a random walk on T_m^* starting at any node different than the root hits the root after a number of steps that is exponentially large in m with high probability. Third, for each matching $M \in \mathcal{M}_n$ we define a quantity $h(M)$ that we call the *height* of M and we argue that, when initialized at M , the algorithm’s dynamics is equivalent to a random walk on the tree $T_{h(M)}^*$ and reaching the stable matching of G_n

corresponds to reaching the root of $T_{h(M)}^*$ (and therefore requires a number of rounds that is exponentially large in $h(M)$ with high probability). Finally, by a counting argument, we show that for any constant $0 < \gamma < 1$ a $1 - O(n2^{-(1-\gamma)n})$ fraction of all the matchings in \mathcal{M}_n have height at least γn , completing the proof of the theorem.

3.10.1 Properties of the matchings in \mathcal{M}_n

Matchings in \mathcal{M}_n enjoy the following structural properties.

Lemma 8 *Let $M \in \mathcal{M}_n$. The following properties hold.*

- (1) *There are a single poor leader $\ell_{i^*(M)}$ and a single unmatched follower $\ell_{j^*(M)}$ in M .*
- (2) $1 \leq i^*(M) \leq j^*(M) \leq n$.
- (3) $(\ell_k, f_k) \in M$ for all $k < i^*(M)$ and all $k > j^*(M)$.
- (4) *Let $\mathcal{J}(M) = \{j_0, j_1, \dots, j_K\}$ be the sorted set of indexes j such that $(\ell_j, f_j) \notin M$.*

Then

- (a) $j_1 = i^*(M)$ and $j_K = j^*(M)$.
- (b) $(\ell_{j_{k+1}}, f_{j_k}) \in M$ for all $k \in \{0, \dots, K-1\}$.

Proof: Property (1). Since $d(M) = \sum_{\ell \in L} d_\ell(M) = 1$, there is a single poor leader $\ell_{i^*(M)}$ in M . Since $c_\ell = 1$ for all $\ell \in L$, each leader $\ell \neq \ell_{i^*(M)}$ is matched to a single follower. It follows that there is a unique unmatched follower $f_{j^*(M)}$.

Property (2). Suppose by contradiction that $i^*(M) > j^*(M)$. Since $N_{\ell_{j^*(M)}} = \{f_1, \dots, f_{j^*(M)}\}$ and $f_{j^*(M)}$ is unmatched, leader $\ell_{j^*(M)}$ is matched to one of the followers in $\{f_1, \dots, f_{j^*(M)-1}\}$. Hence, the $j^*(M) - 1$ leaders $\ell_1, \dots, \ell_{j^*(M)-1}$ are matched to at most $j^*(M) - 2$ out of the $j^*(M) - 1$ followers $f_1, \dots, f_{j^*(M)-1}$, and one of them is necessarily poor, contradicting Property (1).




Figure 3.10. An example of a matching M of G_6 with $d(M) = 1$. M is uniquely determined by the set $\mathcal{S}(M) = \{2, 4, 6\}$, encoding the following: l_2 is not matched, l_4 is matched with f_2 , l_6 is matched with f_4 , f_6 is not matched. $P(M) = l_2, f_2, l_4, f_4, l_6, f_6$ is the unique deficit-decreasing path.

Property (3). We proceed by induction. If $i^*(M) > 1$, then $(l_1, f_1) \in M$ since $N_{l_1} = \{f_1\}$ and l_1 is matched with a follower. Assume that if $i^*(M) > j$ then $(l_k, f_k) \in M$ for all $k \leq j$. If $i^*(M) > j + 1$, then, by the inductive assumption, l_{j+1} can only be matched to f_{j+1} since $N_{l_{j+1}} = \{f_1, \dots, f_{j+1}\}$. This shows that $(l_k, f_k) \in M$ for all $k < i^*(M)$. $k > j^*(M)$ is shown similarly.

Property (4). If $K = 0$ then $M = \{(l_i, f_i) : i \neq i^*(M)\}$, $j^*(M) = i^*(M)$, and properties (4a) and (4b) trivially hold. Now consider $K \geq 1$. Let $\mathcal{S}(M) = \{j_0, j_1, \dots, j_K\}$ be the sorted set of indexes j such that $(l_j, f_j) \notin M$. By property (3), we have that $j_0 = i^*(M)$ and $j_K = j^*(M)$, therefore property (4a) follows. Hence, $(l_{j_2}, f_{j_1}) \in M$ since $(l_k, f_k) \in M$ for all $k \in \{j_1 + 1, \dots, j_2 - 1\}$ by definition of $\mathcal{S}(M)$, and $N_{l_{j_2}} = \{f_1, \dots, f_{j_2}\}$. Property (4b) follows by induction. \square

Lemma 8 states that non-horizontal matching edges do not intersect. In particular, given a matching $M \in \mathcal{M}_n$, the set $\mathcal{S}(M)$ represents the set of (the sorted indexes of) the leaders that are not matched with horizontal edges (see Figure 3.10 for an example), $l_{i^*(M)}$ for $i^*(M) = \min \mathcal{S}(M)$ is the unique unmatched leader, and $f_{j^*(M)}$ for $j^*(M) = \max \mathcal{S}(M)$ is the unique unmatched follower. Recall that $M_n^* = \{(l_k, f_k) : 1 \leq k \leq n\}$ is the unique stable matching of G_n , and let $\mathcal{S}(M_n^*) = \emptyset$. Lemma 8 implies that every matching $M \in \mathcal{M}_n \cup \{M_n^*\}$ is uniquely identified by the set $\mathcal{S}(M)$.




Figure 3.11. The three T_m^* for $m = 5$.

Lemma 9 *The mapping $\mathcal{S}(\cdot)$ from $\mathcal{M}_n \cup \{M_n^*\}$ to $\mathcal{S} = \{A : A \subseteq \{1, \dots, n\}\}$ defined by $M \mapsto \mathcal{S}(M)$ is a bijection.*

Proof: $\mathcal{S}(\cdot)$ is injective as if $M, M' \in \mathcal{M}_n$ and $M \neq M'$ then $\mathcal{S}(M) \neq \mathcal{S}(M')$. $\mathcal{S}(\cdot)$ is surjective, as for $K \leq n - 1$ and $A = \{i_0, \dots, i_K\} \in \mathcal{S}$ such that $1 \leq i_0 < \dots < i_K \leq n$, the matching $M \in \mathcal{M}_n$ such that $\mathcal{S}(M) = A$ is $M = \{(\ell_{i_{k+1}}, f_{i_k}) : 0 \leq k \leq K - 1\} \cup \{(\ell_k, f_k) : k \notin A\}$. \square

3.10.2 The tree T_m^*

Let T_1 be a labeled rooted tree with a singleton node with label 1. Inductively, for $i \leq 2$, let T_i be the labeled rooted tree whose root is labeled with i and its $i - 1$ children are the roots of copies of T_1, \dots, T_{i-1} . Define T_m^* to be the tree with a root with label $m + 1$ whose only child is the root of a copy of T_m (see Figure 3.11). Let r^* be the root of T_m^* . We show that the hitting time of r^* for a random walk on T_m^* starting at any node $u \neq r^*$ is exponential in m with high probability. For a node $u \neq r^*$, we call the edge that connects u to its parent u 's *exit edge*. For any subtree $T_i \subset T_m^*$, let Z_i be the random variable denoting the number of steps that it takes for a walk starting at the root of T_i to exit T_i (that is, to hit the parent of the root of T_i). The following bound holds for Z_i .

Lemma 10 *There exist $\alpha, \gamma > 0$ such that, for all $i \geq 2$,*

$$\Pr[Z_i \geq \gamma \cdot 2^{i/(\alpha \log^2 i)}] \geq 1 - \frac{1}{\log i}.$$

Proof: The proof is the same as Lemma 4 of Chapter 2 and is reported here for completeness. We proceed by induction on i . Let $g(i) = \alpha \log^2 i$ and $f(i) = \gamma \cdot 2^{i/g(i)}$ for some $\alpha, \gamma > 0$. For any $\alpha > 0$ and $i \geq 2$, we can choose $\gamma > 0$ such that $f(i) \leq 1$; therefore, as $Z_i \geq 1$ with probability 1, the claim holds trivially for any $i \leq i^*$, where i^* is a suitably large constant. Now consider any $i \geq i^*$ and suppose the claim holds up to $i - 1$. Every time the walk is on the root of T_i , it exits T_i with probability $1/i$ (since the root of T_i one parent and $i - 1$ children). Therefore, letting E_t be the event that the first t times the walk is on the root of T_i it does *not* exit T_i , we have $\Pr[E_t] \geq 1 - t/i$. Let $t = i/(2 \log i)$, and let D_j , $1 \leq j \leq t$, be the event that, when it is on the root of T_i for the j -th time, the walk moves to the root of one of the subtrees $T_{i-g(i)}, \dots, T_{i-1}$ and takes at least $f(i - g(i))$ steps to exit that subtree. For $1 \leq j \leq t$,

$$\begin{aligned} \Pr[D_j \mid E_t] &\geq g(i)/i \cdot \Pr[Z_{i-g(i)} \geq f(i - g(i))] \\ &\geq g(i)/i \cdot (1 - [\log(i - g(i))]^{-1}), \end{aligned}$$

by the induction hypothesis on $Z_{i-g(i)}$. Let χ_j be the indicator function of D_j for $1 \leq j \leq t$.

The probability that at least two of the events D_j happen, given E_t , is lower bounded by:

$$\begin{aligned}
\Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] &\geq \Pr \left[\sum_{j=1}^{t/2} \chi_j \geq 1, \sum_{j=t/2+1}^t \chi_j \geq 1 \mid E_t \right] \\
&= \Pr \left[\sum_{j=1}^{t/2} \chi_j \geq 1 \mid E_t \right]^2 \geq \left(1 - \prod_{i=1}^{t/2} (1 - \Pr[D_j | E_t]) \right)^2 \\
&\geq \left[1 - (1 - g(i)/i (1 - [\log(i - g(i))]^{-1}))^{t/2} \right]^2 \\
&\geq \left[1 - \exp \left[\frac{-\alpha \log i}{4} \left(1 - \frac{1}{\log(i - g(i))} \right) \right] \right]^2 \geq \left(1 - \frac{1}{i^{\alpha/8}} \right)^2
\end{aligned}$$

where we applied union bound in the third line, and the last step holds for i sufficiently large so that $\log(i - g(i)) \geq 2$. This implies that

$$\Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] \geq \left(1 - \frac{1}{i^{\alpha/8}} \right)^2 \geq 1 - \frac{2}{i^{\alpha/8}},$$

and we conclude that

$$\begin{aligned}
\Pr[Z_i \geq 2 \cdot f(i - g(i))] &\geq \Pr \left[\sum_{j=1}^t \chi_j \geq 2 \right] \\
&\geq \Pr \left[\sum_{j=1}^t \chi_j \geq 2 \mid E_t \right] \Pr[E_t] \geq \left(1 - \frac{2}{i^{\alpha/8}} \right) \left(1 - \frac{t}{i} \right),
\end{aligned}$$

which is greater than $1 - 1/\log i$ for α sufficiently large. The claim follows since $2 \cdot f(i - g(i)) \geq f(i)$. \square

Corollary 3 *The hitting time of r^* of a random walk starting at any node $u \neq r^*$ is $2^{\Omega(n/\log^2 n)}$ with high probability.*

3.10.3 The dynamics of the algorithm starting from $M \in \mathcal{M}_n$

For ease of presentation, we set the parameters of the algorithms to $p = q = 1$ (our result holds in general).

Definition 10 (The height of a matching) Let $M \in \mathcal{M}_n$, $\mathcal{I}(M) = \{i_0, \dots, i_K\}$. The height of M is $h(M) = 0$ if $K = 0$, and $h(M) = i_{K-1} \in \{1 \dots, n-1\}$ if $K \geq 1$.

For $M \in \mathcal{M}_n$, $h(M) > 0$, $\mathcal{I}(M) = \{i_0, \dots, h(M), i_K\}$. For $t \geq 0$, let $M(t)$ be the matching at the beginning of round t . For ease of notation let $\mathcal{I}(t) = \mathcal{I}(M(t))$. For $M \in \mathcal{M}_n$ let

$$\tau^*(M) = \min \{t : M(t) = M_n^* | M(0) = M\}$$

be the number of steps that the algorithm needs to reach the stable matching starting from M .

Note that, with $p = q = 1$, $\tau^*(M) = 1$ for every $M \in \mathcal{M}_n$ such that $h(M) = 0$ (that is, $|\mathcal{I}(M)| = 1$), since according to the algorithm leaders prefer unmatched followers. We are interested in relating $\tau^*(m)$ and $h(M)$ for every matching $M \in \mathcal{M}_n$ such that $h(M) > 0$ (that is, $|\mathcal{I}(M)| > 1$).

We study how the matching evolves over time through the Markov process $\{\mathcal{I}(t) : 0 \leq t \leq \tau^*(M)\}$. Since $\mathcal{I}(M_n^*) = \emptyset$, $\tau^*(M) = \min\{t : \mathcal{I}(t) = \emptyset\}$. The state space of the Markov process is given by the set \mathcal{S} defined in Lemma 9. The transition probabilities are as follows.

Lemma 11 Conditional on $\mathcal{I}(t) = I \in \mathcal{S}$, $|I| > 1$, the transition probabilities at time t are given by

$$\Pr(\mathcal{I}(t+1) = I' | \mathcal{I}(t) = I) = 1/\min I.$$

if $I' \in \{I \cup \{k\} : k < \min I\} \cup \{I \setminus \{\min I\}\}$ and 0 otherwise. Moreover $\Pr(\mathcal{S}(t+1) = \emptyset | \mathcal{S}(t) = \emptyset) = 1$, and $\Pr(\mathcal{S}(t+1) = \emptyset | \mathcal{S}(t) = I) = 1$ for every I such that $|I| = 1$.

Proof: The case of $\mathcal{S}(t) = \emptyset$ corresponds to the stable matching M_n^* , which is an absorbing state for the Markov process. In the case of $|\mathcal{S}(t)| = 1$, we have that $h(M) = 0$, and $p = q = 1$ implies that that $\mathcal{S}(t+1) = \emptyset$.

Consider now $|I| > 1$. Conditional on $\mathcal{S}(t) = I$, the poor leader is $\ell_{\min I}$ and has degree $\min I$ and neighborhood $N_{\min I} = \{f_1, \dots, f_{\min I}\}$, and chooses one of the followers in $N_{\min I}$ uniformly at random. If $\ell_{\min I}$ chooses follower f_k for some $k < \min I$ then the leader ℓ_k becomes poor, since by property (3) of Lemma 8 ℓ_k was matched to f_k in $M(t)$, and we have that $\mathcal{S}(t+1) = I \cup \{k\}$. If instead $\ell_{\min I}$ chooses follower $f_{\min I}$ (matched to $\ell_{\min(I \setminus \min I)}$ in $M(t)$ by property (4) of Lemma 8), then $\mathcal{S}(t+1) = I \setminus \{\min I\}$. \square

For every $M \in \mathcal{M}_n$ such that $h(M) > 0$ and $\mathcal{S}(M) = \{i_0, \dots, i_K\}$, define the matching

$$\mathcal{L}(M) = \{(\ell_j, f_j) : j \neq i_K\},$$

and let

$$\tau(M) = \min\{t : M(t) = \mathcal{L}(M)\}.$$

Note that $h(\mathcal{L}(M)) = 0$ and $\tau^*(M) > \tau(M)$ (in particular, $\tau^*(M) = 1 + \tau(M)$ for $p = q = 1$).

For every matching M such that $|\mathcal{S}(M)| > 1$, let $\mathcal{R}(M)$ be the set of the matchings in \mathcal{M}_n that can be reached from M (after one or multiple steps). By the transition probabilities of Lemma 11,

$$\mathcal{R}(M) = \{M' \in \mathcal{M}_n : I(M') = A \cup \{h(M), i_K\}, A \in \mathcal{A}\} \cup \{\mathcal{L}(M)\},$$

where $\mathcal{A} = \{A \subseteq \{1, \dots, h(M) - 1\}\}$. Observe that every $M' \in \mathcal{R}(M) \setminus \{\mathcal{L}(M)\}$ has height $h(M') = h(M)$. The following lemma characterizes the one-to-one correspondence between matchings in $\mathcal{R}(M)$ and nodes of the tree $T_{h(M)}^*$.

Lemma 12 *Consider the mapping $\omega(\cdot)$ from $\mathcal{R}(M)$ to $T_{h(M)}^*$ defined as follows. Let $\omega(\mathcal{L}(M)) = r$, where r is the root of $T_{h(M)}^*$. For $M' \in \mathcal{R}(M) \setminus \{\mathcal{L}(M)\}$ and $\mathcal{I}(M') = I$, let $\omega(M')$ be the node of $T_{h(M)}^*$ with label $\min I$ and connected to the root with the path of nodes labeled by the sorted indexes in $I \setminus \{\min I\}$. Then $\omega(\cdot)$ is a bijection.*

The proof follows from the construction of $T_{h(M)}^*$ and $\mathcal{I}(\cdot)$.

Lemma 13 *The process $\{\mathcal{I}(t) : 0 \leq t \leq \tau(M) \mid M(0) = M\}$ is equivalent to a random walk on $T_{h(M)}^*$ starting at $\omega(M)$.*

Proof: It suffices to show that the transition probabilities between two matchings $M_1, M_2 \in \mathcal{R}(M)$ are nonzero if and only if the nodes $\omega(M_1)$ and $\omega(M_2)$ are adjacent in $T_{h(M)}^*$. To prove the “only if” direction, assume that $M_1, M_2 \in \mathcal{R}(M)$ are such that there is a nonzero transition probability from M_1 to M_2 (and therefore from M_2 to M_1). Let $\mathcal{I}(M_1) = I_1$ and $\mathcal{I}(M_2) = I_2$. According to the transition probabilities given above, there are two possible cases. In the first case, $I_2 = I_1 \cup \{k\}$ for some $k < \min I_1$, and $\omega(M_2)$ is a child of $\omega(M_1)$. In the second case $I_2 = I_1 \setminus \{\min I_1\}$ and $\omega(M_2)$ is the parent of $\omega(M_1)$. The other direction is similar. \square

To summarize, the number of steps to reach the stable matching of G_n starting from $M \in \mathcal{M}_n$ with $h(M) > 0$ is upper bounded by the time $\tau(M)$ to reach the matching $\mathcal{L}(M)$, and reaching $\mathcal{L}(M)$ is equivalent to reaching the root of $T_{h(M)}^*$ starting from the node $\omega(M)$. By Corollary 3, $\tau(M)$ is exponentially large in $h(M)$ with high probability.

3.10.4 The fraction of matchings $M \in \mathcal{M}_n$ such that $h(M) \geq \gamma n$

Let N be the number of matchings in \mathcal{M}_n . Fix a constant $0 < \gamma < 1$, let

$$\mathcal{M}_\gamma = \{M \in \mathcal{M}_n : h(M) < \gamma n\}$$

and let $N_\gamma = |\mathcal{M}_\gamma|$. For $j = 0, \dots, n-1$, let $N(j)$ be the number of matchings $M \in \mathcal{M}_n$ such that $h(M) = j$. It follows that

$$N = \sum_{j=0}^{n-1} N(j), \quad N_\gamma \leq \sum_{j=0}^{\lceil \gamma n \rceil - 1} N(j).$$

Lemma 14 $N(0) = n$ and $N(j) = (n-j)2^{j-1}$ for all $j = 1, \dots, n-1$.

Proof: $N(0) = n$ since there are n matchings M with $h(M) = 0$ ($\{(\ell_j, f_j) : j \neq k\}$ for $1 \leq k \leq n$). Fix $j \in \{1, \dots, n-1\}$. By Lemma 9, any $M \in \mathcal{M}_n$ with $h(M) = j$ is uniquely identified by $\mathcal{S}(M) = \{i_0, \dots, i_{K-1}, i_K\}$ for some $1 \leq K \leq n-1$ and $i_{K-1} = j$. Since $\mathcal{S}(\cdot)$ is a bijection, to determine $N(j)$ we need to count all subsets of $\{1, \dots, n\}$ of the form $\{i_0, \dots, j, i_K\}$, thus $N(j) = (n-j)2^{j-1}$. \square

For any constant $0 < \gamma < 1$, the fraction of $M \in \mathcal{M}_n$ such that $h(M) < \gamma n$ goes to zero exponentially fast in n

Lemma 15 Fix $0 < \gamma < 1$. Then, $N_\gamma/N = O(n2^{-(1-\gamma)n})$.

Proof: We first compute N .

$$N = \sum_{i=0}^{n-1} N(i) = n + \sum_{i=1}^{n-1} (n-i)2^{i-1} = n + n \sum_{i=0}^{n-2} 2^i - \sum_{i=1}^{n-1} i2^{i-1}.$$

The second sum can be shown (e.g. by induction) to be equal to $(n-1) + (n-2)(2^{n-1} - 1)$. Therefore, $N = \Omega(2^n)$, and for $k = \lceil \gamma n \rceil$, $N_\gamma = O(n2^{\lceil \gamma n \rceil})$, and $N_\gamma/N = O(n2^{-(1-\gamma)n})$.

□

3.11 Acknowledgments

Chapter 3, in full, has been submitted for publication of the material as it might appear in the IEEE Transaction on Network Control Systems. L. Coviello, and M. Franceschetti. The dissertation author was primary investigator and co-author of this paper.

I would like to thank Matthew Jones for helping us with the experiments. I am indebted to Stephen Judd and Michael Kearns for allowing us to modify their software for running our experiments.

Chapter 4

Non-invasive detection of emotional contagion in on-line social networks

4.1 Introduction

In the last decade, the challenge of understanding the spreading and synchrony of human behavior over social networks has attracted the attention of the research community at large. The problem originally arises in the context of the social sciences, but due to the expanding usage of online social networks, it has also attracted the interest of the engineering community with the aim of quantifying these effects using the massive amount of data that these networks generate. Studies have included the diffusion of news and “memes” [140]; cascades in communication platforms, networked games, microblogging services [83]; health-related phenomena such as obesity and smoking [51, 52]; emotional states like happiness and depression [74, 183]; purchase of online products [14, 158]; clicking online advertisements, and joining online recreational leagues and store purchases [82].

Studies based on observational data pose an inherent difficulty for causal inference because social contacts may have similar behavior as a result of at least two processes: homophily (the tendency of similar individuals to group together) or influence [150, 152]. Controlled experiments allow us to disentangle influence effects from homophily both in

the laboratory [75] and online [43, 44, 91, 34], but they are often limited in scale and lack external validity. Large scale experiments have been shown to be feasible in the context of political participation [34], product adoption [14, 158] and emotional influence [135], but are often impractical or require very close collaboration with private companies.

Moreover, the experimental change in the users' experience required by some of these studies recently came under scrutiny because of questions about the ethics involved. Some people criticized [84] a large-scale study [135] of emotional contagion on Facebook in which the researchers changed the content shown to some users in order to study their reaction. Similar criticisms were directed at the online dating website OkCupid for experimenting with their platform in order to understand how individuals react to each other [216]. These recent events call for the development of alternative, non-experimental methods to study human behavior at large scale,

Our work is an attempt to compensate for the shortcomings of existing experimental and observational approaches, using a method to detect and quantify influence via instrumental variable regression. We studied text-based expression in massive social networks, developed a model of emotional contagion of semantic expression, and validated it on the content posted by a large sample of Facebook users over a period of three years.

In this paper, we show how our model can also be applied to different and possibly heterogeneous data from other social networking platforms, and to contexts other than emotional expression. Our approach is fully non-experimental: it is based only on observational data and, as a result, it does not alter users' experience. It also guarantees respect for user privacy: for our study, individuals' information and posts were never visible to researchers and resided on secure servers where Facebook stores user data, and were analyzed only at an aggregate level. The study was reviewed for ethics and approved in advance by the Institutional Review Board at the University of California, San Diego.

Focusing on the mathematical model and on the engineering methodology employed, this paper reviews and complements our previous work. Our *individual-level* model assumes that a person’s usage of words in a semantic category is a linear function of temporal and individual baseline effects, exogenous variables like news, the stock market, or the weather, and endogenous variables – corresponding to the usage of given semantic categories in posts written by the person’s social contacts, referred to as “friends”. The *reciprocal causality* between the endogenous variables of the model makes it difficult to obtain consistent and unbiased estimates of social influence. We therefore proceed in two steps. First, we aggregate the model on a geographical basis by averaging over all people who are in the same city, obtaining a model based on the same coefficients as the individual-level model but with a much smaller number of observations. Second, we deal with the problem of reciprocal causality by estimating the model using *instrumental variable* regression, a method pioneered in economics [9]. This method relies on the availability of an exogenous variable – called an *instrument* – that affects the endogenous variables (friends’ posts) but does not directly induce a change in the subject’s posts, called the *dependent variable*. In general, valid instruments might be unavailable, or they might lack sufficient power to predict changes in the endogenous variable. In our work, we considered *rainfall* experienced by friends as the instrument, using data made available by the National Climatic Data Center¹, which proved to be a robust predictor of emotional expression. Upon finding a relationship between friend’s rainfall and their expression, we can assume the former affects the latter as the opposite direction is unlikely. Our method first computes the effect that friends’ rainfall (the instrument) has on friends’ posts (the endogenous variables). Then, it evaluates the corresponding effect of the rainfall-induced change in friends’ expression on the person’s posts (the dependent variable).

¹NCDC, <http://www.ncdc.noaa.gov>

In order to obtain consistent estimates, the instrument must satisfy the *exclusion restriction* [9]. This posits that, controlling for all other variables, the instrument (friends' rainfall) must not directly affect the dependent variable. An implication of this restriction is that the instrument must also be uncorrelated with the exogenous variable experienced by the subject (subject's rainfall), otherwise the model might only be estimating how a subject's rainfall affects her own expression. Therefore, to break any correlation between a subject's rainfall and friends' rainfall, we restricted our analysis to observations for which it did not rain in the subject's city. Once this is applied, the subject's rainfall is constant in the dataset and therefore it does not correlate with friends' rainfall. Moreover, breaking the correlation between user' and friends' rainfall solves the potential issue of the geographic similarity of the weather in close-by cities. As a result, we must also focus exclusively on social ties between individuals in different cities. Note that individuals in different cities likely do not interact face-to-face, but they can reach each other via multiple communication media, such as the telephone, email, and social networking websites. Therefore, any influence detected between them is unlikely to be caused by physical interaction and would suggest that remote communication plays an important role in spreading semantic expression.

Our method allows us to determine what semantic categories are susceptible to influence between social contacts by estimating how an individual's usage of a semantic category is affected by her friends' usage of the *same* category. We can then use the estimates for each semantic category to rank them from the most to the least likely to spread.

Moreover, our method allows us to determine the *relationship* between different semantic categories, by estimating how an individual's usage of one category is altered by her friends' usage of a *different* category. This will help us to understand whether the usage of a semantic category fosters or inhibits the usage of other categories. We

also show that expression of positive affect inhibits expression of negative affect and vice versa.

Finally, our model allows us to compute the *cumulative effect* a person has on her friends. Although the effect on any one social contact will be small, each person typically has many social contacts, so the total expected effect of a single act of expression may alter the expression of several other people. Here, we show how to use our model to quantify this *multiplier effect* on posts within the same semantic category and on posts in different categories.

4.1.1 Related work

Our work is related to a growing body of literature on influence and diffusion in networks, whose goal is to characterize how behaviors and information spread from person to person. Online social networks are becoming increasingly popular as research environments and sources of data for these investigations. For example, the content posted by people online has been used to identify which people or topics are influential in social networking websites [20] and in the blogosphere [4]. It has also been used to study which network attributes and sharing behaviors make people influential [45], which topics (e.g., represented by hashtags) diffuse in a more persistent way [182], and even to study the structure of diffusion cascades on different communication platforms [83]. Large-scale experimental studies have isolated the role of the network in the diffusion of information [21], emotional expression [135] and behaviors [14, 34]. However, homophily has been shown to play a similarly important role, and scholars have devoted their attention to distinguishing between the two phenomena and to comparing the size of their effects [152, 44, 11, 12].

Our work is related to the econometric literature on instrumental variables. Instrumental variables have been proposed as a tool to infer causal effects from observa-

tional data [9]. This approach has been applied to a variety of contexts, such as labor economics [10], the study of the causal effect of education on earning [40], program evaluation [105], the characterization of neighborhood effects [131], and the impact of microfinance [157]. However, valid instruments can be difficult to find [35], and scholars have warned against the risks of using “weak” instruments that do not predict variation in the endogenous variable with sufficient precision [198].

A large body of research studies text meaning by analyzing patterns of words or grammar [201, 58, 68]. However, the performance of most traditional classification methods relies on sufficient text length, as in the case of bag-of-words or kernel-based methods [112, 144]. The analysis of short text from microblogging services (such as Twitter or Facebook) requires new approaches [197, 176, 111], which in some cases leverage metadata (e.g., user’s information) or the content of related posts.

Although we mainly focus on the engineering aspects of the detection and measurement of peer influence in semantic expression, our work is also related to sociolinguistics. The full understanding of language in a society requires us to consider the social network in which the language is embedded, intended as the set of relationships and interactions between its individuals [213]. Scholars have argued that speech patterns might depend on the looseness and tightness of the social network [64]. Our model formulation allows us to take tie-strength between individuals into account. Different approaches have been proposed to quantify tie-strength in online social networks [80, 217], and future research should investigate whether strong ties play a major role in the spread of semantic expression.

4.2 Variables of the model

We consider a set T of distinct days. For each day $t \in T$, let $S(t)$ be the population on day t , and let $n(t) = |S(t)|$ be their number. To apply our method, we assume that

individuals can be geolocated at the level of cities. For each city g let $S_g(t)$ be the set of individuals in city g on day t and let $n_g(t) = |S_g(t)|$. In general, one might consider different time and geographic resolution. We assume resolution at the level of days and cities.

4.2.1 Quantifying the semantic of text-based expression

Several methods can be used to *quantify* semantic expression of the content posted by individuals (see discussion in Section 4.9). We referred to the semantic categories defined by the Linguistic Inquiry and Word Count (LIWC) 2007 [171], a word classification tool widely used in the social sciences and in psychology research [172, 207, 153, 173, 89]. The LIWC contains several classes of processes, each of which contains one or more semantic categories, pertaining to affective processes, perceptual processes, biological processes, social processes, and personal concerns. We consider the categories for positive and negative affective processes. In general, a larger set C of semantic categories can be considered by our method.

For user i and day t , let $U_i(t)$ be the set of status updates posted by i on day t , and let $u_i(t) = |U_i(t)|$ be its cardinality. Let $u_i^{(p)}(t)$ be the number of status updates in $U_i(t)$ that contain at least one word from the “positive emotion” category defined by LIWC 2007 [171]. Similarly, let $u_i^{(n)}(t)$ be the number of status updates in $U_i(t)$ that contain at least one word from the “negative emotion” category.

Note that a single status update might contain both a negative word and a positive word, therefore contributing to both $u_i^{(p)}(t)$ and $u_i^{(n)}(t)$. Moreover, our analysis simply considers raw matching of positive and negative words, without making any attempt to identify expressions like negations or sarcasm.

We measure emotion in two ways based on these definitions: (i) the rate of status updates that contain words, (ii) the rate of status updates that contain negative words.

Consider a user i and a day t such that $u_i(t) \neq 0$. The positive rate of user i on day t is defined as

$$y_i^{(p)}(t) = \frac{u_i^{(p)}(t)}{u_i(t)},$$

that is, the fraction of status updates with at least one positive word. Note that $0 \leq y_i^{(p)}(t) \leq 1$.

Similarly, the negative rate of user i on day t is defined as

$$y_i^{(n)}(t) = \frac{u_i^{(n)}(t)}{u_i(t)},$$

that is, the fraction of status updates with at least one negative word. Note that $0 \leq y_i^{(n)}(t) \leq 1$.

By averaging these quantities over all users in city g , we obtain the average positive rate and negative rate of that city. Let S_g be the set of n_g users i in city g such that $u_i(t) \neq 0$. That is,

$$\begin{aligned} \bar{y}_g^{(p)}(t) &= \frac{1}{n_g} \sum_{i \in S_g} y_i^{(p)}(t), \\ \bar{y}_g^{(n)}(t) &= \frac{1}{n_g} \sum_{i \in S_g} y_i^{(n)}(t), \end{aligned}$$

Table 4.1 shows mean values for each of these emotion variables, and Table 4.2 shows the aggregates by city.

The variables $\bar{Y}_g^{(p)}(t)$ and $\bar{Y}_g^{(n)}(t)$ of the model in equation (4.4) are given by

$$\begin{aligned} \bar{Y}_g^{(p)}(t) &= \sum_i y_i^{(p)}(t) \frac{1}{n_g} \sum_{j \in S_g} \frac{1}{\delta_j(t)} a_{ij}(t), \\ \bar{Y}_g^{(n)}(t) &= \sum_i y_i^{(n)}(t) \frac{1}{n_g} \sum_{j \in S_g} \frac{1}{\delta_j(t)} a_{ij}(t). \end{aligned}$$

Table 4.1. Summary statistics for each emotional and meteorological variable, computed considering one observation for each city-day pair.

Summary of Emotion and Meteorological Variables				
	Mean	Standard Deviation	Minimum	Maximum
Positive rate	0.407	0.0445	0.116	0.614
Negative rate	0.213	0.0329	0.0388	0.440
Weather posts	0.0653	0.0347	0	0.527
Rainfall indicator	0.257	0.437	0	1

4.2.2 Exogenous control variable

Our method relies on the availability of an exogenous variable that affects the semantic expression of a person’s friends but not (directly) the semantic expression of the person. We call this variable the “instrument.” Our model characterizes how a change in the instrument induces a change in friends’ semantic expression, and how the induced change predicts a change in the person’s semantic expression.

There are many sources of exogenous variation in the world, but we chose rainfall as the instrument, relying on data from the National Climatic Data Center (NCDC, <http://www.ncdc.noaa.gov>). For each city g we consider the NCDC station closest to it, and let $\bar{x}_g(t) = 1$ if that station recorded rainfall on day t , and zero otherwise. For each subject $i \in S_g(t)$, let $x_i(t) = \bar{x}_g(t)$, that is, a binary indicator variable of rainfall in city g . We focus on rainfall as the instrument for several reasons. First, its geographical resolution lends itself to the analysis of our geographically aggregated model. Second individuals in the same city tend to experience the same weather on a given day. We show it is a robust instrument in the sense that it captures enough variation of the endogenous explanatory variable (friends’ emotional expression). Other meteorological variables would have been a valid alternative. The identification of valid instruments is challenging and finding a systematic way to characterize them is key to apply our method to more

Table 4.2. Number of rainy days and number of total observed days for each of the 100 cities in the dataset.

Summary of Rainfall in each City											
City Code	Num Days	Rainy Days	City Code	Num Days	Rainy Days	City Code	Num Days	Rainy Days	City Code	Num Days	Rainy Days
ABQ	120	1180	CVG	379	1180	LEX	236	1180	PIT	378	1180
ANA	424	1180	DAL	257	1180	LGB	404	1180	PLA	386	1180
ANC	278	1180	DEN	217	1180	LNK	274	1180	RAL	94	1180
ATL	139	1180	DFW	350	1180	LRD	358	1180	RDU	241	1180
AUR	236	1180	DTT	566	1180	MCI	280	1180	RIV	84	1180
AUS	243	1180	ELP	132	1180	MEM	131	1180	RNO	377	1180
AWO	94	1180	EWR	217	1180	MES	303	1180	ROC	113	1180
BFL	438	1180	FAT	137	1180	MGM	102	1180	SAN	357	1180
BHM	303	1180	FWA	148	1180	MIA	446	1180	SAT	384	1180
BNA	207	1180	GAR	341	1180	MKE	558	1180	SBD	402	1180
BOI	344	1180	GEU	132	1180	MOD	132	1180	SCK	116	1180
BOS	460	1180	GKY	389	1180	MSN	408	1180	SDL	213	1180
BTR	420	1180	GSP	347	1180	MSP	257	1180	SEA	519	1180
BUF	226	1180	HIA	214	1180	MSY	450	1180	SFO	461	1180
BWI	207	1180	HND	519	1180	NHE	341	1180	SJC	370	1180
CAK	372	1180	HNL	84	1180	NYC	523	1180	SMF	443	1180
CHD	342	1180	HOU	373	1180	OAK	382	1180	SNA	183	1180
CHI	243	1180	HTS	323	1180	OKC	243	1180	SNP	177	1180
CHU	407	1180	ICT	153	1180	OMA	444	1180	STL	536	1180
CLE	131	1180	IND	256	1180	ORF	299	1180	TOL	196	1180
CLT	397	1180	JAX	267	1180	ORL	296	1180	TPA	305	1180
CMH	390	1180	JCY	414	1180	PDX	369	1180	TUL	334	1180
COS	388	1180	LAS	283	1180	PHL	92	1180	TUS	400	1180
CPK	365	1180	LAX	362	1180	PHX	311	1180	VIB	320	1180
CRP	562	1180	LBB	134	1180	PIE	154	1180	WAS	388	1180

general contexts.

For the instrumental variable regression described below we will make use of the variable

$$\bar{X}_g(t) = \sum_i x_i(t) \frac{1}{n_g} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t). \quad (4.1)$$

In particular, if user i is in city h then $x_i(t) = \bar{x}_h(t)$ (the user's own weather is the same as the average weather of all users in the same city).

4.2.3 Social network information

For each day $t \in T$, and subjects $i, j \in S(t)$, let $a_{i,j}(t) \in [0, 1]$ be the strength of the relationship from i to j on day t , which need not be symmetric. Also, let $\delta_i(t) = \sum_{j \in S(t)} a_{i,j}(t)$. We let $a_{i,j}(t) \in \{0, 1\}$, where $a_{i,j}(t) = 1$ denotes that i and j were friends on day t . In this case $\delta_i(t)$ is the degree of subject i on day t (that is, the total number of friends of the subject). Allowing $a_{i,j}(t)$ to have any value between zero and one would allow to assess the role of tie-strength.

4.3 Data

Our period of observation starts on January 1st 2009 and ends on March 31st 2012, for a total of 1185 consecutive days. Data for five days of 2009 (March 4th, June 24th, August 15th, September 13th, November 11th) was not available at the time of analysis, so we consider the remaining 1180 days.

Data were collected from the Facebook online social network, and data were analyzed in aggregate within Facebook’s data centers. Researchers did not access any personal information.

For each day in the period of observation, we consider all Facebook users in the 100 most populous US cities, and their status updates. Table 4.3 reports the list of the cities, each paired with the corresponding three-letter code used in the figures (airport codes in most cases). In particular, the subpopulation of Facebook users in a given city contains all users that (i) chose English as the language in which they view the website, (ii) selected United States as Country in their profile settings, (iii) can be matched to city g by IP-based geographic location. We build separate user pools for different days to allow us to take user mobility into account, since on any particular day a user might travel or move to a new city.

For each Facebook user, we measured emotion using all status updates as explained above. We also measured social contacts for each day in the observation period, letting $a_{ij}(t) = 1$ for all pairs of users i and j who were “friends” with one another on day t , and 0 otherwise.

Table 4.4 summarizes our sample size by showing mean and standard deviation of the daily number of users, number of status updates, and friendship ties.

Figure 4.1 shows temporal and geographical variation in emotions expressed by Facebook users in 2011, and a representation of between-city friendship ties.

Table 4.3. List of the 100 cities in the dataset.

List of US cities and codes					
Code	City	Code	City	Code	City
ABQ	Albuquerque, NM	GAR	Garland, TX	OKC	Oklahoma City, OK
ANA	Anaheim, CA	GEU	Glendale, AZ	OMA	Omaha, NE
ANC	Anchorage, AK	GKY	Arlington, TX	ORF	Norfolk, VA
ATL	Atlanta, GA	GSP	Greensboro, NC	ORL	Orlando, FL
AUR	Aurora, CO	HIA	Hialeah, FL	PDX	Portland, OR
AUS	Austin, TX	HND	Henderson, NV	PHL	Philadelphia, PA
AWO	Arlington, VA	HNL	Honolulu, HI	PHX	Phoenix, TX
BFL	Bakersfield, CA	HOU	Houston, TX	PIE	St Petersburg, FL
BHM	Birmingham, AL	HTS	Huntington, WV	PIT	Pittsburgh, PA
BNA	Nashville, TN	ICT	Wichita, KS	PLA	Plano, TX
BOI	Boise, ID	IND	Indianapolis, IN	RAL	Raleigh, NC
BOS	Boston, MA	JAX	Jacksonville, FL	RDU	Durham, NC
BTR	Baton Rouge, LA	JCY	Jersey City, NJ	RIV	Riverside, CA
BUF	Buffalo, NY	LAS	Las Vegas, NV	RNO	Reno, NV
BWI	Baltimore, MD	LAX	Los Angeles, CA	ROC	Rochester, NY
CAK	Akron, OH	LBB	Lubbock, TX	SAN	San Diego, CA
CHD	Chandler, AZ	LEX	Lexington, KY	SAT	San Antonio, TX
CHI	Chicago, IL	LGB	Long Beach, CA	SBD	San Bernardino, CA
CHU	Chula Vista, CA	LNK	Lincoln, NB	SCK	Stockton, CA
CLE	Cleveland, OH	LRD	Laredo, TX	SDL	Scottsdale, AZ
CLT	Charlotte, NC	MCI	Kansas City, MO	SEA	Seattle, WA
CMH	Columbus, OH	MEM	Memphis, TN	SFO	San Francisco, CA
COS	Colorado Springs, CO	MES	Mesa, CA	SJC	San Jose, CA
CPK	Chesapeake, VA	MGM	Montgomery, AL	SMF	Sacramento, CA
CRP	Corpus Christi, TX	MIA	Miami, FL	SNA	Santa Ana, CA
CVG	Cincinnati, OH	MKE	Milwaukee, WI	SNP	St Paul, MN
DAL	Dallas, TX	MOD	Modesto, CA	STL	St Louis, MO
DEN	Denver, CO	MSN	Madison, WI	TOL	Toledo, OH
DFW	Fort Worth, TX	MSP	Minneapolis, MN	TPA	Tampa, FL
DTT	Detroit, MI	MSY	New Orleans, LA	TUL	Tulsa, OK
ELP	El Paso, TX	NHE	North Hempstead, NY	TUS	Tucson, AZ
EWR	Newark, NJ	NYC	New York, NY	VIB	Virginia Beach, VA
FAT	Fresno, CA	OAK	Oakland, CA	WAS	Washington, DC
FWA	Fort Wayne, IN				

Table 4.4. Summary statistics of the dataset. For each day in the period of observation (a set of 1180 days from January 2009 to March 2012) all Facebook users that are English-speakers and geolocated within the 100 most populous US cities are included. Assuming that each user posts either one or zero status updates on a day, the average number of status updates per user per day is $\alpha = 0.206$.

Quantity	Mean	Standard Dev.
Number of users (daily)	9,903,993	3,447,776
Number of users who updated (daily)	2,042,996	775,162
Number of friendships (daily)	52,787,239	25,118,462

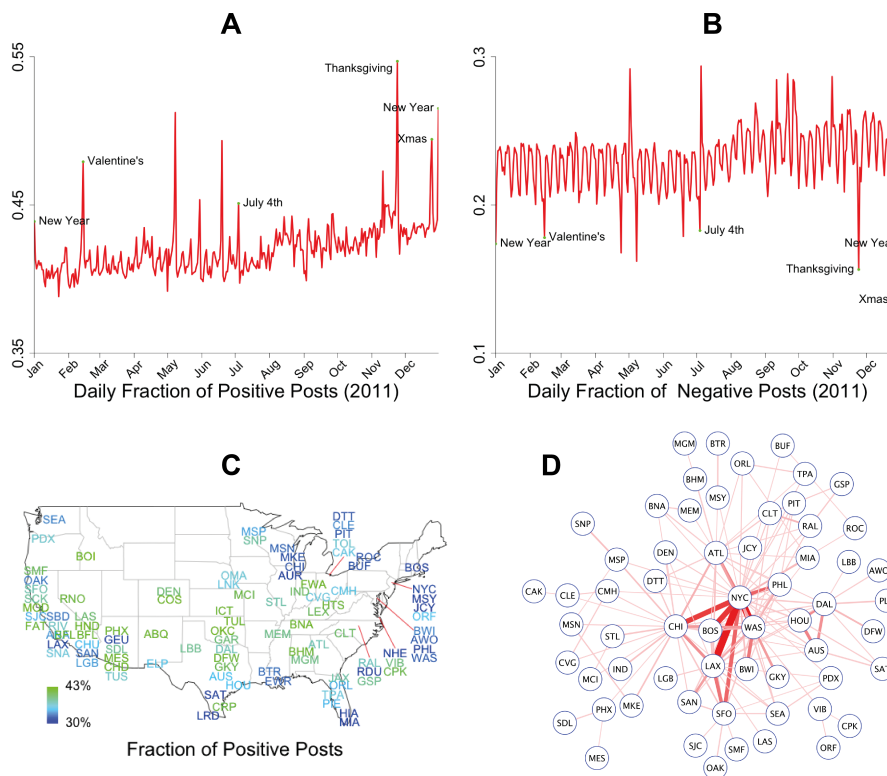


Figure 4.1. Temporal and geographical variation in emotions expressed by Facebook users in 2011 as measured by (a) the fraction of status updates containing positive emotion words; (b) the fraction of status updates containing negative emotion words. Extreme values are noted for holidays. (c) A map of the U.S. with approximate locations of the 100 most populous cities (represented by airport code) and their average fraction of posts with positive emotions (blue is less and green is more). (d) Network of between-city ties for all pairs of cities with at least 50,000 friendships. Darker, thicker lines indicate more friendship ties (maximum = 1,210,769).

4.4 A model of emotional contagion

Let $y_i(t)$ be the emotional expression of individual i at time t . Let $a_{ij}(t)$ be the strength of the relationship from individual i to individual j at time t . Note that $a_{ij}(t)$ need not be symmetric (i may perceive a stronger relationship with j than j does with i), and it allows for temporal variations. Let $\delta_i(t) = \sum_j a_{ij}(t)$ be the degree of individual i at time t .

In the simplest case, $a_{ij}(t)$ can take binary values, 1 designating that a relationship

between i and j exists at time t , 0 designating that it does not. Under this assumption, $\delta_j(t)$ is simply the number of her social contacts of i at time t .

Suppose there are three kinds of exogenous factors that affect emotion. First, there are factors that are time-varying and affect everyone equally (like holidays, for example). We denote these with a fixed effect $\theta(t)$ for each time period t . Second, there are factors that are time-invariant and specific to an individual (such as a person's baseline personality). We denote these with a fixed effect f_j for each individual j . Third, some factors are *both* time-varying and specific to an individual (like the weather). We denote these by $x_j(t)$ for each individual j and time period t .

In addition, suppose there is an endogenous factor that affects the emotion of each individual in proportion to the strength of the relationship between j and her social contacts. That is, each individual j is affected by the specific emotion on day t of each individual i to whom she is connected.

Assuming a memoryless model where individuals influence each other only within a time period t and not across time periods, we can specify a *linear model* for the emotion y of individual j on day t :

$$y_j(t) = \theta(t) + f_j + \beta x_j(t) + \gamma \frac{1}{\delta_j(t)} \sum_i a_{ij}(t) y_i(t) + \varepsilon_j(t), \quad (4.2)$$

where β indicates the strength and direction of influence of the time-varying exogenous factor, γ indicates the strength and direction of social influence, and $\varepsilon_j(t)$ includes unobserved variables and noise.

Observe that the model in equation (4.2) assumes that influence is averaged over all social contacts and therefore inversely proportional to the cumulative weight $\delta_j(t)$ of all j 's relationships. If $a_{ij}(t)$ takes binary values, then this implies that the influence from i to j is inversely proportional to the number of j 's social contacts. This assumption

is based on the idea that an individual with many social contacts is less likely to be influenced by each single contact i than an individual with few social contacts.

We are interested in estimating the value of the influence factor γ , which is difficult due to the inherent *feedback* present in the process of emotional contagion. Correlation in emotions may not only be the result from pairwise mutual influence, but also from cycles in the social network. For example, i might influence k 's emotional expression, which in turn affects j 's emotional expression, and so on. We address the inherent endogeneity of contagion in Section 4.6 by using instrumental variable regression [9].

A second difficulty here is the large size of our data set. We would like to apply our model to the longitudinal content generated by millions of users with billions of friends over hundreds of days. We address this difficulty in Section 4.5, where we propose a method to estimate the individual-level parameter γ using aggregated data. A key to this method is to identify a unit of analysis in which many individuals within the same subpopulation are affected by the same exogenous variables. For example, individuals i and j may be in the same city g and therefore experience the same weather, traffic conditions, sporting event outcomes, and so on. Or they may be in different cities g and h , in which case their different exposures to exogenous factors may help us to identify how one person affects another. In our aggregated model, we leverage these between-unit social ties to consider how a factor in city g affects individual i , which in turn affects individual j who was not exposed directly to that factor because she is in city h . In other words, if it rains on you in New York, does it make your friends in San Diego less happy?

4.5 Aggregating the model

The model in equation (4.2) can be computationally demanding in big data sets, since there is one observation for each individual-time pair. We therefore simplify the

model further by averaging equation (4.2) over all n_g individuals in a given subpopulation S_g who are in city g .

$$\frac{1}{n_g} \sum_{j \in S_g} y_j(t) = \frac{1}{n_g} \sum_{j \in S_g} \left(\theta(t) + f_j + \beta x_j(t) + \gamma \frac{1}{\delta_j(t)} \sum_i a_{ij}(t) y_i(t) + \varepsilon_j(t) \right). \quad (4.3)$$

We can change the notation to make things clearer.

Let $\bar{y}_g(t) = \frac{1}{n_g} \sum_{j \in S_g} y_j(t)$ be the average emotion at time t for all individuals in subpopulation S_g .

Let $\bar{f}_g(t) = \frac{1}{n_g} \sum_{j \in S_g} f_j$ be the average individual fixed effects for all individuals in subpopulation S_g (this is therefore a city-level fixed effect).

Let $\bar{x}_g(t) = \frac{1}{n_g} \sum_{j \in S_g} x_j(t)$ be the average exogenous variable at time t for all individuals in subpopulation S_g (this is therefore a city-level exogenous variable).

Let $\bar{Y}_g(t) = \frac{1}{n_g} \sum_{j \in S_g} \frac{1}{\delta_j(t)} \sum_i a_{ij}(t) y_i(t)$. We can exchange the ordering of the summations and write

$$\bar{Y}_g(t) = \sum_i y_i(t) \frac{1}{n_g} \sum_{j \in S_g} \frac{1}{\delta_j(t)} a_{ij}(t)$$

Observe that the term $\frac{1}{n_g} \sum_{j \in S_g} \frac{1}{\delta_j(t)} a_{ij}(t)$ represents the average strength of the relationship between i and an individual in city g . Therefore, $\bar{Y}_g(t)$ represents the average emotional influence at time t on an individual in city g .

The model in equation (4.3) can now be written as

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \beta \bar{x}_g(t) + \gamma \bar{Y}_g(t) + \bar{\varepsilon}_g(t) \quad (4.4)$$

where $\bar{\varepsilon}_g(t) = \frac{1}{n_g} \sum_{j \in S_g} \varepsilon_j(t)$ is a city-specific error for all individuals j who are in city g . Assuming $\varepsilon_j(t)$ independent normally distributed with zero mean and variance σ^2 , $\bar{\varepsilon}_g(t)$

will also be normally distributed with mean 0, but it will have a city-specific variance σ^2/n_g . Notice that this indicates the variance is inversely proportional to the number of individuals in a city. As we describe below, we can use the equation for the variance explicitly to weight each observation in the model.

4.6 Model estimation

We are interested in estimating the parameters of the model in equation (4.4), which is simply an aggregated restatement of the individual-level model in equation (4.2). To recap, this model is:

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \beta \bar{x}_g(t) + \gamma \bar{Y}_g(t) + \bar{\varepsilon}_g(t),$$

and we are primarily interested in estimating the effect of emotional contagion (γ). The dependent variable $\bar{y}_g(t)$ is the average emotion of users in city g on day t , the independent variable $\bar{Y}_g(t)$ is the average emotion of the friends of these users, $\bar{x}_g(t)$ is a binary indicator variable for rainfall in city g , and $\theta(t)$ and \bar{f}_g are fixed effects for each day and each city.

Note that we can estimate γ for contagion of either positive and negative emotion, and we can also see if these two emotions tend to inhibit one another by estimating the effect of friends' positive emotion on users' negative emotion and vice versa.

An observation period of 1180 days and a set of 100 cities results in a model with 118,000 observations, each corresponding to a city-day pair. The parameters that need to be estimated are the coefficients β and γ , 1180 fixed effects for the days, and 100 fixed effects for the cities.

Since one of the explanatory variables of the model in equation (4.4), $\bar{Y}_g(t)$, is an endogenous variable (i.e. it is correlated both to the dependent variable $\bar{y}_g(t)$ and to

the error term $\bar{\epsilon}_g(t)$), ordinary least squares regression would yield biased coefficient estimates. We therefore use *instrumental variable regression* [9].

Instrumental variable regression is an estimation method that can produce consistent and unbiased estimates when one of the explanatory variables is correlated with the error terms in the model equation. This is the case when there is *reciprocal causality* from the dependent variable to an explanatory variable (in our case, users affect their friends and vice versa), when one or more relevant explanatory variables are omitted from the model, or when the covariates are affected by measurement errors. If an *instrument* is available that predicts the endogenous variable, then consistent and unbiased estimates can be obtained. In a linear model, an instrument for an endogenous explanatory variable v is a variable z that does not appear in the model equation, is correlated with v (conditional on all the exogenous explanatory variables) and is not correlated with the error term [9].

In our model, an instrument for the endogenous explanatory variable $\bar{Y}_g(t)$ is an exogenous variable z that is not correlated to the error term in equation (4.4), that is $Cov(z, \hat{\epsilon}_g(t)) = 0$, and is partially correlated to $\bar{Y}_g(t)$ when controlling for the other exogenous explanatory variables. In the context of our model, we can write:

$$\bar{Y}_g(t) = \theta'(t) + \bar{f}'_g + \beta_2 \bar{x}_g(t) + \beta_1 z + v_g(t), \quad (4.5)$$

where $v_g(t)$ is an error term not correlated to any regressors and $\theta'(t)$ and \bar{f}'_g are separately estimated time and subpopulation fixed effects.

Equation (4.5) can be seen as the linear projection of $\bar{Y}_g(t)$ on the space of all the exogenous variables. Substituting equation (4.5) into equation (4.4) yields:

$$\bar{y}_g(t) = (\theta(t) + \gamma \theta'(t)) + (\bar{f}_g + \gamma \bar{f}'_g) + (\beta + \gamma \beta_2) \bar{x}_g(t) + \gamma \beta_1 z + \bar{\epsilon}'_g(t), \quad (4.6)$$

where the error term is uncorrelated with all the explanatory variables.

We use the variable $\bar{X}_g(t)$ defined in equation (4.1) as the instrument (z) for $\bar{Y}_g(t)$, as it is uncorrelated with the error term in equation (4.4) and it is partially correlated to $\bar{Y}_g(t)$ (see Tables 4.5 to 4.8 for details). Specifically, we utilize rainfall experienced by the friends of users in city g to predict the emotion of those friends since it directly affects their mood.

The procedure above is equivalent to estimating the model in equation (4.4) using two stage least-squares (2SLS) regression. The first stage regression estimates a model of the form

$$\bar{Y}_g(t) = \theta'(t) + \bar{f}'_g + \beta_1 \bar{X}_g(t) + \beta_2 \bar{x}_g(t) + \varepsilon'_g(t). \quad (4.7)$$

The second stage regression uses the predicted values $\bar{Y}_g^{pred}(t)$ from the first stage to estimate the model

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \beta \bar{x}_g(t) + \gamma \bar{Y}_g^{pred}(t) + \bar{\varepsilon}_g(t). \quad (4.8)$$

Finally, recall that the variance of the $\bar{\varepsilon}_g(t)$ error term is proportional to $\frac{1}{n_g}$ where n_g is the number of individuals in a city. We therefore weight each observation by the corresponding value of n_g .

A key assumption of instrumental variables regression is the exclusion restriction – the instrument must not directly influence the dependent variable. In our case, some of the users' friends are experiencing the same weather as the users because they are in the same city. Therefore, in order to break any possible correlation between friends' rainfall $\bar{X}_g(t)$ and users' rainfall $\bar{x}_g(t)$, we only consider observations for city-day pairs (g, t) such that $\bar{x}_g(t) = 0$ (that is, it did not rain in city g on day t). This results in dropping 30,300

observations, for a total of 87,700 remaining observations. Conditional on $\bar{x}_g(t) = 0$, equations (4.7) and (4.8) can be respectively written as

$$\bar{Y}_g(t) = \theta'(t) + \bar{f}'_g + \beta_1 \bar{X}_g(t) + \varepsilon'_g(t), \quad (4.9)$$

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \gamma \bar{Y}_g^{pred}(t) + \bar{\varepsilon}_g(t). \quad (4.10)$$

Note that since $\bar{x}_g(t) = 0$, there is no rainfall for either the user or the user's friends who are in the same city. This means that the instrument $\bar{X}_g(t)$ now depends *only on friends who are in different cities* (not in city g).

Results of the first and second stage regressions are shown in Figure 4.2.

Tables 4.5 to 4.8 report the estimates (with standard errors, t-statistics, 95% confidence intervals, and diagnostic statistics) for the first and second stage of the 2SLS regression for the model in equation (4.4) (fixed effects estimates are not reported due to their number). The estimates of the emotional transmission parameter γ from the second stage regression are always significantly different than zero, and the two positive coefficients support the hypothesis of contagion. When users' friends post positive status updates, it increases their own positive updates. When users' friends post negative status updates, it increases their own negative updates. At the same time, the two negative coefficients for γ support the idea that opposite moods have an inhibitory effect. When users' friends post positive status updates, it decreases their own negative updates. When users' friends post negative status updates, it decreases their own positive updates.

In order to assess the quality of the estimates obtained via instrumental variable regression, we also compute diagnostic statistics.

First, we need to verify that the model is not underidentified. The Kleinbergen-Paap rk LM statistic allows to test the null hypothesis of underidentification [128], and all of our tests reject the null (test statistics are reported in the caption underneath each

table).

Second, we need to verify that the instruments are good predictors of the endogenous explanatory variable in the first-stage regression (otherwise the instruments are considered *weak*). Weak instruments would cause poor predicted values in the first-stage regression (for example, little variation) and presumably poor estimation in the second-stage regression. To ensure the instruments are not weak, the Cragg-Donald Wald F statistic must exceed the critical threshold suggested by Stock and Yogo [200].

For robustness, we also tested a version of the model in equation (4.4) that only considers observations for city-day pairs (g, t) such that $\bar{x}_g(t) = 1$ (that is, it rained in city g on day t). This results in dropping 87,700 observations, for a total of 30,300 remaining observations. Tables 4.9 to 4.12 report the estimates (with standard errors, t-statistics, 95% confidence intervals, and diagnostic statistics) for the first and second stage of the 2SLS regression, which are substantially the same as the ones in Tables 4.5 to 4.8, with overlapping 95% CI. These results suggest that the users' own experience of rainfall does not affect emotional contagion online.

4.6.1 Placebo test

If our procedure is correctly estimating social influence, we would not expect to be able to predict users' emotion using future friends' weather and emotion. Here, we test a placebo model by using the same instrumental variables procedure described above to estimate the effect of *future* friends' rainfall on users' emotion *today*. We arbitrarily choose $t + 30$ as a point in time far enough in the future that friends' rainfall then will not be correlated with friends' rainfall at time t . We then modify the equation in (4.4) to shift the independent variable forward by 30 days:

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \beta \bar{x}_g(t) + \gamma \bar{Y}_{g,t+30} + \bar{\epsilon}_g(t), \quad (4.11)$$

Table 4.5. Instrumental variable regression estimates: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 25.507 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 324.053, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Emotion measure: positive rate (non rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.0119	0.00207	-5.75	0.000	-0.0160	-.00781
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	1.752	0.122	14.39	0.000	1.514	1.991

Table 4.6. Instrumental variable regression estimates: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 24.598 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 505.398, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Emotion measure: negative rate (non rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.0116	0.00195	5.97	0.000	0.00776	0.0155
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	1.288	0.0486	26.53	0.000	1.193	1.383

Table 4.7. Instrumental variable regression estimates: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 25.507 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 324.053, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

How friends' positive rate affects users' negative rate (non rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.0119	0.00207	-5.75	0.000	-0.0160	-0.00781
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	-1.255	0.227	-5.52	0.000	-1.701	-0.809

Table 4.8. Instrumental variable regression estimates: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 24.598 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 505.398, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

How friends' negative rate affects users' positive rate (non rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.0116	0.00195	5.97	0.000	0.00776	0.0155
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	-1.798	0.271	-6.62	0.000	-2.330	-1.266

Table 4.9. Instrumental variable regression estimates: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations). The Kleibergen-Paap rk LM statistic is 13.531 ($p = 0.0002$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 78.189, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Emotion measure: positive rate (rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.00985	0.00268	-3.68	0.000	-0.0152	-0.00454
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	1.794	0.233	7.70	0.000	1.338	2.251

Table 4.10. Instrumental variable regression estimates: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations). The Kleibergen-Paap rk LM statistic is 11.333 ($p = 0.0008$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 102.297, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Emotion measure: negative rate (rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.00973	0.00281	3.47	0.001	0.00416	0.0153
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	1.473	0.134	10.97	0.000	1.210	1.736

Table 4.11. Instrumental variable regression estimates: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations). The Kleibergen-Paap rk LM statistic is 13.531 ($p = 0.0002$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 78.189, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

How friends' positive rate affects users' negative rate (rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.00985	0.00268	-3.68	0.000	-0.0152	-0.00454
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	-1.456	0.475	-3.06	0.002	-2.387	-0.524

Table 4.12. Instrumental variable regression estimates: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 1$ are considered (30,300 total observations). The Kleibergen-Paap rk LM statistic is 11.333 ($p = 0.0008$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 102.297, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

How friends' negative rate affects users' positive rate (rainy days)						
Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.00973	0.00281	3.47	0.001	0.00416	0.0153
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	-1.816	0.550	-3.30	0.001	-2.895	-0.738




Figure 4.2. Models estimates. (a) Difference in emotional expression between days with and without rain. Estimates derived from first stage regressions of each measure of emotion on a binary measure of rainfall. (b) Estimates of emotional contagion between friends from the second stage of an instrumental variables regression from four separate models. The results show that rain affects emotional expression, both positive and negative posts are contagious, and positive posts tend to inhibit negative posts and vice versa. All models include fixed effects for city and day, average friends weather in other cities, and standard errors clustered by city and day. Vertical bars show 95% confidence intervals.

We conduct instrumental variable regression using $\bar{X}_{g,t+30}$ as an instrument for $\bar{Y}_{g,t+30}$. Tables 4.13 to 4.16 report the estimates (with standard errors, t-statistics, 95% confidence intervals, and diagnostic statistics) for the first and second stage of the 2SLS regression. The estimates of γ from the second stage regression are not statistically significant and they are much lower in magnitude than those estimated for the model in equation (4.4).

4.6.2 Controlling for topic contagion

One concern is that our estimates of emotional contagion are actually estimates of topic contagion. Friends who post more negatively when it rains may be posting about the weather itself, and users may respond with their own statuses about weather.

Table 4.13. Placebo test: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations). The Kleibergen-Paap *rk* LM statistic is 28.711 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 265.910, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.11) (non rainy days)						
Emotion measure: positive rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}_{t+30}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}_{t+30}	-0.0118	0.00191	-6.19	0.000	-0.0156	-0.00804
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}_{t+30}^{(p)}$	-0.112	0.177	-0.63	0.526	-0.459	0.235

Table 4.14. Placebo test: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations). The Kleibergen-Paap *rk* LM statistic is 26.552 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 458.685, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.11) (non rainy days)						
Emotion measure: negative rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}_{t+30}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}_{t+30}	0.0120	0.00188	6.42	0.000	0.00832	0.0158
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}_{t+30}^{(n)}$	-0.185	0.126	-1.46	0.143	-0.432	0.0627

Table 4.15. Placebo test: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations). The Kleibergen-Paap *rk* LM statistic is 28.711 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 265.910, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.11) (non rainy days)						
Friends' positive rate to users' negative rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}_{t+30}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}_{t+30}	-0.0118	0.00191	-6.19	0.000	-0.0156	-0.00804
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}_{t+30}^{(p)}$	0.188	0.121	1.55	0.120	-0.0493	0.425

Table 4.16. Placebo test: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ and $\bar{x}_{g,t+30} = 0$ are considered (67,493 total observations). The Kleibergen-Paap *rk* LM statistic is 26.552 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 458.685, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.11) (non rainy days)						
Friends' negative rate to users' positive rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}_{t+30}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}_{t+30}	0.0120	0.00188	6.42	0.000	0.00832	0.0158
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}_{t+30}^{(n)}$	0.110	0.172	0.64	0.521	-0.226	0.447

This would not undermine our statistical results on contagion, but it might change our interpretation if we discovered that topics were driving the similarity in word choice by users and friends.

To address this issue, we created a dictionary of weather terms based on a meteorological glossary supplied by NOAA (<http://www.erh.noaa.gov/box/glossary.htm>). We then crowdsourced this dictionary to approximately 100 students, postdocs, and professors asking for additional suggestions. The resulting list is not exhaustive, but we expect it will allow us to detect most status updates that are on a weather-related topic. The full list of terms can be found in Table 4.17.

Recall that $U_i(t)$ represents the status updates of user i on day t , and let $u_i^{(w)}(t)$ be the number of status updates in $U_i(t)$ that contain at least one word from our dictionary of weather terms. If $u_i(t) \neq 0$, let $w_i(t) = u_i^{(w)}(t)/u_i(t)$ be the fraction of i 's status updates related to weather, and let $\bar{w}_g(t) = \frac{1}{n_g} \sum_{i \in S_g} w_i(t)$ be the average over city g . We can now use this variable to control for the tendency to post status updates about the weather by adding it to equation (4.4):

$$\bar{y}_g(t) = \theta(t) + \bar{f}_g + \lambda \bar{w}_g(t) + \gamma \bar{Y}_g(t) + \bar{\epsilon}_g(t). \quad (4.12)$$

Tables 4.18 to 4.21 report the estimates (with standard errors, t-statistics, 95% confidence intervals, and diagnostic statistics) for the first and second stage of the 2SLS regression. The negative estimates for λ suggest that increased usage of weather words is generally associated with decreased emotional expression. However, the relationships are weak and sometimes insignificant, and more importantly, the estimates for the emotional transmission parameter γ remain substantially the same as the estimates from model (4.4) without controls. These results indicate that posting on the topic of weather is not driving the relationship in use of emotional words between users and their friends.

Table 4.17. Terms used to identify status updates on the topic of weather.

aerovane air airstream altocumulus altostratus anemometer anemometers anticyclone
 anticyclones arctic arid aridity atmosphere atmospheric autumn autumnal balmy
 baroclinic barometer barometers barometric blizzard blizzards blustering blustery
 blustery breeze breezes breezy brisk calm celsius chill chilled chillier chilliest
 chilly chinook cirrocumulus cirrostratus cirrus climate climates cloud cloudburst
 cloudbursts cloudier cloudiest clouds cloudy cold colder coldest condensation contrail
 contrails cool cooled cooling cools cumulonimbus cumulus cyclone cyclones damp damp
 damper damper dampest dampest degree degrees deluge dew dews dewy doppler downburst
 downbursts downdraft downdrafts downpour downpours dried drier dries driest drizzle
 drizzled drizzles drizzly drought droughts dry dryline fall fahrenheit flood flooded
 flooding floods flurries flurry fog fogbow fogbows fogged fogging foggy fogs forecast
 forecasted forecasting forecasts freeze freezes freezing frigid frost frostier
 frostiest frosts frosty froze frozen gale gales galoshes gust gusting gusts gusty
 haboob haboobs hail hailed hailing hails haze hazes hazy heat heated heating heats
 hoarfrost hot hotter hottest humid humidity hurricane hurricanes ice iced ices icing
 icy inclement landspout landspouts lightning lightnings macroburst macrobursts maelstrom
 mercury meteorologic meteorologist meteorologists meteorology microburst microbursts
 microclimate microclimates millibar millibars mist misted mists misty moist moisture
 monsoon monsoons mugginess muggy nexrad nippy NOAA nor'easter nor'easters noreaster
 noreasters overcast ozone parched parching pollen precipitate precipitated precipitates
 precipitating precipitation psychrometer radar rain rainboots rainbow rainbows raincoat
 raincoats rained rainfall rainier rainiest raining rains rainy sandstorm sandstorms
 scorcher scorching searing shower showering showers skiff sleet slicker slickers slush
 slushy smog smoggier smoggiest smoggy snow snowed snowier snowiest snowing snowmageddon
 snowpocalypse snows snowy spring sprinkle sprinkles sprinkling squall squalls squally
 storm stormed stormier stormiest storming storms stormy stratocumulus stratus
 subtropical summer summery sun sunnier sunniest sunny temperate temperature tempest thaw
 thawed thawing thaws thermometer thunder thundered thundering thunders thunderstorm
 thunderstorms tornadic tornado tornadoes tropical troposphere tsunami turbulent twister
 twisters typhoon typhoons umbrella umbrellas vane warm warmed warming warms warmth
 waterspout waterspouts weather wet wetter wettest wind windchill windchills windier
 windiest windspeed windy winter wintry wintry

Table 4.18. Controlling for topic contagion: effect of friends positive emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap *rk* LM statistic is 17.750 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald *F* statistic is 169.315, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.12) (non rainy days)						
Emotion measure: positive rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	<i>t</i>	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.00888	0.00199	-4.46	0.000	-0.0128	-0.00492
Users' weather rate \bar{w}	-0.0186	0.00322	-5.78	0.000	-0.0250	-0.0122
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(p)}$	Coefficient	Error	<i>t</i>	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	1.205	0.0974	12.37	0.000	1.0140	1.396
Users' weather rate \bar{w}	-0.0399	0.00301	-13.25	0.000	-0.0458	-0.0340

Table 4.19. Controlling for topic contagion: effect of friends negative emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 15.799 ($p = 0.0001$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 199.681, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.12) (non rainy days)						
Emotion measure: negative rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.00749	0.00175	4.29	0.000	0.00403	0.0110
Users' weather rate \bar{w}	0.0252	0.00496	5.09	0.000	0.0154	0.0351
SECOND STAGE		Standard			95% Confidence Interval	
Users' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	1.509	0.0986	15.30	0.000	1.315	1.702
Users' weather rate \bar{w}	-0.0157	0.00305	-5.14	0.000	-0.0217	-0.00971

4.7 Quantifying the total effect of a user on her friends

Consider a user j and assume she posts a single status update during day t . For presentation, we consider negative emotions and compare the case in which j 's status update contains a negative word ($y_j(t) = 1$) versus the case it does not ($y_j(t) = 0$). We estimate the additional number of negative status updates posted by j 's friends conditional on $y_j(t) = 1$ versus $y_j(t) = 0$.

According to the individual level model (4.2), the emotional contagion from j to i is given by $c_{ijt} = \gamma a_{ij}(t) y_j(t) / \delta_i(t)$ for each user i who posted on day t (we assume that each user i posted either one or zero status updates). Conditional on $y_j(t) = 1$ and

Table 4.20. Controlling for topic contagion: effect of friends positive emotion on user negative emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 17.750 ($p = 0.0000$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 169.315, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.12) (non rainy days)						
Friends' positive rate to users' negative rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	-0.00888	0.00199	-4.46	0.000	-0.0128	-0.00492
Users' weather rate \bar{w}	-0.0186	0.00322	-5.78	0.000	-0.0250	-0.0122
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(p)}$	-1.274	0.302	-4.22	0.000	-1.866	-0.681
Users' weather rate \bar{w}	-0.00134	0.0113	-0.12	0.905	-0.0234	0.0207

on $y_j(t) = 0$ respectively, this term is

$$c_{ijt}^{(1)} = \gamma a_{ij}(t) / \delta_i(t),$$

$$c_{ijt}^{(0)} = 0.$$

Summing over all users i who posted on day t , the total emotional contagion of user j conditional on $y_j(t) = 1$ is

$$C_j^{(1)}(t) = \sum_i c_{ijt}^{(1)} = \gamma \sum_i a_{ij}(t) / \delta_i(t),$$

while conditional on $y_j(t) = 0$ it is $C_j^{(0)}(t) = 0$. The difference in number of negative status updates posted by j 's friends conditional on $y_j(t) = 1$ versus $y_j(t) = 0$ can be

Table 4.21. Controlling for topic contagion: effect of friends negative emotion on user positive emotion. Observations such that $\bar{x}_g(t) = 0$ are considered (87,700 total observations). The Kleibergen-Paap rk LM statistic is 15.799 ($p = 0.0001$) suggesting the regression is not underidentified [128]. The Cragg-Donald Wald F statistic is 199.681, which exceeds the critical thresholds suggested by Stock and Yogo [200] to ensure the instruments are not weak. All statistics are robust to heteroskedasticity, autocorrelation, and clustering.

Model in Equation (4.12) (non rainy days)						
Friends' negative rate to users' positive rate - Instrument: binary indicator of rainfall						
FIRST STAGE		Standard			95% Confidence Interval	
Friends' emotion $\bar{Y}^{(n)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' rainfall \bar{X}	0.00749	0.00175	4.29	0.000	0.00403	0.0110
Users' weather rate \bar{w}	0.0252	0.00496	5.09	0.000	0.0154	0.0351
SECOND STAGE		Standard			95% Confidence Interval	
User' emotion $\bar{y}^{(p)}$	Coefficient	Error	t	$P > t $	Low	High
Friends' emotion $\bar{Y}^{(n)}$	-1.427	0.368	-3.88	0.000	-2.149	-0.706
Users' weather rate \bar{w}	-0.0264	0.0150	-1.76	0.078	-0.0557	0.00300

therefore quantified as

$$F_j(t) = C_j^{(1)}(t) - C_j^{(0)}(t) = \gamma \sum_i a_{ij}(t) / \delta_i(t) = \gamma A_j(t),$$

where $A_j(t) = \sum_i a_{ij}(t) / \delta_i(t)$ constitutes a measure of how influential user j is. In words, j 's cumulative effect on her friends is proportional to the coefficient of emotional contagion γ and her influence $A_j(t)$.

Observe that $A_j(t)$ can be computed exactly for each j and t . And note that this measure is increasing in the number of friends (more friends means more people might be influenced) and decreasing in the number of friends those friends have (if a friend has more friends, the user will on average have less influence on that friend).

The average user's effect $\bar{F}(t)$ can be computed as the average individual effect

Table 4.22. Average cumulative effect of a user on her friends. \bar{D}_t is the cumulative number of additional status updates posted by a user's friends and containing an emotion word (either positive or negative, depending on the variable choice) when the user posts an additional status updates containing an emotion word (either positive or negative, depending on the variable choice).

Average user emotional contagion effect – estimates and 95% CI				
User's emotion	Friend's emotion	\bar{D}_t	Lo95%	Hi95%
Positive rate	Positive rate	1.752	1.514	1.991
Negative rate	Negative rate	1.288	1.193	1.383
Positive rate	Negative rate	-1.255	-1.701	-0.809
Negative rate	Positive rate	-1.798	-2.330	-1.266

over all n users,

$$\begin{aligned}\bar{F}(t) &= \frac{1}{n} \sum_j F_j(t) = \gamma \frac{1}{n} \sum_j A_j(t) = \gamma \frac{1}{n} \sum_j \sum_i a_{ij}(t) \frac{1}{\delta_i(t)} \\ &= \gamma \frac{1}{n} \sum_i \frac{1}{\delta_i(t)} \sum_j a_{ij}(t) = \gamma \frac{1}{n} \sum_i \frac{1}{\delta_i(t)} \delta_i(t) = \gamma.\end{aligned}$$

The average user's effect $\bar{F}(t)$ and 95% CI for all four choices of emotions (user's positive/negative rate, friends' positive/negative rate) are shown in Table 4.22, and correspond directly to the estimates of γ in Tables 4.5 to 4.8. In other words, the γ coefficients themselves are estimates of the total effect a user has on all her friends.

4.8 How rain affects friends in other cities

Here we compute the cumulative effect that rain in one city has on all friends of users in that city who are in *different* cities. This allows us to answer the question: if it rains in New York, how many additional users in other cities post negative status updates as a result?

Consider the 2SLS model given by equations (4.9) and (4.10) for day t and city

g ,

$$\begin{aligned}\bar{Y}_g(t) &= \boldsymbol{\theta}'(t) + \bar{f}'_g + \beta_1 \bar{X}_g(t) + \boldsymbol{\varepsilon}'_g(t), \\ \bar{y}_g(t) &= \boldsymbol{\theta}(t) + \bar{f}_g + \gamma \bar{Y}_g^{pred}(t) + \bar{\boldsymbol{\varepsilon}}_g(t).\end{aligned}$$

Suppose that other cities are indexed by $h \neq g$ and let $\bar{X}_g^{(h,1)}(t)$, $\bar{Y}_g^{(h,1)}(t)$, $\bar{y}_g^{(h,1)}(t)$ respectively denote $\bar{X}_g(t)$, $\bar{Y}_g(t)$, $\bar{y}_g(t)$ conditional on $\bar{x}_h(t) = 1$ (that is, rainfall in city h). Similarly, let $\bar{X}_g^{(h,0)}(t)$, $\bar{Y}_g^{(h,0)}(t)$, $\bar{y}_g^{(h,0)}(t)$ be the same quantities conditional on $\bar{x}_h(t) = 0$. Using this notation, we can derive the following relationships:

$$\begin{aligned}\bar{X}_g^{(h,1)}(t) - \bar{X}_g^{(h,0)}(t) &= \sum_{i \in \mathcal{S}_h} \frac{1}{n_g} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t), \\ \bar{Y}_g^{(h,1)}(t) - \bar{Y}_g^{(h,0)}(t) &= \beta_1 \left(\bar{X}_g^{(h,1)}(t) - \bar{X}_g^{(h,0)}(t) \right) = \beta_1 \sum_{i \in \mathcal{S}_h} \frac{1}{n_g} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t), \\ \bar{y}_g^{(h,1)}(t) - \bar{y}_g^{(h,0)}(t) &= \gamma \beta_1 \left(\bar{X}_g^{(h,1)}(t) - \bar{X}_g^{(h,0)}(t) \right) = \gamma \beta_1 \sum_{i \in \mathcal{S}_h} \frac{1}{n_g} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t).\end{aligned}$$

Observe that $\bar{y}_g^{(h,1)}(t) - \bar{y}_g^{(h,0)}(t)$ is the difference in emotion of the average user in city g conditional on $\bar{x}_h(t) = 1$ versus $\bar{x}_h(t) = 0$. Assuming that each user posts either one or zero status updates on day t , $n_g(\bar{y}_g^{(h,1)}(t) - \bar{y}_g^{(h,0)}(t))$ is the additional number of negative status updates posted in city g conditional on $\bar{x}_h(t) = 1$ versus $\bar{x}_h(t) = 0$, where n_g is the number of users in city.

Fix a day t , let $\bar{I}_h(t)$ be the cumulative number of negative status updates posted in all cities different than h conditional on $\bar{x}_h(t) = 1$ versus $\bar{x}_h(t) = 0$, that is the *indirect*

of rain in city h . This can be computed by summing the effect on each city $g \neq h$,

$$\begin{aligned}\bar{I}_h(t) &= \sum_{g \neq h} n_g \left(\bar{y}_g^{(h,1)}(t) - \bar{y}_g^{(h,0)}(t) \right) = \gamma \beta_1 \sum_{g \neq h} n_g \sum_{i \in \mathcal{S}_h} \frac{1}{n_g} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t) \\ &= \gamma \beta_1 \sum_{i \in \mathcal{S}_h} \sum_{g \neq h} \sum_{j \in \mathcal{S}_g} \frac{1}{\delta_j(t)} a_{ij}(t) = \gamma \beta_1 \sum_{i \in \mathcal{S}_h} \sum_{j \notin \mathcal{S}_h} \frac{1}{\delta_j(t)} a_{ij}(t).\end{aligned}$$

For a user i in city h , $\sum_{j \notin \mathcal{S}_h} a_{ij}(t) / \delta_j(t)$ is the sum of the inverse degrees of i 's friend who are in a different city, and represents a measure of i 's *influence* outside city h . The indirect effect or rain $\bar{I}_h(t)$ is therefore proportional to the total influence from users in city h to their friends in other cities.

For each city h , we let \bar{I}_h be the average of $\bar{I}_h(t)$ over all days t .

The confidence interval of \bar{I}_h is computed from the confidence interval for the product $\gamma \beta_1$, as the other terms can be exactly computed. To compute a confidence interval on the product $\gamma \beta_1$ we cannot simply multiply the estimates of γ and β_1 in Table 4.6 derived from the model in equation (4.4) using two-stage least-squares regression, because they might be correlated. We therefore use bootstrap sampling by independently generating 100 bootstrap samples of our data set. Each bootstrap sample is generated by first selecting 1180 days uniformly at random with replacement, and then selecting 100 cities uniformly at random with replacement for each of the 1180 selected days. For each bootstrap sample, we estimate the model in equation (4.4) using two-stage least-squares regression, and we compute the product between β_1 (from the first-stage regression) and γ (from the second-stage regression). We then compute the mean and 95% CI of the estimates of $\gamma \beta_1$ from all bootstrap samples.

Tables 4.23 and 4.24 show the *direct* effect \bar{D}_h (with 95% CI) of rain in city h on status updates posted by users in that city, computed by multiplying the number of users n_h by the coefficient β_1 from the first stage in the 2SLS regression. The tables also show

the *indirect* effect \bar{I}_h (with 95% CI) for each city h of rain on status updates posted by users in other cities.

Similar results can be obtained by considering the effect on the number of either positive or negative posts, and using positive or negative emotions as the variable $\bar{Y}_g(t)$ in the first stage regression. The size and direction of \bar{I}_h and \bar{D}_h would depend on the magnitude and sign of $\gamma\beta_1$ and β_1 respectively.

Figure 4.3 plots the direct and indirect effect on the cities considered in the dataset.




Figure 4.3. Total number of negative posts generated by a day of rainfall within a city (direct) and in other cities via contagion (indirect). Blue colors indicate higher indirect/direct effect ratio. Larger labels indicate larger population

Table 4.23. Indirect and direct effect on negative emotion of rain in a city (the 50 most populous cities out of 100). The direct effect of rain in a city g is the cumulative number of status updates posted by the users in city g , as an effect of rain versus the lack of rain in city g (keeping everything else fixed). The indirect effect of rain in a city g is the cumulative number of status updates posted by the friends of users in city g who do not live in city g , as an effect of rain versus the lack of rain in city g . The table shows estimates and 95% CI

Indirect and direct effect of rain in a city – estimates and 95% CI (Number of negative posts)							
City Code	Population (US Census 2010)	Indirect effect			Direct effect		
		\bar{I}_g	Lo95%	Hi95%	\bar{D}_g	Lo95%	Hi95%
NYC	8175133	712.14	626.23	806.32	1550.42	1023.81	2071.68
LAX	3792621	666.96	586.5	755.16	1136.34	750.37	1518.38
CHI	2695598	573.32	504.16	649.14	1637.96	1081.62	2188.65
WAS	601723	494.37	434.73	559.75	878.38	580.04	1173.7
ATL	420003	477.59	419.98	540.75	992.6	655.46	1326.32
DAL	1197816	445.36	391.63	504.25	714.63	471.9	954.89
HOU	2100263	352.2	309.71	398.77	881.18	581.88	1177.44
SAN	1307402	299.02	262.95	338.56	536.04	353.97	716.26
AUS	790390	283.22	249.05	320.67	443.31	292.74	592.35
SFO	805235	278.62	245.01	315.47	419.23	276.84	560.18
ORL	238300	262.76	231.06	297.51	530.92	350.59	709.42
PHX	1445632	252.38	221.94	285.76	432.45	285.57	577.85
PHL	1526006	250.07	219.9	283.14	831.29	548.94	1110.77
BOS	617594	226.16	198.87	256.07	678.99	448.37	907.27
LAS	583756	215.53	189.53	244.03	367.55	242.71	491.12
TPA	335709	203.99	179.38	230.96	407.76	269.26	544.85
CLT	731424	202.33	177.92	229.09	441.28	291.4	589.65
BWI	620961	201.65	177.32	228.32	510.88	337.36	682.64
SEA	608660	199.65	175.57	226.05	450.34	297.38	601.75
MSP	382578	189.01	166.21	214	487.64	322.01	651.58
MIA	399457	184.1	161.89	208.44	511.4	337.7	683.34
BNA	601222	181.93	159.98	205.99	340.53	224.87	455.02
SAT	1327407	180.27	158.52	204.11	402.74	265.95	538.14
DEN	600158	174.49	153.44	197.57	342.77	226.35	458.01
DTT	713777	169.29	148.86	191.67	722.21	476.91	965.02
CMH	787033	165.41	145.45	187.28	385.78	254.75	515.48
VIB	437994	164.3	144.48	186.03	261.33	172.57	349.19
RAL	403892	155.61	136.84	176.19	331.11	218.64	442.43
PDX	583776	151.09	132.86	171.07	368.13	243.1	491.9
IND	820445	145.44	127.89	164.67	426.14	281.4	569.42
STL	319294	140.84	123.85	159.47	428.03	282.65	571.94
PIT	305704	138.34	121.65	156.63	423.86	279.89	566.36
JAX	821784	128.47	112.97	145.46	313.3	206.88	418.63
CVG	296943	125.75	110.58	142.38	349.55	230.82	467.07
MKE	594833	123.44	108.55	139.76	356.36	235.32	476.17
CLE	396815	121.64	106.97	137.73	290.14	191.59	387.69
MCI	459787	121.57	106.9	137.64	361.38	238.64	482.88
MSY	343829	119.4	104.99	135.18	217.86	143.86	291.1
DFW	741206	114.47	100.66	129.6	143.96	95.06	192.36
SMF	466488	109.26	96.08	123.7	259.26	171.2	346.42
MEM	646889	101.05	88.86	114.41	281.99	186.21	376.8
COS	416427	85.02	74.76	96.26	164.61	108.7	219.95
TUS	520116	85	74.74	96.24	183.67	121.29	245.42
OKC	579999	82.09	72.19	92.95	197.24	130.25	263.55
BHM	212237	80.15	70.48	90.74	206.74	136.52	276.25
ROC	210565	79.8	70.18	90.36	231	152.54	308.66
GSP	269666	79.29	69.72	89.77	148.43	98.02	198.34
BUF	261310	78.59	69.11	88.99	240.16	158.59	320.91
OMA	408958	77.38	68.04	87.61	198.7	131.21	265.51
HNL	337256	76.94	67.66	87.11	146.49	96.73	195.74

Table 4.24. Indirect and direct effect on negative emotion of rain in a city (the 50 least populous cities out of 100). The direct effect of rain in a city g is the cumulative number of status updates posted by the users in city g , as an effect of rain versus the lack of rain in city g (keeping everything else fixed). The indirect effect of rain in a city g is the cumulative number of status updates posted by the friends of users in city g who do not live in city g , as an effect of rain versus the lack of rain in city g . The table shows estimates and 95% CI

Indirect and direct effect of rain in a city – estimates and 95% CI (Number of negative posts)							
City Code	Population (US Census 2010)	Indirect effect			Direct effect		
		\bar{I}_g	Lo95%	Hi95%	\bar{D}_g	Lo95%	Hi95%
MSN	233209	75.29	66.21	85.25	157.48	103.99	210.42
TUL	391906	73.27	64.43	82.96	182.92	120.79	244.42
LGB	462257	70.16	61.7	79.44	82.38	54.4	110.08
BTR	229493	69.06	60.73	78.19	144.26	95.26	192.76
GKY	365438	68.67	60.38	77.75	77.07	50.89	102.98
ORF	242803	66.44	58.43	75.23	69.23	45.72	92.51
SDL	217385	64.1	56.37	72.58	56.15	37.08	75.03
SIC	945942	63.62	55.95	72.04	124.93	82.5	166.93
LEX	295803	62.26	54.75	70.49	166.71	110.08	222.76
MES	439041	58.33	51.29	66.04	70.82	46.76	94.63
SNP	285068	58.13	51.12	65.82	56.62	37.39	75.65
AWO	207627	58.11	51.1	65.79	48.02	31.71	64.16
CAK	199110	54.97	48.34	62.24	127.26	84.03	170.04
ICT	382368	53.72	47.24	60.82	172.3	113.78	230.23
OAK	390724	52.66	46.31	59.62	63	41.6	84.17
PLA	259841	52.58	46.23	59.53	51.25	33.84	68.47
TOL	287208	52.5	46.17	59.44	139.47	92.1	186.36
ABQ	545852	51.48	45.27	58.29	124.55	82.25	166.42
LNK	258379	48.66	42.79	55.09	113.83	75.17	152.11
LBB	229573	45.6	40.1	51.63	83.81	55.34	111.98
CPK	222209	44.53	39.16	50.42	48.25	31.86	64.47
CHD	236123	42.98	37.8	48.67	44.24	29.21	59.11
FAT	494665	40.7	35.79	46.08	114.99	75.94	153.66
JCY	247597	40.27	35.41	45.6	79.44	52.46	106.15
FWA	253691	39.61	34.83	44.85	112.57	74.33	150.42
ELP	649121	38.59	33.94	43.69	89.57	59.15	119.68
MGM	205764	38.23	33.62	43.29	83.39	55.07	111.43
CRP	305215	36.17	31.81	40.95	72.06	47.58	96.28
BOI	205671	34.14	30.02	38.66	98.93	65.33	132.19
ANC	291826	33.83	29.75	38.3	82.87	54.72	110.73
PIE	244769	33.8	29.72	38.27	36.59	24.16	48.89
HND	257729	33.54	29.49	37.97	32.71	21.6	43.71
RNO	225221	30.89	27.16	34.97	67.57	44.62	90.29
BFL	347483	27.67	24.33	31.33	84.18	55.59	112.49
RIV	303871	25.28	22.23	28.62	48.15	31.79	64.33
GAR	226876	23.98	21.08	27.15	26.58	17.55	35.52
SCK	291707	19	16.71	21.51	40.11	26.49	53.6
HTS	189992	18.27	16.07	20.69	49	32.36	65.48
ANA	336265	18.18	15.99	20.59	30.24	19.97	40.41
CHU	243916	17.9	15.74	20.27	18.21	12.03	24.34
MOD	201165	15.42	13.56	17.46	37.16	24.54	49.65
GEU	226721	14.99	13.18	16.97	13.86	9.15	18.52
EWR	277140	13.44	11.81	15.21	27.85	18.39	37.21
HIA	224669	11.27	9.91	12.76	12.02	7.94	16.06
LRD	236091	9.46	8.31	10.71	27.13	17.92	36.25
SNA	324528	8.65	7.61	9.8	17.44	11.52	23.31
SBD	209924	7.51	6.6	8.5	15.98	10.55	21.36
NHE	226322	5.35	4.71	6.06	6.01	3.97	8.03
RDU	228330	1.22	1.07	1.38	28.31	18.69	37.82
AUR	325078	0.31	0.27	0.35	3.87	2.55	5.17

4.9 Discussion

We proposed a rigorous method based on mathematical modeling and instrumental variable regression to detect and quantify contagion of semantic expression in online social networks using observational data. First, our method allows us to determine what semantic categories are susceptible to peer influence between social contacts. In particular, we showed that a person's post expressing positive or negative emotion can cause his or her friends to generate one to two additional posts expressing the same emotion. Second, it allows us to estimate a signed relationship between different categories, characterizing how an increase in the usage of a semantic category by an individual alters the usage of another by her social contacts. Third, our model allows us to quantify the cumulative effect that a person has on all her social contacts.

One potential concern is the instrument's weakness [198]—rainfall has only a small effect in our analysis, but this does not harm the validity of our conclusions because it is the precision, and not the size of the estimate that matters. In the dataset we used, built from content posted by millions of users, even a small effect is statistically significant and robust to a multitude of statistical tests against instrument weakness.

Our method limits inference to influence *between* subpopulations (individuals in different cities). Drawing conclusions about influence *within* a subpopulation (individuals in the same city) using observational data requires either the identification of a valid instrument or the definition of a different approach. This is an avenue of future research.

There are, of course, some limitations in inferring causality from observational data, and robust instruments may not always be available. Our model provides an alternative method when a large scale experiment is infeasible and researchers must rely on observational data. In an experiment, one would directly control the state of some people in order to track changes in their friends' outcomes (semantic expression, in

our case). With the proposed approach, which constitutes a “natural experiment”, the instrument (rainfall, in our case) constitutes a source of variation that affects some people directly (those experiencing it) but can predict changes in their social contacts who do not directly experience it. Moreover, our method can be easily applied to massive datasets (thanks to aggregation), and allows us to perform multiple analyses regarding several outcomes.

We advocate for the involvement of the engineering community in the development of non-experimental methods of causal inference. On the one hand, it is an open question how methods based on instrumental variable regression generalize to different contexts (especially contagion within a population) and how to build instruments in a systematic way. On the other hand, although instrumental variables might provide interesting answers, researchers should also develop and propose alternative techniques.

4.10 Acknowledgments

Chapter 4, in full, is a reprint of the material as it appears in the Proceedings of the IEEE, Vol. 102, Issue 12, pp. 1911–1921, 2014, L. Coviello, J. Fowler, and M. Franceschetti, and in PLoS ONE, 9(3), e90315, 2014, L. Coviello, Y. Sohn, A. Kramer, C. Marlow, M. Franceschetti, N. Christakis, and J. Fowler. The dissertation author was the primary investigator and co-author of these papers.

Chapter 5

Approximating physical encounter with friendship to predict epidemic outbreaks

5.1 Introduction

The forecast and mitigation of epidemics is a central theme in public health [104, 55, 92, 73, 146, 71, 72]. Events such as the recent ebola epidemic constantly drive the attention and resources of governments, institutions such as the World Health Organization, and the research community [86, 93, 178, 136, 145, 154]. The study of infectious processes on real-world networks is of interests to diverse disciplines, and similar models have been proposed to characterize the spread of information, behaviors, cultural norms, innovation, as well as the diffusion of computer viruses [85, 180, 209, 143, 170]. Therefore, epidemiologists, computer scientists and social scientist have joint forces in the study of contagion phenomena. Due to the impossibility to study the spread of infectious diseases through controlled experiments, modeling efforts have prevailed [99, 170, 132, 125, 169]. Recently, advancements in computation tools determined the emergence of data-driven simulations in the study of epidemic outbreaks and dynamical processes in general [210].

Physical encounter is the preferred vehicle for the spread of infectious deceases,

and occurs when two individuals meet at the same location at the same time. Therefore, detailed information about said encounters is fundamental to monitor and contain outbreaks. Different sources of data can serve as proxies of physical encounter – checkins on social networking platforms [49, 164, 163], traffic records [23, 203, 202], phone call records [166, 87, 101], wifi and RFID wearable sensors data [189, 42, 199], geographical and non-geographical information shared online [19, 48].

Human mobility and encounter present high spatial and temporal regularity and predictability [39, 87, 196, 203]. However, information about physical encounter, or proxies of it, is not always available, or it might be expensive and unpractical to collect (as in the case of sensor technologies [189, 42, 199]), prone to errors (as in the case of survey data [181, 65]), and privacy-sensitive [139, 28, 27, 60]. Therefore, the researcher needs to rely on the specific information in her possession, which might come in the form of self-reported social relationships between the individuals in a population (e.g., family, professional, friendship ties), or explicit ties formed by the users of a social network (e.g. Facebook friendship, follower-followee relationships on Twitter). Both real-world and online social relationships have been shown to predict the diffusion of behaviors and other phenomena [51, 52, 43, 13, 14]. Moreover, location data from cell phone records and online social networks has shown that social relationships can partially explain the patterns of human mobility [50].

The spread of an infection over a real-world network is composed by two processes. On the one hand, the dynamics of the network, whose structure changes over time according to the contacts between individuals. On the other hand, the dynamics on the network, whose state is represented by the set of infected nodes. When these dynamics operate at comparable time scale, the time-varying nature of the network cannot be ignored [90, 192], aggregating the dynamics of the edges into a static version of the network introduces bias [102, 174], and specifically devised control strategies are

necessary [142].

Given an infection transmitted by physical encounter, the present work is an attempt to quantify how accurately and within which limits said process can be predicted and approximated if the researcher has only access to explicit relationship ties between individuals (i.e., friendship). We consider a dataset from Yelp (www.yelp.com), a popular crowd-sourced online review service, in which users write reviews to restaurants, bars and other types of business. We use the time in which reviews to a business are posted as a proxy of physical encounter between users. That is, we assume that two individuals encountered on a given day if they both wrote a review to the same restaurant on that day. Users of Yelp can also form friendship ties between each other.

Physical encounter defines a *time-varying* network between users, over which an infection can spread over time. Such network has a fixed set of users, and the link between two users is activated when they have a contact. As a comparison, we define a time-varying network based on friendship, which has a fixed set of links between users (i.e., friendship ties) and where a user is activated whenever she writes a review. On such network, the infection can spread from a user to a friend only if they both write a review, even to different businesses, on the same day.

How well can an infection on the *friendship* network approximate an infection on the *encounter* network? We consider a susceptible-exposed-infected process (SEI process [8]), in which susceptible individuals can become infected when they are exposed to infected individuals, and individuals never recover from infection. To answer our question, we compare how the process spreads on the two networks both at a macroscopic and a microscopic level.

At the *macroscopic* level, we observe how the infection spreads over time on the two different networks. Two are the quantities of interests, the final size of the population infected over a given period to time, and the time needed to infect a fixed

fraction of the entire population. Differences between the friendship and encounter networks arise for their different connectivity, firstly because friendship determines a denser structure than physical encounter. In particular, friendship ties predict faster spread of the infection, larger infected population and earlier detection than physical encounter. Given an alarm threshold in terms of the infected fraction of the population, friendship predicts significantly more infections above the threshold than physical encounter. Despite these quantitative differences, the qualitative observations are similar for the two networks. Friendship appears a valid approximation to predict infections driven by physical encounter, if anything, with earlier detection and more likely alarm.

At the *microscopic* level, we focus on the set of nodes that are infected during an infection process. We consider seeds that are present in both the friendship and encounter networks and, on each network, we consider independent runs of the infection process starting at a given seed. We observe that, on average, the sets of nodes infected on the two networks show little overlap, suggesting that the two types of ties are substantially different. In addition, the stochasticity of the infection introduces larger unpredictability between independent runs on different networks than on the same network. Our results suggest that, given a set of initially infected seeds, friendship does not provide accurate prediction of the individuals at risk if the infection is driven by physical encounter.

Since seminal work on the structure and growth of complex networks [215, 26, 69], interdisciplinary research has shown that biological networks, social networks and the Internet are governed by similar rules [5, 110, 161, 32], and share similar structure [81, 162, 168]. Local connectivity and notions of centrality [76, 76, 41] reflect into nodes having different structural [6, 170, 54] and influential [127, 24, 25] importance. In this light, a contrast between global and local properties of complex networks emerges in the present paper. On the one hand, the global structural properties shared by two networks based on friendship and physical encounter result in similar infection dynamics

at a macroscopic level. On the other hand, differences in local connectivity determine striking differences in the dynamics at the microscopic level, making the similarity only superficial.

5.1.1 Outline

Section 5.2 describes the dataset and defines the notion of physical encounter. Section 5.3 formally defines the time-varying and static networks determined by friendship and encounter. Section 5.4 introduces the SEI process, defines the metrics to measure its spread, and describes the sensor selection mechanisms. Infection processes on time-varying networks are analyzed in Section 5.5 and Section 5.6. First, Section 5.5, takes macroscopic look at the process, focusing on the size of the infected population and the time to infect a desired fraction of the population (Section 5.5.1), and on the detection via sensors (Section 5.5.2). Then, Section 5.6, considers the process from a microscopic perspective, analyzing the sets of nodes infected on the two different networks, starting from the same seed. Infection processes on static networks are analyzed in Section 5.7 and Section 5.8, which, following an approach similar to that for time-varying network, take a macroscopic and a microscopic look at the process, respectively.

5.2 Dataset

We consider the Yelp Dataset Challenge dataset¹, which contains 1,569,264 reviews and 495,107 tips to 61,184 businesses (in 10 cities around the world) posted by 366,715 users over a period spanning over than 10 years. Each review is accompanied by rich metadata including the user who wrote it, the reviewed business, the day in which the review was written. Similarly, each tip has metadata including the user who wrote tip it, the date and the target business. Other than writing reviewed to businesses, Yelp users

¹www.yelp.com/dataset_challenge

can also form friendship ties between each other, in the same way as they would do in other social networking platforms. Each users in the dataset is accompanied by metadata listing his Yelp friends. Time information about the formation of friendship ties is not available.

The dataset allows us to build networks representing for friendship and physical encounter between users. We use such networks to simulate the spread of infection processes.

When friendship is the driver of a spreading phenomenon (such as a new opinion), then such phenomenon spreads over a network dictated by friendship ties. Friendship defines *friendships network*. Of all users, 174,100 have at least one friend, with an average number of friends per user, or friend degree, of 7.03 (14.8 restricted to non-singletons). The friend degree distribution presents a power law shape with cutoff at degree of about 1000 (Figure 5.1, circles).

Other phenomena, such as viral outbreaks, require physical contacts between individuals (rather than friendship) in order to spread. Therefore, we are interested in physical encounters between users. Strictly speaking, two users encountered on a given day if they visit to the same business on that day and at the same time. However, the dataset does not allow to precisely determine if two users encountered, as exact data about users checkins in business is not available. We use reviews as a proxy of physical encounter. In particular, we assume that two users encountered on a given day if they wrote a review to the same business on that day. This constitutes an approximation to real physical encounter, which requires users to *visit* (rather then review) a business at about the same time. We justify our approximation of physical encounter under the assumption that the date of a review is a proxy of the date of the visit to the reviewed business. Moreover, the element which spreads over a network (e.g., a virus or an opinion) does not necessarily require strict and timely physical contact between two particular individuals.




Figure 5.1. Inverse Cumulative Distribution Function of friends degree and encounter degree. Left: all non-singleton users. Center: users in the giant component of the network defined by all encounters. Right: users in the giant component of the network defined by all friendship ties.

For example, in the context of our dataset, after an infected user visits a business, the virus can infect customers which are not included in the dataset, and from them can infect another user who visits the business in a later moment.

Encounters defines *encounter networks*. In the dataset, 143,780 users have at least one encounter (who reviewed the same business as the user on the same day as the user), with an average number of encounters per user, or encounter degree, of 1.51 (3.9 restricted to non-singletons). The distribution of the encounter degree has a power law shape with cutoff at degree of about 100 (Figure 5.1, triangles).

Figure 5.2 shows a heat map of friend degree and encounter degree of users. Despite friend degree and encounter degree are correlated (Pearson product-moment correlation 0.3416, $p\text{-value} < 2.2 \cdot 10^{-16}$), the sets of the friends and of the encounters of an individual significantly differ on average. Figure 5.3 shows the cumulative distribution function of the Jaccard similarity of the set of friends and the set of encounters of all users in the dataset (left panel), of all users in the giant component of the network defined by all friendship ties (center panel), and of all users in the giant component of the network defined by all encounters (right panel). Considering the 72,786 users with at least one




Figure 5.2. Heat map of friend degree and encounter degree of all users with at least one friend and one encounter (friend degree and encounter degree are limited to 200 in the plot).

friend and one encounter, the average Jaccard similarity of their encounter and friend set is 0.01716, with only 9,527 of them with a value different than zero. Looking at the giant component of the network defined by all friendship ties, the users with nonzero encounters have average Jaccard similarity of their encounter and friend set of 0.1306, with only 9,022 users with a nonzero value. Looking at the giant component of the network defined by all encounters, the users with nonzero friends have average Jaccard similarity of their encounter and friend set of 0.112, with only 8,278 users with a nonzero value. In general, the sets of encounters and of friends of a user can significantly different and often have empty intersection. Therefore, as we will see, although processes driven by friendship and physical encounter evolve in a qualitatively similar way, the sets of infected nodes can be very dissimilar even if the seed of the infection is the same in the two cases.




Figure 5.3. Percentile plot of the Jaccard similarity of the set all user’s friends and the set of all user’s encounters. Left: all non-singleton users. Center: users in the giant component of the network defined by all encounters. Right: users in the giant component of the network defined by all friendship ties.

5.3 Static and time-varying networks

Let U be the set of users, $F \subseteq U \times U$ be the set of friendship ties, B the set of businesses, T be the set of days, $R \subseteq U \times B \times T$ be the set of reviews and tips (which we will refer to as reviews). For each user $u \in U$ let $F_u \subset U$ be the set of friends of u . Therefore $F = \cup_{u \in U} \{(u, v) : v \in F_u\}$. Each review (or tip) $r \in R$ is a triple (u, b, t) where $u \in U, b \in B, t \in T$.

For each $t \in T$, let $U(t) = \{u \in U : (u, b, t) \in R \text{ for some } b \in B\}$, that is, the set of users who wrote a review on day t . We refer to $U(t)$ as the active users on day t . For each $t \in T$, let $F(t) = \{(u, v) \in F : u, v \in U(t)\}$, that is, the set of friendship ties between users in $U(t)$. Observe that friendship ties are not associated to temporal information (i.e., the time when the edge formed is unknown).

For each $t \in T$ and $u \in U(t)$, let

$$E_u(t) = \{v \in U(t), v \neq u : (u, b, t) \in R \text{ and } (v, b, t) \in R \text{ for some } b \in B\} \subseteq U,$$

that is, the encounters of user u on day t . Let $E_u = \cup_{t \in T} E_u(t) \subseteq U$ be the set of encounters

of u , $E(t) = \cup_{u \in U} E_u(t) \subseteq U \times U$ be the set of encounters on day t , and $E = \cup_{t \in T} E(t) \subseteq U \times U$ be the set of encounters.

We define static and time-varying networks of friends and encounters. Static networks have a set of edges that is independent of time, whereas the edges of a time-varying network are activated at specific times. Let $N_F = (U, F)$ be the static friendship network. Let $N_E = (U, E)$ be the static encounter network. That is, the static networks are defined by the set of all friendship ties and all encounter ties, respectively. For each $t \in T$, let $N_F(t) = (U, F(t))$ be the friendship network of day t , defined by the friendship ties between the active users on day t (observe that the set of nodes is U rather than $U(t)$). For each $t \in T$, let $N_E(t) = (U, E(t))$ be the encounter network of day t , defined by encounters on day t . The time-varying networks defined sequences $\{N_F(t)\}_{t \in T}$ and $\{N_E(t)\}_{t \in T}$, called the friendship time-varying network and the encounters time-varying network.

As we consider processes spreading between connected nodes, connectedness is the key property of the networks. For static networks, we restrict our attention to giant components. Users outside giant components form small components whose dynamics are not relevant. The giant component of the static friendship network N_F includes 168,923 users (whereas the second largest component has 8 users). The giant component of the static encounter network N_E includes 113,187 users (whereas the second largest component has 23 users). For time-varying networks, we restrict our attention to a set T of 1,469 consecutive days ranging from January 1st, 2011 to January 8th, 2015 included, as the number of daily reviews before year 2011 is smaller. The encounter time-varying network contains the 133,038 users who had at least one encounter during T . The friendship time-varying network contains the 41,664 users who reviewed a business during T and also had a friend who reviewed a business during T .

Different networks might be relevant for the spread of different phenomena. As an

example, if a spreading process is driven by friendship (and the time in which friendship ties formed is irrelevant) then the static network N_F is the most natural network model. On the other hand, encounters are inherently associated to time, and if a spreading process is driven by physical encounters then the time-varying network $\{N_E(t)\}_{t \in T}$ is the most natural network model. In addition, different information, and therefore different networks, might be available to the researcher. For example, in specific situations, friendship ties might be the only available information, and the researcher would have to rely on those. In other situation, friendship information might be unavailable and the researcher might have to rely on checkin information. Therefore, it is our interest to compare all the networks defined above. However, static and time-varying networks (say, N_E and $\{N_E(t)\}_{t \in T}$) are not directly comparable, due to the different processes through which they are defined. In what follows, we will either compare the static networks N_F and N_E , or the time-varying networks $\{N_F(t)\}_{t \in T}$ and $\{N_E(t)\}_{t \in T}$.

5.4 Infection dynamics

Given a set of nodes \mathcal{V} , a set of edges $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ and a set of time indices \mathcal{T} , let $\{N(t)\}_{t \in \mathcal{T}}$ be a sequence of networks, where $N(t) = (\mathcal{V}, \mathcal{E}(t))$ with $\mathcal{E}(t) \subseteq \mathcal{E}$. For static networks, $\mathcal{E}(t) = \mathcal{E}$ for all of $t \in \mathcal{T}$.

We consider a Susceptible-Infected (SI) model, in which nodes never recover after being infected.

Let $\mathcal{I}(t)$ denote the set of infected nodes at time t , of cardinality $I(t)$. The infection starts at time $t = 0$ from a set $\mathcal{I}(0)$ of infected seeds.

Consider any $t > 0$. The infection spreads from the set of already infected nodes $\mathcal{I}(t-1)$ as follows. For each non-infected node $v \in \mathcal{V} \setminus \mathcal{I}(t-1)$, let $d_v(t) = |\{u \in \mathcal{I}(t-1) : (u, v) \in \mathcal{E}(t)\}|$, that is, the number of neighbors of v at time t which are infected at time $t-1$. Let $B(t) = \{v \in \mathcal{V} \setminus \mathcal{I}(t-1) : d_v(t) > 0\}$, that is, the set of

susceptible nodes at time t . Each node $v \in B(t)$ gets infected with probability $\beta d_v(t)$, where $\beta \in [0, 1]$ is the rate of infection.

For $\beta \rightarrow 1$, the infection process is deterministic in the sense that all non-infected neighbors at time t of the infected nodes at time $t - 1$ become infected at time t . For finite values of β , the infection spreads in a stochastic way.

For the time-varying networks defined above, we take $\mathcal{T} = T$, that is, the set of days in the date set. The infection will propagate for $|T|$ time steps, resulting in an infected population $\mathcal{I}(|T|)$. For static networks, we take $\mathcal{T} = [0, \infty)$ and let the infection propagate until $\mathcal{I}(t) = \mathcal{V}$ (i.e., until the entire population is infected).

5.4.1 Infection time

We consider a sequence of networks $\{N(t)\}_{t \in \mathcal{T}}$, where $N(t) = (\mathcal{V}, \mathcal{E}(t))$, and the SEI process with rate β starting from a seed $\mathcal{I}(0)$. For each $\alpha \in [0, 1]$ let

$$\tau(\alpha) = \min\{t : I(t)/|\mathcal{V}| \geq \alpha\}.$$

$\tau(\alpha)$ is a random variable and represents the first time in which an α -fraction of \mathcal{V} are infected (once $\mathcal{I}(0)$ is fixed, $\tau(\alpha)$ is a degenerate random variable for $\beta = 1$). For time-varying networks, $\tau(\alpha) = \infty$ for $\alpha > \mathcal{I}(|T|)/|\mathcal{V}|$.

We will also consider the number, rather than the fraction, of infected nodes. For each $M \in [0, |\mathcal{V}|]$, let

$$t(M) = \min\{t : I(t) \geq M\},$$

The random variable $t(M)$ denotes the first time in which at least M nodes are infected. For time-varying networks, Let $t(M) = \infty$ for $M > \mathcal{I}(|T|)$.

5.4.2 Seed selection

In a static network, seeds are chosen at random and without replacement. In a time-varying network, the infection can start propagating at the first time t in which there is an edge between an infected seed and a non-infected node, that is, at time

$$t_0(\mathcal{I}(0)) = \min\{t : \exists(u, v) \in \mathcal{E}(t) \text{ for some } u \in \mathcal{I}(0), v \in \mathcal{V} \setminus \mathcal{I}(0)\}.$$

As a remark, for $\beta < 1$, it is possible that no node is infected at time t_0 . Seeds are selected uniformly at random and without replacement among all nodes $v \in \mathcal{V}$ such that $t_0(\{v\}) \leq 500$, that is, nodes that have a neighbor in the time-varying network by time $t = 500$.

5.4.3 Detection time with sensors

In real scenarios, it might be infeasible to monitor all nodes in the network. Constraints of different nature (e.g., budget, physical, privacy) might limit the researchers to monitor a subset $S \subset \mathcal{V}$ of all nodes, referred to as sensors. At each time t , let $\mathcal{I}_S(t) = \mathcal{I}(t) \cap S$ be the set of infected sensors, and $I_S(t)$ be its cardinality. Assuming as before that the network and the set of seeds are given, for each $\alpha \in [0, 1]$ let

$$\tau_S(\alpha) = \min\{t : I_S(t)/|S| \geq \alpha\}.$$

That is, $\tau_S(\alpha)$ represents the first time in which an α -fraction of the sensors S are infected. For time-varying networks, Let $t(\alpha) = \infty$ for $\alpha > I_S(|T|)/|S|$.

5.4.4 Sensor selection

We consider two types of sensor selection, random sensors and friend sensors, defined as follows. Let m be a fixed parameter. A set S of random sensors is obtained

by selecting m nodes from \mathcal{V} uniformly at random and without replacement. A set S of friend sensors is obtained in two steps. First, S is initialized as the empty set, and a set S_0 of random nodes is obtained by selecting m users from \mathcal{V} uniformly at random and without replacement. Then, for each node $u \in S_0$, a *friend* $v \in \mathcal{V}$ is selected uniformly at random from F_u (i.e., from the set of friends of u) and added to S . We require each friend sensor to be in \mathcal{V} and to be friend of a node in S_0 . We remark that, even for encounter networks, friend sensors are selected on the basis of friendship rather than encounter. We make this assumption because explicit relationships (such as friendship, family or professional ties) might be accessible or inferable in a real setting in which the researcher has to select a set of sensors. Observe that, in both cases of random sensors and friend sensors, the size of the resulting set S might be smaller than m .

Given the fact that, on average, people have fewer friends than their friends have (also know and the friendship paradox [70]), randomly sampled friends are more connected than randomly sampled individuals and are shown to provide earlier detection of phenomena spreading over complex networks [53, 79].

5.5 Infection detection – Time-varying networks

In this section, we consider SEI processes on the two time-varying networks $\{N_F(t)\}_{t \in T}$ and $\{N_E(t)\}_{t \in T}$, started from a single seed, $\mathcal{I}(0) = \{s\}$. We will focus on $\beta = 1$ (i.e., infection is certain) in order to study how the structure of the two different networks affects the spreading process. We remark that, despite the set of nodes of the networks is potentially the same (i.e., all users in the dataset), the definitions of friendship and encounter result in different sets of nodes.

5.5.1 Single seed – Infection rate

With $\beta = 1$, we perform 10,000 simulations on each time-varying network. In each simulation, a single seed is selected uniformly at random between all nodes s such that $t_0(\{s\}) \leq 500$. That is, in the case of the friendship (respectively, encounter) network, we consider potential seeds that have an edge in $N_F(t)$ (respectively, $N_E(t)$) for some $t \leq 500$. As infections on time-varying networks spread for a limited number of time steps, we require them to start early enough.

Each simulation i is therefore associated to a seed s_i and, as $\beta = 1$, the first time in which a node other than s_i is infected is

$$t_0(s_i) = \min\{t : \exists (s_i, v) \in \mathcal{E}(t) \text{ for some } v \neq s_i\} \in [1, 500],$$

We refer to $t_0(s_i)$ as the starting time of the infection. Let $t_F(s_i)$ be the last time in which a node is infected in an infection starting from s_i (i.e., the time after which the size of the infected population stops increasing). It holds that $t_F(s_i) \leq \max T$. Moreover, $t_F(s_i)$ depends on the sequence of the edge sets in the time-varying network. At time $t_F(s_i)$, the infection reaches its peak, infecting a fraction $r(s_i) \in [0, 1]$ of the population.

The final infection $r(s_i)$ decreases with increasing infection starting time $t_0(s_i)$, for both the friendship network (OLS, coefficient $-4.255 \cdot 10^{-4}$, p-value $< 2 \cdot 10^{-16}$, intercept 0.451, p-value $< 2 \cdot 10^{-16}$) and the encounter network (OLS, coefficient $-3.922 \cdot 10^{-4}$, p-value $< 2 \cdot 10^{-16}$, intercept 0.757, p-value $< 2 \cdot 10^{-16}$). Instead, $t_0(s_i)$ does not predict $t_F(s_i)$ for either the friendship network (OLS, coefficient -0.03701 , p-value 0.376) or the encounter network (OLS, coefficient -0.03662 , p-value 0.381). This suggests that the networks remain connected over time and therefore infections that start earlier do not stop earlier.

Due to higher connectivity, the final rate of infection $r(s_i)$ is on average 31.5%

higher on the friendship network than on the encounter network (OLS, 0.3149, p-value < $2 \cdot 10^{-16}$, when controlling for $t_0(s_i)$), see Figure 5.4 (right panel). Also, the time $t_F(s_i)$ of maximum infection is reached on average 79 time steps later on the friendship network than on encounter network (OLS, 79.19, p-value < $2 \cdot 10^{-16}$, when controlling for $t_0(s_i)$), see Figure 5.4 (left panel).

The fraction of infected nodes increases linearly over time in both networks (see Figure 5.5). In particular, we consider all infections that infected at least 1% of the total population (7,888 out of 10,000 simulations in the encounter network, and 9,100 in the friendship network). The infection spreads faster in the friendship network (OLS, slope 0.06209, p-value < $2 \cdot 10^{-16}$) than in the encounter network (OLS, slope 0.03187, p-value < $2 \cdot 10^{-16}$), with a significantly different slope difference (OLS, interaction coefficient of $7.57 \cdot 10^{-3}$, p-value < $2 \cdot 10^{-16}$). Moreover, even if an infection starts at time $t \leq 500$, it still might take a while to infect a significant amount of the population (see Figure 5.5). There is, therefore, a period of “incubation” during which the fraction of the infected population remains very low.

5.5.2 Single seed – Sensor monitoring

Instead of monitoring the entire population, in each run of the SEI process, we consider a random set of sensors composed by 1% of the population. Sensors are selected in the two ways described above: random sensors and friend sensors (where the selection is based on friendship rather than encounter, even when considering a process spreading on the encounter network). We perform 10,000 simulations on each time-varying network and each sensor type, setting $\beta = 1$ (i.e., infection is certain). In each simulation, a single seed is selected uniformly at random between all nodes s_i such that $t_0(s_i) \leq 500$.

Let $r_S(s_i)$ denote the final infection of the sensors (which we have to use instead of $r(s_i)$ to monitor the infection). Also $r_S(s_i)$ linearly decreases with increasing infection

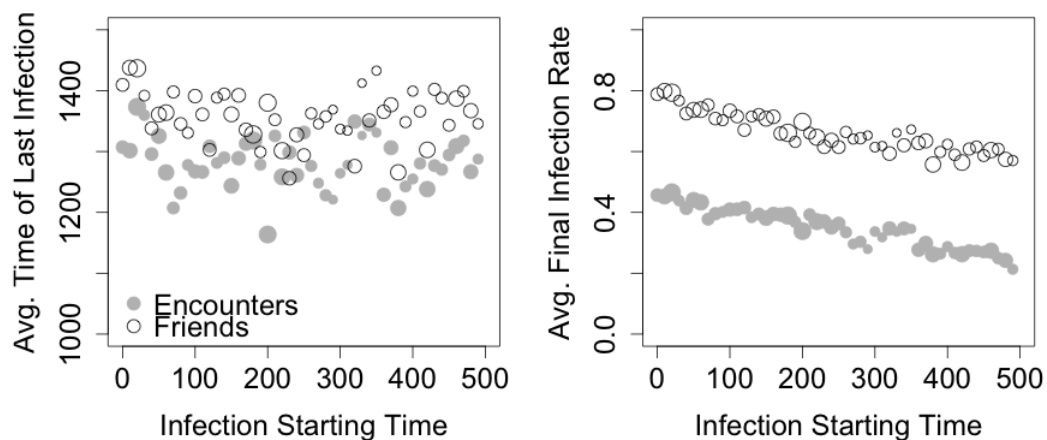


Figure 5.4. Susceptible-exposed-infected process on the friendship time-varying network (white circles) and encounter time-varying network (grey circles), $\beta = 1$ (certain infection). 10,000 simulations are run on each network, each with a single seed s_i selected at random among all nodes such that $t_0(s_i) \leq 500$. The x -axis represents the infection start time $t_0(s_i)$, rounded to the lower multiple of 10. Point size is proportional to the number of observations for the corresponding value of the x -axis. Left: Average of the last time of infection $t_F(s_i)$ (i.e., the time at which the peak of the infection is reached) with respect to $t_0(s_i)$, for both the friendship and encounter networks. Right: Average of the final infection $r(s_i)$ with respect to $t_0(s_i)$, for both the friendship and encounters networks.




Figure 5.5. Fraction of infected nodes over time, for the friendship time-varying network (left) and the encounter time-varying network (right). Each SEI process (with $\beta = 1$) is started from a single seed s_i selected at random among all nodes such that $t_0(s_i) \leq 500$. For each network, 60 simulations that infected at least 1% of the population are considered. Colors are not meaningful.

start time (Figure 5.6, left). On average, friend sensors predict an infection rate 9.5% higher than random sensors (OLS, coefficient 0.0953, $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time $t_0(s_i)$ and type of network). As random sensor constitute a random sample of the population, their infection reflects the infection of the entire population. Instead, friends sensors are more connected than average nodes (friend paradox) and therefore their larger infection constitutes an overestimation of the infection of the population. Such overestimation can be beneficial for early detection of an outbreak. The overestimation effect is larger on the encounter network (OLS, coefficient 0.1197, $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time) than on the friendship network (OLS, coefficient 0.0709, $p\text{-value} < 2 \cdot 10^{-16}$, when controlling for infection starting time). However, the sensor type does not significantly affect the slope of the observed linear decrease (OLS: interaction between infection starting time and sensor type, $2.299 \cdot 10^{-5}$, $p\text{-value} 0.21$). We also observe that, on the friendship network,

the $r_S(s_i)$ is on average 29% higher than on the encounter network (OLS, coefficient -0.2935 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time and type of network). This effect is larger for random sensors (OLS, coefficient 0.2809 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time) than for friend sensors (OLS, coefficient 0.2130 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time).

When restricting our attention to simulations which infected at least 10% of the sensors (on the encounter network, 7,669 with random sensors, 7,781 with friend sensors, on the friendship network, 9,140 with random sensors, 9,109 with friend sensors), on average, the 10% infection of friends sensors is reached 128 time units earlier than the 10% infection of random sensors (OLS, coefficient -128.0 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time and type of network). For the same consideration as above, friend sensors offer earlier detection with respect to the 10% infection of the entire population. This underestimation effect is larger on the encounter network (OLS, coefficient -197.3 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time) than on the friendship network (OLS, coefficient -69.2 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time). Also in this case, the sensor type does not affect the slope of the observed linear increase (OLS: interaction between infection starting time and sensor type, $-2.302 \cdot 10^{-3}$, $p\text{-value} 0.892$). We also observe that, on the friendship network, the infection of 10% of the sensors requires on average 302 units of time less than on the encounter network (OLS, coefficient -302.6 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time and type of sensors). This effect is larger for random sensors (OLS, coefficient -3.672 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time) than friend sensors (OLS, coefficient -2.384 , $p\text{-value} < 2 \cdot 10^{-16}$, controlling for infection starting time).

Figure 5.7 plots the fraction of simulations that reached a target sensors' infection versus the infection starting time (values of the target: 10%, 25%, 50%, 75%, 80%, 85%).

We refer to the infections that reached a given target as successful (for the given target). For targets of 10% and 20% (top plots) the observations are the same as above. For a target of 50% (middle left plot), on the encounter network (circles), the fraction of successful infection decreases more steeply for random sensors (grey) than friend sensors (white). with the former, the fraction of successful infections approaches zero for infection starting time above $t = 350$. This effect is not observed in the case of the friendship network (triangles) for target of 50%. For a target of 75% (middle right plot), we observe a similar effect also on the friendship network, on which the success rate decreases faster with random sensors (approaching zero for infection starting times above $t = 400$). On the encounter network, there is no successful infection of random sensors, whereas some successful infection of friends sensors happens for infection starting time before $t = 100$. For targets of 80% and 85% (bottom plots), the observations are similar.




Figure 5.6. Infection detection with random sensors and friend sensors on the friendship and encounter time-varying networks. 10,000 simulations, with $\beta = 1$, are run on each network and for each sensor type. Each simulation starts with a seed s_i selected at random among all nodes such that $t_0(s_i) \leq 500$. Sensor size is 1% of the population. The x -axis represents the infection start time $t_0(s_i)$, rounded to the lower multiple of 10. Point size proportional to the number of observations for the corresponding value of the x -axis. Left: average final sensor infection versus infection start time, for the friendship network (triangles) and the encounter network (circles), with random sensors (grey) and friend sensors (white). Right: average time to infect 10% of the sensors versus infection start time, considering only the simulations in which at least 10% of the sensors are infected.

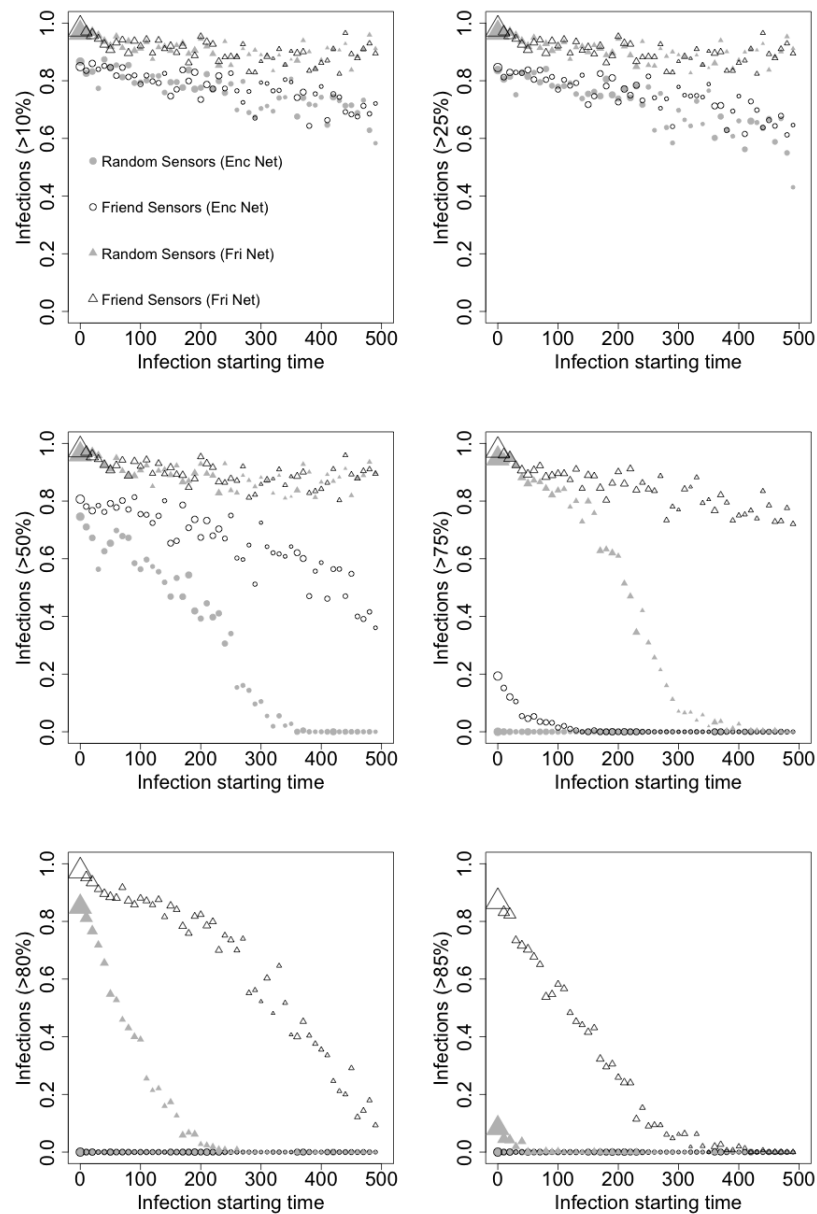


Figure 5.7. Fraction of simulations that reached a target sensors' infection versus the infection starting time, for different targets, for the encounter (circles) and friendship time-varying networks (triangles), using random sensors (grey) and friend sensors (white). 10,000 simulations, with $\beta = 1$, are run on each network and for each sensor type. Each simulation starts with a seed s_i selected at random among all nodes such that $t_0(s_i) \leq 500$. Sensor size is 1% of the population. The x -axis represents the infection start time $t_0(s_i)$, rounded to the lower multiple of 10. Point size proportional to the number of observations for the corresponding value of the x -axis.

5.6 The infected population – Time-varying networks

In the previous section, we observed how infection processes spread on the friendship and encounter time-varying networks, at a macroscopic level, by focusing on quantities such as the size of the final infection and the infection detection time. In this section, we take a microscopic look at the process, moving our attention on the sets of nodes that become infected. That is, we will consider seed nodes that are present in both the friendship and the encounter network, and we will compare the sets of nodes that become infected in processes starting at the same seed but evolving on the two different networks. First, we will consider the case of certain infection ($\beta = 1$), in which the dynamics are dictated by the network structure only. As we will observe, the structural differences between the two networks result in substantially different sets of infected nodes. Then, we will consider the case of stochastic infection ($\beta < 1$). Randomness introduces a certain amount of unpredictability in the spread of the infection, and two runs of the process on the same network starting from the same seed can result in different sets of infected nodes. However, we will observe that the unpredictability within a given network is substantially lower than the unpredictability between the two different networks.

Despite infections spread similarly on the encounter and friendship network at the macroscopic level (at least qualitatively), our results show that the same is not true at the microscopic level. That is, the individuals predicted to be at risk are different according to the two networks. This observation is particularly relevant if we consider that the experimenter has to rely on the social network information available to her, regardless of the network on which the phenomenon of interest spreads. Such information might come in the form of implicit ties formed by individuals (e.g., friendship, work relationship), or in the form of records of human behavior across time (e.g., reviews, checkins, sharing).

Similarly, different phenomena spread in different ways. On the one hand, physical encounter is the vehicle over which diseases spread. On the other hand, ideas, news and digital content might spread over the ties of a social network.

5.6.1 Certain infection ($\beta = 1$)

In order to study how the network structure affects the infection process at a microscopic level, we consider infections with $\beta = 1$ starting from seeds that are present in both the friendship and encounter network. Fixed a seed s_i , we consider infection processes that spread independently on the two networks. The infection from a seed s_i spreads deterministically. It begins at the first time step $t_0(s_i)$ in which s_i is connected to at least another node in the network, it continues up to time T , and the set of infected nodes is non-decreasing over time. Given a target set size m , we will compare the sets of the first m infected nodes in the encounter and the friendship network, respectively. As we consider a discrete process, there might be no time t at which exactly m nodes are infected. Therefore, we allow for some flexibility on the number of infected nodes, requiring that *at least* (rather than exactly) m nodes are infected. In particular, let $\mathcal{I}_E(t; s_i)$ and $\mathcal{I}_F(t; s_i)$ denote the set of infected nodes in the encounter and the friendship network at time t , respectively. Let $I_E(t; s_i)$ and $I_F(t; s_i)$ be their cardinality. Let

$$t_E(m; s_i) = \min\{t \in T : I_E(t; s_i) \geq m\}$$

$$t_F(m; s_i) = \min\{t \in T : I_F(t; s_i) \geq m\}$$

be the minimum time at which at least m nodes are infected on the encounter and friendship network, respectively. $t_E(m; s_i)$ or $t_F(m; s_i)$ are undefined if m nodes never get infected in the corresponding network. $t_E(m; s_i)$ and $t_F(m; s_i)$ can differ between each other and need not be both defined.

If $t_E(m; s_i)$ is defined, then the corresponding infected set is

$$\mathcal{I}_E^*(m; s_i) = \mathcal{I}_E(t_E(m; s_i)).$$

If $t_F(m)$ is defined, then the corresponding infected set is

$$\mathcal{I}_F^*(m; s_i) = \mathcal{I}_F(t_F(m; s_i)).$$

When they are both defined, $\mathcal{I}_E^*(m; s_i)$ and $\mathcal{I}_F^*(m; s_i)$ can be compared by means of their Jaccard similarity

$$J(m; s_i) = \frac{|\mathcal{I}_E^*(m; s_i) \cap \mathcal{I}_F^*(m; s_i)|}{|\mathcal{I}_E^*(m; s_i) \cup \mathcal{I}_F^*(m; s_i)|},$$

by means of the fraction of $\mathcal{I}_E^*(m; s_i)$ contained in $\mathcal{I}_F^*(m; s_i)$ (that is, the nodes infected in the encounter network that are also infected in the friendship network),

$$P_E(m; s_i) = \frac{|\mathcal{I}_E^*(m; s_i) \cap \mathcal{I}_F^*(m; s_i)|}{|\mathcal{I}_E^*(m; s_i)|},$$

and by means of the fraction of $\mathcal{I}_F^*(m; s_i)$ contained in $\mathcal{I}_E^*(m; s_i)$ (that is, the nodes infected in the friendship network that are also infected in the encounter network),

$$P_F(m; s_i) = \frac{|\mathcal{I}_E^*(m; s_i) \cap \mathcal{I}_F^*(m; s_i)|}{|\mathcal{I}_F^*(m; s_i)|}.$$

If either $\mathcal{I}_E^*(m; s_i)$ or $\mathcal{I}_F^*(m; s_i)$ is not defined, then the metrics above are not defined.

It holds that

$$\begin{aligned}
0 &\leq \frac{P_E(m; s_i) \cdot P_F(m; s_i)}{P_E(m; s_i) + P_F(m; s_i)} \\
&\leq J(m; s_i) \leq \max\{P_E(m; s_i), P_F(m; s_i)\} \\
&\leq 1.
\end{aligned}$$

Experiments. We ran 5000 pairs simulations of the SEI process. For each simulation pairs, a single seed is selected at random among all nodes s_i such that $t_0(s_i) \leq 500$ in both the encounter and the friendship time-varying network (that is, we consider nodes that have neighbors on both networks by time $t = 500$). For each choice of the seed, we separately run two infection processes: one on the encounter network and one on the friendship network. For target set size $m \in \{500, 1000, 2000, 5000, 10000, 20000\}$ and each of the 5000 seeds s_i , the metrics $J(m; s_i), P_E(m; s_i), P_F(m; s_i)$ are considered when $\mathcal{J}_E^*(m; s_i)$ and $\mathcal{J}_F^*(m; s_i)$ are defined.

Results are plot in Figure 5.8. Observations for a given value of m constitute a block on the x -axis (larger m corresponds to x positions on the right) and are represented with the same color. For a fixed value of m , relative x positions are irrelevant. The three panels represent the metrics defined above, $J(m; s_i), P_E(m; s_i)$ and $P_F(m; s_i)$ (which are between zero and one). For a given metric and m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations. Combined, the low values of the metrics highlight the structural differences between the two networks.

However, the intersections between the infected sets on the encounter and friendship networks are far from random. To compare the intersection of the infected sets with the intersection of random sets, we derive an upper bound for the expected size of the intersection of random sets, as follows. If we consider only the $n = 24,251$ nodes

which appear in both networks, then two random sets of size m_1 and m_2 have expected intersection size $m_1 m_2 / n$. If all nodes in each network can be selected, the expected size of the intersection of random sets is smaller, and the quantity above is an upper bound. With $m_1 = \mathcal{I}_E^*(m; s_i)$ and $m_2 = \mathcal{I}_F^*(m; s_i)$, and considering each value of the target m separately, t-tests reject hypothesis that the size of the intersections of the infected sets on the two networks has the same mean as the size of the intersection of random sets (p-values $< 2.2 \cdot 10^{-16}$).




Figure 5.8. The three panels show the metrics $J(m; s_i)$, $P_E(m; s_i)$, $P_F(m; s_i)$ for 5000 simulation pairs, each with a random choices of a single seeds, and different values of the target set size m . For each seed, simulations are run separately on the friendship and encounter network. On the x - axis, observations for a given value of m form a block with a constant color (within the block, the x position is irrelevant). We only consider pairs (m, s_i) for which the metrics are defined. For a given metric and each value m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations.

5.6.2 Stochastic infection ($\beta < 1$)

In the previous section, setting $\beta = 1$ (certain infection) allowed us to isolate the effect of the network structure from any randomness other than the selection of the seed. In this section, we consider $\beta \leq 1$ (stochastic infection) and show that the structural difference between friendship and encounter networks introduces additional unpredictability about the sets of infected nodes with respect to the unpredictability of the random infection. In particular, given a seed s_i that is present in both networks, for different target infection size m , we compare the sets of the first m infected nodes between runs of the infection process on the two different networks as well as within the same network.

Fixed a seed s_i , let $\mathcal{I}_{E_1}(t; s_i)$ and $\mathcal{I}_{E_2}(t; s_i)$ denote the set of infected nodes at time t in two separate infection processes on the encounter network starting at s_i . $\mathcal{I}_{F_1}(t; s_i)$ and $\mathcal{I}_{F_2}(t; s_i)$ are similarly defined by considering the friendship network. Let $I_{E_1}(t; s_i)$, $I_{E_2}(t; s_i)$, $I_{F_1}(t; s_i)$, $I_{F_2}(t; s_i)$ be their cardinality. For $j = 1, 2$, let

$$t_{E_j}(m; s_i) = \min\{t \in T : I_{E_j}(t; s_i) \geq m\}$$

$$t_{F_j}(m; s_i) = \min\{t \in T : I_{F_j}(t; s_i) \geq m\}$$

be the minimum time at which at least m nodes are infected in the corresponding process. $t_{E_j}(m; s_i)$ or $t_{F_j}(m; s_i)$ is undefined if m nodes never get infected in the corresponding process.

If $t_{E_j}(m; s_i)$ is defined, then the corresponding infected set is

$$\mathcal{I}_{E_j}^*(m; s_i) = \mathcal{I}_{E_j}(t_{E_j}(m; s_i)).$$

If $t_{F_j}(m)$ is defined, then the corresponding infected set is

$$\mathcal{I}_{F_j}^*(m; s_i) = \mathcal{I}_F(t_{F_j}(m; s_i)).$$

When all values $\mathcal{I}_{E_j}^*(m; s_i)$ and $\mathcal{I}_{F_k}^*(m; s_i)$ for $j, k \in \{1, 2\}$ are defined, we define the following measures of Jaccard similarity,

$$\begin{aligned} J_{E_j, F_k}(m; s_i) &= \frac{|\mathcal{I}_{E_j}^*(m; s_i) \cap \mathcal{I}_{F_k}^*(m; s_i)|}{|\mathcal{I}_{E_j}^*(m; s_i) \cup \mathcal{I}_{F_k}^*(m; s_i)|}, \\ J_{E_1, E_2}(m; s_i) &= \frac{|\mathcal{I}_{E_1}^*(m; s_i) \cap \mathcal{I}_{E_2}^*(m; s_i)|}{|\mathcal{I}_{E_1}^*(m; s_i) \cup \mathcal{I}_{E_2}^*(m; s_i)|}, \\ J_{F_1, F_2}(m; s_i) &= \frac{|\mathcal{I}_{F_1}^*(m; s_i) \cap \mathcal{I}_{F_2}^*(m; s_i)|}{|\mathcal{I}_{F_1}^*(m; s_i) \cup \mathcal{I}_{F_2}^*(m; s_i)|}. \end{aligned}$$

$J_{E_j, F_k}(m; s_i)$ is the similarity between the infected sets (for target m) in two infection processes on the different networks. In the analyses we will consider $J_{E_1, F_1}(m; s_i)$ only. $J_{E_1, E_2}(m; s_i)$ (resp., $J_{F_1, F_2}(m; s_i)$) is the similarity between the infected sets (for target m) in the two infection processes on the encounter (resp., friendship) network.

When all values $\mathcal{I}_{E_j}^*(m; s_i)$ and $\mathcal{I}_{F_k}^*(m; s_i)$ for $j, k \in \{1, 2\}$ are defined, we also define the following measures of precision,

$$\begin{aligned} P_{E_j, F_k}(m; s_i) &= \frac{|\mathcal{I}_{E_j}^*(m; s_i) \cap \mathcal{I}_{F_k}^*(m; s_i)|}{|\mathcal{I}_{E_j}^*(m; s_i)|}, \\ P_{F_j, E_k}(m; s_i) &= \frac{|\mathcal{I}_{F_j}^*(m; s_i) \cap \mathcal{I}_{E_k}^*(m; s_i)|}{|\mathcal{I}_{F_j}^*(m; s_i)|}, \\ P_{E_1, E_2}(m; s_i) &= \frac{|\mathcal{I}_{E_1}^*(m; s_i) \cap \mathcal{I}_{E_2}^*(m; s_i)|}{|\mathcal{I}_{E_1}^*(m; s_i)|}, \\ P_{F_1, F_2}(m; s_i) &= \frac{|\mathcal{I}_{F_1}^*(m; s_i) \cap \mathcal{I}_{F_2}^*(m; s_i)|}{|\mathcal{I}_{F_1}^*(m; s_i)|}. \end{aligned}$$

For target m , $P_{E_j, F_k}(m; s_i)$ is the fraction of nodes infected in the process with index

j in the encounter network that are also infected in the process with index k in the encounter network. $P_{F_j, E_k}(m; s_i)$ has a similar interpretation, by inverting the role of the two networks. In the analyses we will consider $P_{E_1, F_1}(m; s_i)$ and $P_{F_1, E_1}(m; s_i)$ only. $P_{E_1, E_2}(m; s_i)$ (resp., $P_{F_1, F_2}(m; s_i)$) is the fraction of nodes infected in the first simulation in the encounter (resp., friendship) network that are also infected in the second infection on the same network.

Experiments. We ran 5000 groups of simulations of the SEI process with $\beta = 0.5$. For each simulation, a single seed is selected at random among all nodes s_i such that $t_0(s_i) \leq 500$ in both the encounter and the friendship time-varying network (that is, we consider nodes that have neighbors on both networks by time $t = 500$). For each choice of the seed, we separately run two infection processes on the encounter network and two infection processes on the friendship network. Therefore, each seed selection is associated to four simulations (referred to as E_1, E_2, F_1, F_2). For target set size $m \in \{500, 1000, 2000, 5000, 10000, 20000\}$ and each of the 5000 seeds s_i , we consider the metrics above when they are defined.

Figure 5.9 plots the Jaccard similarity measures $J_{E_1, F_1}(m; s_i)$ (top-left panel), $J_{E_1, E_2}(m; s_i)$ (top-right panel), $J_{F_1, F_2}(m; s_i)$ (bottom panel). Observations for a given value of m constitute a block on the x -axis (larger values of m correspond to x positions on the right) and are represented with the same color. For a fixed value of m , relative x positions are irrelevant. For a given metric and each value m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations.

For all values of m , two-sample t-tests support the hypotheses that $J_{E_1, F_1}(m; s_i)$ has smaller average than $J_{E_1, E_2}(m; s_i)$ and $J_{F_1, F_2}(m; s_i)$, and that $J_{E_1, E_2}(m; s_i)$ has smaller average than $J_{F_1, F_2}(m; s_i)$ (p-values $< 2.2 \cdot 10^{-16}$).

A comparison between $J_{E_1, F_1}(m; s_i)$, $J_{E_1, E_2}(m; s_i)$, and $J_{F_1, F_2}(m; s_i)$ is not straight-

Table 5.1. Single seed infection on the time-varying networks. Jaccard similarity measures: empirical upper bounds, average of original measures, average of rescaled measures.

m	J_{E_1,F_1}^U	J_{E_1,E_2}^U	J_{F_1,F_2}^U	\bar{J}_{E_1,F_1}	\bar{J}_{E_1,E_2}	\bar{J}_{F_1,F_2}	J_{E_1,F_1}	J_{E_1,E_2}	J_{F_1,F_2}
500	0.110	0.459	0.786	0.365	0.259	0.563	0.040	0.119	0.443
1000	0.129	0.481	0.808	0.417	0.343	0.652	0.053	0.165	0.527
2000	0.161	0.507	0.797	0.442	0.450	0.737	0.071	0.228	0.588
5000	0.199	0.550	0.764	0.544	0.593	0.839	0.108	0.327	0.641
10000	0.225	0.592	0.749	0.697	0.681	0.910	0.157	0.403	0.682
20000	0.218	0.590	0.739	0.876	0.795	0.966	0.1918	0.469	0.714

forward for the lack of an upper bound for $J_{E_1,F_1}(m; s_i)$. There are $n_I = 31,735$ nodes in the intersection of the friendship and encounter network and $n_U = 142,967$ nodes in their union. Therefore, for large values of target m , $J_{E_1,F_1}(m; s_i)$ is upper bounded by $n_I/n_U = 0.2219$. A bound that is independent of s_i cannot be derived for general values of m , for which $J_{E_1,F_1}(m; s_i)$ is not constrained to have small values. However, $J_{E_1,E_2}(m; s_i)$ and $J_{F_1,F_2}(m; s_i)$ can be as large as 1 for all values of m .

To take this into account, we define a rescaled version of the Jaccard similarity,

$$\bar{J}_{E_1,F_1}(m; s_i) = \frac{J_{E_1,F_1}(m; s_i)}{J_{E_1,F_1}^U(m)},$$

where $J_{E_1,F_1}^U(m) = \max_{s_i} J_{E_1,F_1}(m; s_i)$ is an empirical upper bound for $J_{E_1,F_1}(m; \cdot)$, and the maximum is taken over all experiments. We similarly define rescaled measures $\bar{J}_{E_1,E_2}(m; s_i)$, and $\bar{J}_{F_1,F_2}(m; s_i)$, considering the empirical upper bounds $J_{E_1,E_2}^U(m)$ and $J_{F_1,F_2}^U(m)$. Table 5.1 reports the empirical upper bounds, and the averages of the original and rescaled measures of Jaccard similarity. Two-sample t-tests support the hypothesis that $\bar{J}_{E_1,F_1}(m; s_i)$ has a larger average than $\bar{J}_{E_1,E_2}(m; s_i)$ for all considered values of m (p-values smaller than 0.0078), whereas the null hypothesis of equal mean is not rejected for $m = 2000$. For all values of m , two-sample t-tests support the hypotheses that $\bar{J}_{E_1,F_1}(m; s_i)$ and $\bar{J}_{E_1,E_2}(m; s_i)$ have a smaller average than $\bar{J}_{F_1,F_2}(m; s_i)$ (p-values $< 2.2 \cdot 10^{-16}$).

Figure 5.10 plots the precision measures $P_{E_1, F_1}(m; s_i)$, $P_{E_1, E_2}(m; s_i)$, $P_{F_1, F_2}(m; s_i)$ in the top-left, top-right and bottom panels respectively. Observations for a given value of m constitute a block on the x -axis (larger m correspond to x positions on the right) and are represented with the same color. For a fixed value of m , relative x positions are irrelevant. For a given metric and each value m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations.

For all values of m , two-sample t-tests support the hypotheses that $P_{E_1, F_1}(m; s_i)$ and $P_{F_1, E_1}(m; s_i)$ have smaller average than both $P_{E_1, E_2}(m; s_i)$ and $P_{F_1, F_2}(m; s_i)$, and that $P_{E_1, E_2}(m; s_i)$ has smaller average than $P_{F_1, F_2}(m; s_i)$ (p-values $< 2.2 \cdot 10^{-16}$).

As before, a comparison between the precision measures is not possible for the lack of a straightforward upper bound for $P_{E_1, F_1}(m; s_i)$ and $P_{F_1, E_1}(m; s_i)$. There are $n_E = 133,038$ and $n = 41,664$ nodes in the encounter and friendship networks, respectively, and $n_I = 31,735$ nodes in their intersection. Therefore, for large values of m , $P_{E_1, F_1}(m; s_i)$ and $P_{F_1, E_1}(m; s_i)$ are upper bounded by $n_I/n_E = 0.2385$ and $n_I/n_F = 0.7617$, respectively. Bounds that are independent of s_i cannot be derived for general values of m . However, $P_{E_1, E_2}(m; s_i)$ and $P_{F_1, F_2}(m; s_i)$ can be as large as 1 for all values of m .

To take this consideration into account, we define rescaled version of the precision measures, for example,

$$\bar{P}_{E_1, F_1}(m; s_i) = \frac{P_{E_1, F_1}(m; s_i)}{P_{E_1, F_1}^U(m)},$$

where $P_{E_1, F_1}^U(m)$ is an empirical upper bound obtained taking the maximum over all simulations. $\bar{P}_{F_1, E_1}(m; s_i)$, $\bar{P}_{E_1, E_2}(m; s_i)$ and $\bar{P}_{F_1, F_2}(m; s_i)$ are similarly defined.

Table 5.2 reports the averages of the original and rescaled precision metrics. For all values of m , two-sample t-tests support the hypothesis that $\bar{P}_{F_1, F_2}(m; s_i)$ has a larger average than all other precision measures. For $m \in \{500, 1000, 20000\}$, two-sample

Table 5.2. Single seed infection on the time-varying networks. Precision measures: average of original and rescaled measures.

m	\bar{P}_{E_1, F_1}	\bar{P}_{F_1, E_1}	\bar{P}_{E_1, E_2}	\bar{P}_{F_1, F_2}	P_{E_1, F_1}	P_{E_1, F_1}	P_{E_1, E_2}	P_{F_1, F_2}
500	0.377	0.399	0.314	0.648	0.079	0.075	0.200	0.579
1000	0.428	0.461	0.407	0.733	0.102	0.101	0.266	0.658
2000	0.476	0.477	0.518	0.806	0.132	0.132	0.349	0.715
5000	0.582	0.584	0.660	0.883	0.194	0.194	0.469	0.766
10000	0.731	0.730	0.749	0.939	0.268	0.268	0.557	0.804
20000	0.895	0.895	0.853	0.980	0.321	0.321	0.633	0.833

t-tests support the hypotheses that $\bar{P}_{E_1, F_1}(m; s_i)$ and $\bar{P}_{F_1, E_1}(m; s_i)$ have larger average than $\bar{P}_{E_1, E_2}(m; s_i)$ (all p-values < 0.001). For $m \in \{500, 1000, 20000\}$, two-sample t-tests support the hypotheses that $\bar{P}_{E_1, F_1}(m; s_i)$ and $\bar{P}_{F_1, E_1}(m; s_i)$ have smaller average than $\bar{P}_{E_1, E_2}(m; s_i)$ (all p-values < 0.002). The null hypothesis that $\bar{P}_{E_1, F_1}(m; s_i)$ and $\bar{P}_{F_1, E_1}(m; s_i)$ have equal average is rejected only for $m \in \{500, 1000\}$, for which the former has larger average (p-values < $1e - 10$).




Figure 5.9. Single seed infection on time-varying networks: Jaccard similarity. The three panels show the metrics $J_{E_1, F_1}(m; s_i)$ (top-left), $J_{E_1, E_2}(m; s_i)$ (top-right) and $J_{F_1, F_2}(m; s_i)$ (bottom), for 5000 random choices of a single seeds, and different values of the target set size m . For each seed, two simulations on the friendship network and two simulations on the encounter network are run separately. The top-left panel considers, for each of the 5000 seeds, a pair of simulations on the two different networks. The top-right panel considers the 5000 pairs of simulations ran on the encounter network. The bottom panel considers the 5000 pairs of simulations ran on the friendship network. On the x -axis, observations for a given value of m form a block with a constant color (within the block, the x position is irrelevant). We only consider pairs (m, s_i) for which the metrics are defined. For a given metric and each value m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations.

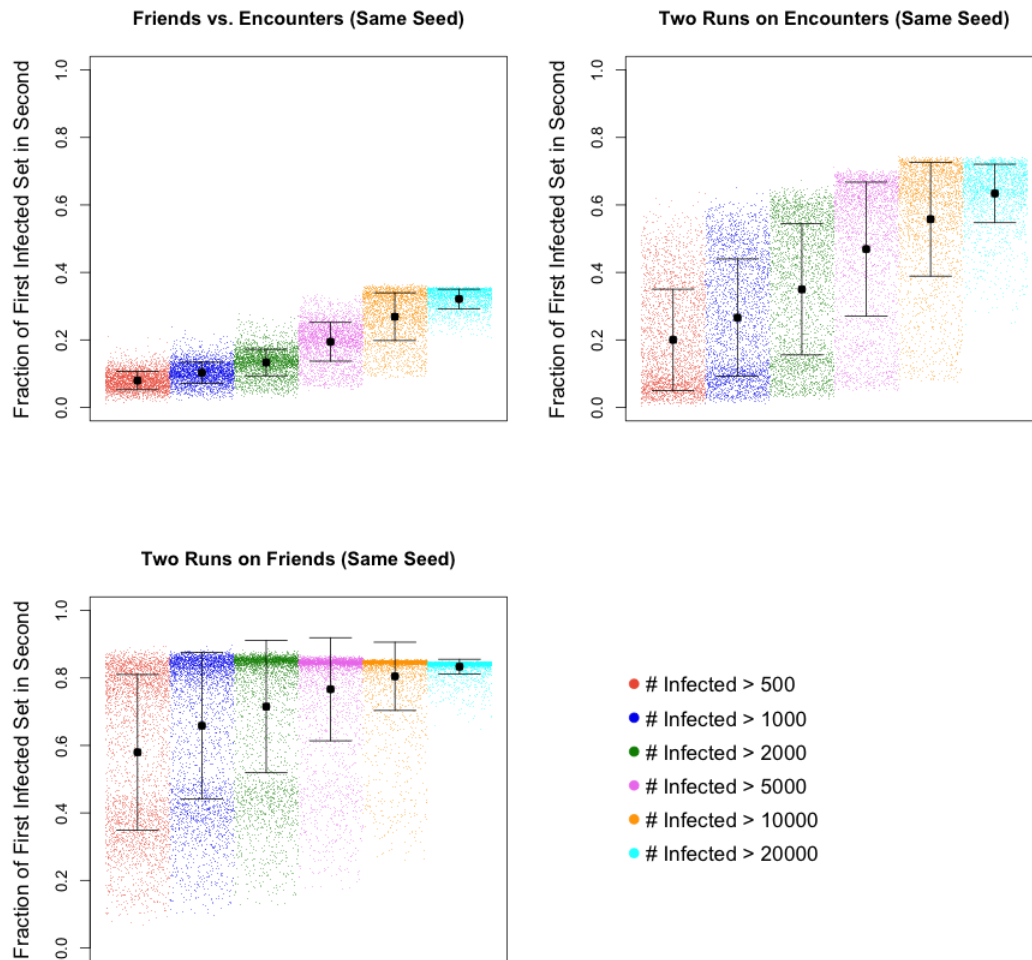


Figure 5.10. Single seed infection on time-varying networks: precision measures. The three panels show the metrics $P_{E_1, F_1}(m; s_i)$ (top-left), $P_{E_1, E_2}(m; s_i)$ (top-right) and $P_{F_1, F_2}(m; s_i)$ (bottom), for 5000 random choices of a single seeds, and different values of the target set size m . For each seed, two simulations on the friendship network and two simulations on the encounter network are run separately. The top-left panel considers, for each of the 5000 seeds, a pair of simulations on the two different networks. The top-right panel considers the 5000 pairs of simulations ran on the encounter network. The bottom panel considers the 5000 pairs of simulations ran on the friendship network. On the x -axis, observations for a given value of m form a block with a constant color (within the block, the x position is irrelevant). We only consider pairs m of s_i for which the metrics are defined. For a given metric and each value m , the black point represents the average of the metric over all observations such that the metric is defined, and the bars represent standard deviations.

5.7 Infection detection – Static networks

Let $G_E = (U_E, E_{U_E})$ and $G_F = (U_F, F_{U_F})$ be the giant components of the encounter and friendship static networks, respectively. For simplicity, we refer to such giant components as the static networks. U_E (resp. U_F) represents the set of nodes in the encounter (resp. friendship) giant component, and E_{U_E} (resp. F_{U_F}) represents the set of encounter (resp. friendship) ties restricted to U_E (resp. U_F). We have that $u_E = |U_E| = 113,187$ and $u_F = |U_F| = 168,923$, and $u_U = |U_E \cup U_F| = 210,899$. We consider the SEI process on the static networks G_F and G_E .

5.7.1 Single seed – Infection rate

We perform 5,000 simulations on each static network, setting $\beta = 0.01$. In each simulation, a single seed s_i is selected uniformly at random between all nodes in the corresponding network. Given that in a SEI process nodes never recover from infection, the entire population eventually becomes infected for each $\beta > 0$ and for each seed s_i . Recall that, for $0 \leq \alpha \leq 1$, $\tau(\alpha)$ represents the first time in which a α -fraction of the population is infected (for ease of notation, we omit the dependency on s_i). In this section, we study how the infection grows over time, that is, how $\tau(\alpha)$ grows with α .

Figure 5.11 relates the degree of the infection seed (i.e., encounter and friend degree) to the time $\tau(\alpha)$ to reach infection targets of $\alpha \in \{0.5\%, 1\%, 5\%, 10\%\}$. Top and bottom panels consider the SEI process on the encounter and friendship network, respectively. The x -axis show either the encounter degree (left panels) or the friend degree (right panels) of the seed (with degree at most 25).

In general, for all targets $\alpha \in \{0.5\%, 1\%, 5\%, 10\%\}$, increasing encounter (reps. friendship) degree is related to an initial steep decrease in the infection time on the encounter (reps. friendship) network, that then smooths out when the degree surpasses a

threshold.

In the encounter network (compare Figure 5.11, top-left panel), encounter degree larger than 10 results in a four-fold decrease of the infection time with respect to degree one, for all values of α (two-sample t-tests, means 188 and 42 for $\alpha = 0.5\%$, 191 and 45 for $\alpha = 1\%$, 201 and 55 for $\alpha = 5\%$, 209 and 62 for $\alpha = 10\%$, $p\text{-value} < 2.2 \cdot 10^{-16}$ for all α). The decrease of the infection time is slow for degree larger than 15 (OLS, restricted to seed with encounter degree larger than 15, degree coefficient -0.300 for all α , $p\text{-value} < 5.57e - 11$). The effect of the seed's friend degree on the infection speed on the encounter network is limited (degree coefficient -0.14 for all α , $p\text{-value} < 2.48e - 8$; compare Figure 5.11, top-right panel).

In the friendship network (compare Figure 5.11, bottom-right panel), friend degree larger than 5 results in a six-fold decrease of the infection time with respect to degree one, for all values of α (two-sample t-tests, means 134.97 and 20.63 for $\alpha = 0.5\%$, 135.60 and 21.23 for $\alpha = 1\%$, 137.48 and 23.11 for $\alpha = 5\%$, 139.37 and 25.01 for $\alpha = 10\%$, $p\text{-value} < 2.2 \cdot 10^{-16}$ for all α). The decrease of the infection time is slow for degree larger than 10 (OLS, restricted to seed with encounter degree larger than 10, degree coefficient -0.0547 for all α , $p\text{-value} < 9.63e - 14$). Larger encounter degree is not related to an equally steep decrease of the infection speed on the friendship network (compare Figure 5.11, bottom-left panel), despite its effect is somewhat (degree coefficient -1.270 for all α , $p\text{-value} < 2.08e - 7$), likely due to the low average encounter degree (mean 2.594).

If we look at how the infection grows over time, we observe an initial “incubation” period, during which the infected population is very small, followed by an explosion of the infection. Figure 5.11 plots the percentage of the infected population over time (up to 25%) for a sample of 60 randomly selected seeds on the encounter (right panel) network and 60 randomly selected seeds on the friendship network (left panel). Overall, on the

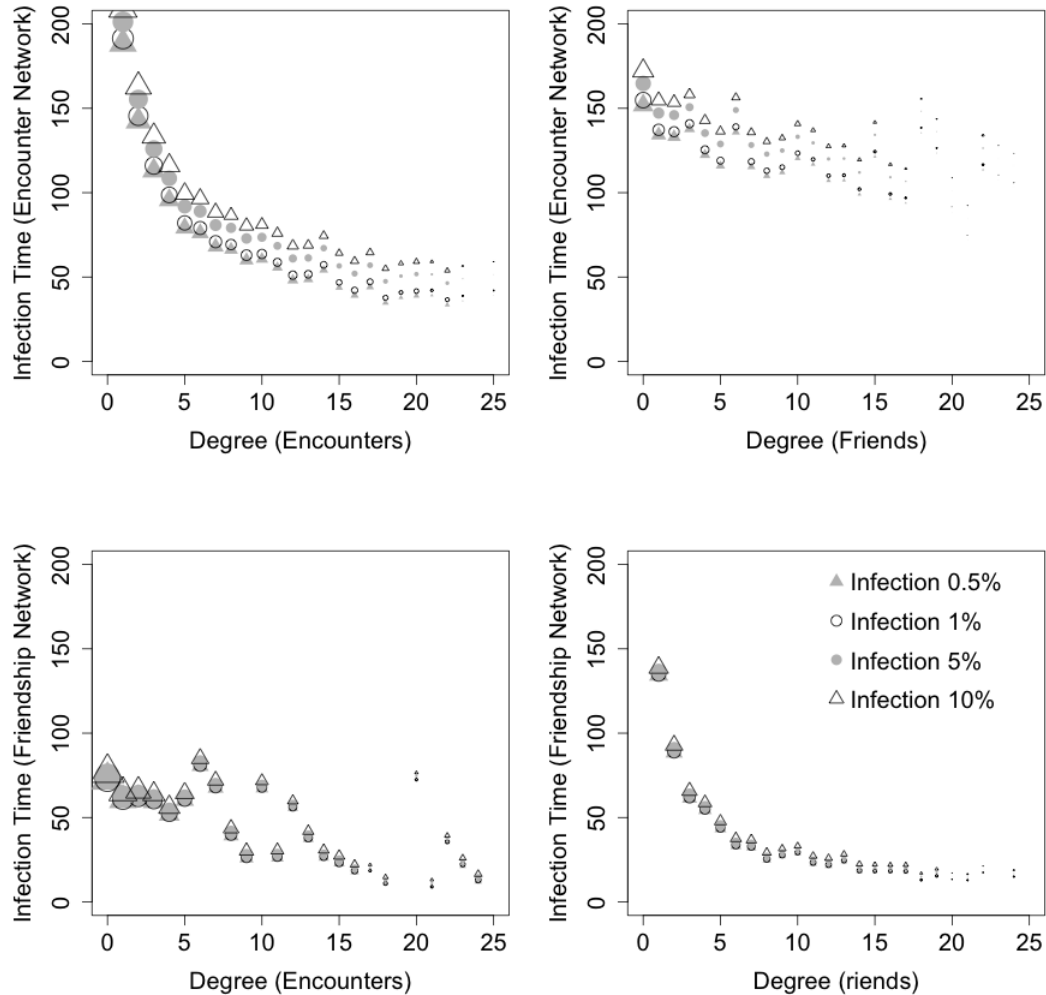


Figure 5.11. Infection speed versus degree. The plots relate the encounter and friendship degree of a seed node with the infection speed for different target infections α . 5000 simulations with $\beta = 0.01$ are run per network, selecting a seed uniformly at random for each simulation. The left panels relate encounter degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. The right panels relate friends degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. Point size are proportional to the logarithm of the number of observations for each degree.

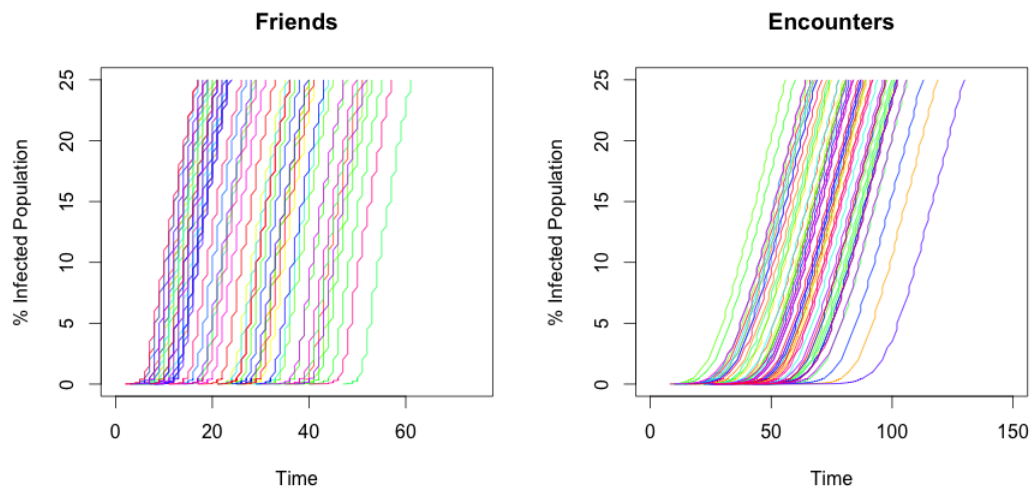


Figure 5.12. Growth of the infection over time. 60 simulations with $\beta = 0.01$ are shown for the friendship network (left) and for the encounter network (right). For each simulation, a seed is selected uniform at random and the infection starts at time $t = 0$. Colors are not meaningful. An initial “incubation” period, during which the infection spreads from the seed to its first neighbors, is followed by an explosion of the infection.

friendship network, an infection starting from a single seed takes on average 59.44 time units to infect an initial 0.01% of the population (about 17 nodes), with more connected nodes requiring less time (OLS, degree coefficient -0.273 , $p\text{-value} < 2.2 \cdot 10^{-16}$). On the encounter network, an infection starting from a single seed takes on average 107.78 time units to infect an initial 0.01% of the population (about 12 nodes), with more connected nodes requiring less time (OLS, degree coefficient -3.384 , $p\text{-value} < 2.2 \cdot 10^{-16}$).

The incubation period is determined by the stochasticity of the infection process, represented by the parameter β , and it is in large part constituted by the time required by the seed to infect the first neighbor. Indeed, the first infection happens, on average, after 39.68 time units in the friendship network (decreasing with degree, OLS, -0.1608 , $p\text{-value} 4.46e - 11$), and after 53.07 time units in the encounter network (decreasing with degree, OLS, -2.011 , $p\text{-value} < 2 \cdot 10^{-16}$).

5.7.2 Single seed – Sensor monitoring

Instead of monitoring the entire population, in each run of the SEI process, we consider a random set of sensors composed by 1% of the population. Sensors are selected in the two ways described above: random sensors and friend sensors (where the selection is based on friendship rather than encounter, even when considering a process spreading on the encounter network). We perform 5,000 simulations on each static network and each sensor type, setting $\beta = 0.01$ (stochastic infection). In each simulation, a single seed is selected uniformly at random between all nodes in the network.

Figure 5.13 plots the average time to detect a 5% infection of the sensors versus the seed degree, on the encounter network (top panels) and friendship network (bottom panels). The x -axis shows either the encounter degree (left panels) or the friend degree (right panels) of the seed (degree at most 25).

On the encounter network (compare Figure 5.13, top panels), friends sensors guarantee earlier detection than random sensors. The average detection time for a 5% infection of friends sensors (135.36 time units) is smaller than that of random sensors (141.66 time units, t-test, p-value 0.00487). The average detection time for a 10% infection of friends sensors (139.96 time units) is smaller than that of random sensors (149.13 time units, t-test, p-value $4.109 \cdot 10^{-5}$). The average detection time for a 25% infection of friends sensors (151.17 time units) is smaller than that of random sensors (168.51 time units, t-test, p-value $9.913e - 15$). The earlier detection provided by friend sensors over random sensors is not statistically significant for targets of 0.05% and 1% infection.

On the friendship network (compare Figure 5.13, bottom panels), despite friend sensors provide a lower average detection time than random sensors, the difference is not statistically significant for any target infection rate.

The results above are driven by the stochastic incubation time needed to get the infection started, driven by the parameter β , as we observed in the previous section. In order to control for such randomness, we perform 5,000 additional simulations for each time-varying network and sensor type, setting $\beta = 1$ (certain infection). This choice allows to study the effect of the structural properties of the selected sensors on the infection detection time. Friend sensor provide faster detection of the infection both on the friendship and the encounter network, and for all targets α .

Figure 5.14 plots the average time to detect a 25% infection of the sensors versus the seed degree, on the encounter network (top panels) and friendship network (bottom panels). The x -axis shows either the encounter degree (left panels) or the friend degree (right panels) of the seed (degree at most 25). On the encounter network (compare Figure 5.14, top panels), the average detection time for a 0.5% infection of friends sensors is 3.3516 time units (versus 3.7426 for random sensors), for a 1% infection is 3.5488 (versus 3.9894), for a 5% infection is 4.1192 (versus 4.5660), for a 10% infection is 4.4008 (versus 4.8526), for a 25% infection is 4.8922 (versus 5.3340), and all value are statistically significant (t-tests, p -values $< 2 \cdot 10^{-16}$). On the friendship network (compare Figure 5.14, bottom panels), the average detection time for a 0.5% infection of friends sensors is 3.3516 time units (versus 3.7426 for random sensors), for a 1% infection is 2.4748 (versus 2.8432), for a 5% infection is 2.6004 (versus 2.9804), for a 10% infection is 2.9286 (versus 3.6092), for a 25% infection is 3.5142 (versus 3.9056), and all value are statistically significant (t-tests, p -values $< 2 \cdot 10^{-16}$).




Figure 5.13. Sensor infection monitoring versus seed degree. The plots show the average time to infect 5% of the sensors versus the degree of the infection seed. 5000 simulations with $\beta = 0.01$ are run per network and per sensor type. For each simulation, a seed is selected uniformly at random, and the sensor size is 1% of the total population. The left panels relate encounter degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. The right panels relate friends degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. Point size are proportional to the logarithm of the number of observations for each degree.

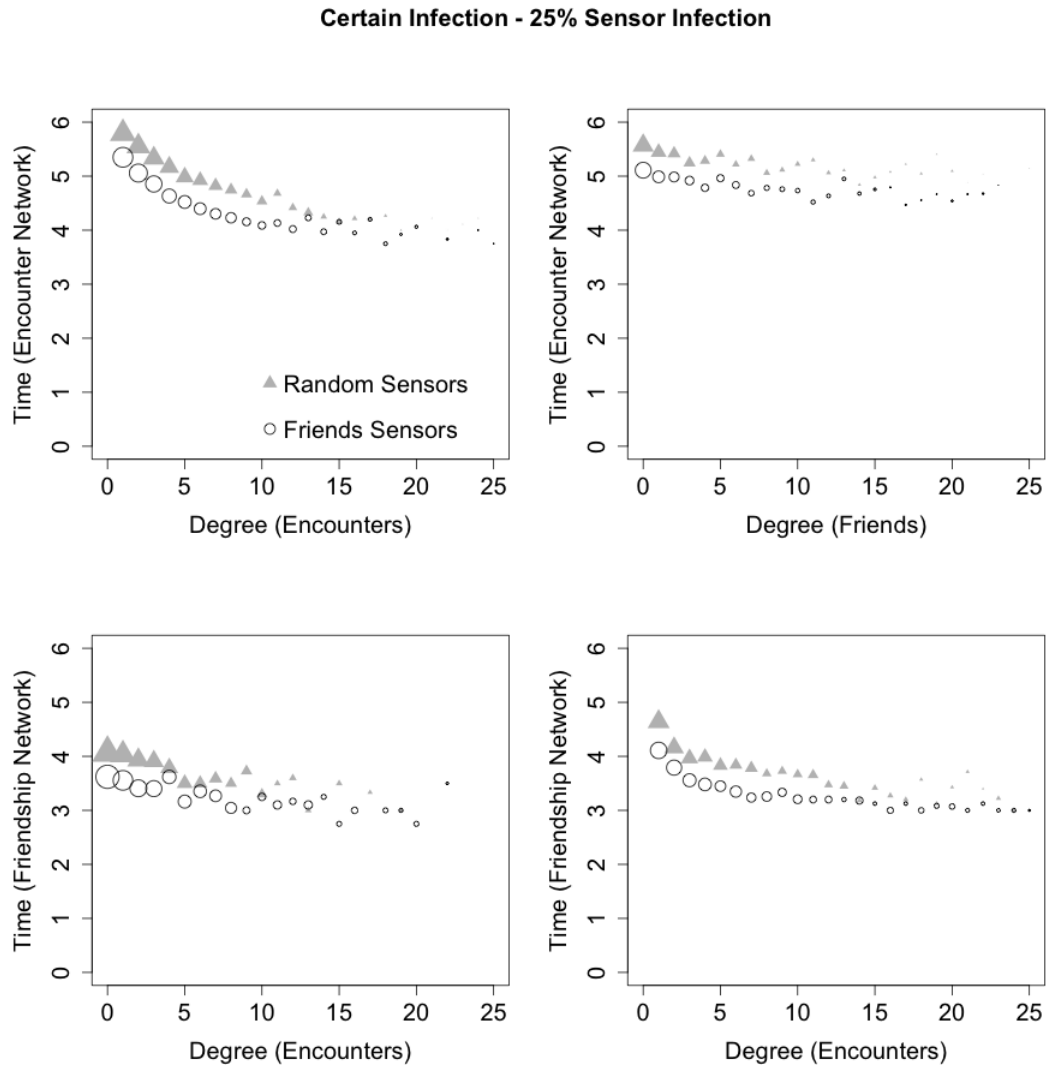


Figure 5.14. Sensor infection monitoring versus seed degree. The plots show the average time to infect 25% of the sensors versus the degree of the infection seed. 5000 simulations with $\beta = 1$ (certain infection) are run per network and per sensor type. For each simulation, a seed is selected uniformly at random, and the sensor size is 1% of the total population. The left panels relate encounter degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. The right panels relate friends degree of the seed with infection on the encounter network (top) and friendship network (bottom), for degree of at most 25. Point size are proportional to the logarithm of the number of observations for each degree.

5.8 The infected population – Static networks

In the previous section, we analyzed how infection processes spread on the encounter and friendship static networks at a macroscopic level, focusing on the infection detection time and on the infection growth over time. In this section, we take a microscopic look at the process, moving our attention to the sets of nodes that become infected. In particular, we consider seed nodes that are present in both the friendship and encounter network, and we will compare the sets of nodes that become infected in processes starting at the same seed but evolving on the two different networks. By comparing several independent runs of the infection process starting at each seed, we will observe that the unpredictability within a given network is substantially lower than the unpredictability between the two different networks.

We ran 10,000 groups of simulations of the SEI process with $\beta = 0.01$ (stochastic infection). For each group of simulation, a single seed is selected at random among all nodes s_i in the intersection of the two networks ($u_I = |U_E \cap U_F| = 71,211$). For each choice of the seed, we separately run two infection processes on the encounter network and two infection processes on the friendship network.

The Jaccard similarity measures $J_{E_j, F_k}(m; s_i)$ for $j, k \in \{1, 2\}$, $J_{E_1, E_2}(m; s_i)$ and $J_{F_1, F_2}(m; s_i)$, and the precision measures $P_{E_j, F_k}(m; s_i)$ for $j, k \in \{1, 2\}$, $P_{E_1, E_2}(m; s_i)$, and $P_{F_1, F_2}(m; s_i)$ are defined as in the case of time-varying networks. Observe that, as all nodes eventually become infected in a SEI process on a static network, these quantities are defined for each s_i and $m < n$, where n is the number of nodes in the network. We consider target set size $m \in \{500, 1000, 2000, 5000, 10000, 20000\}$.

Figure 5.15 plots the Jaccard similarity measures $J_{E_1, F_1}(m; s_i)$, $J_{E_1, E_2}(m; s_i)$ and $J_{F_1, F_2}(m; s_i)$ in the top-left, top-right and bottom panels respectively. Figure 5.16 plots the precision measures $P_{E_1, F_1}(m; s_i)$, $P_{E_1, E_2}(m; s_i)$ and $P_{F_1, F_2}(m; s_i)$ in the top-left, top-right

and bottom panels respectively. Observations for a given value of m constitute a block on the x -axis (larger m corresponds to x positions on the right) and are represented with the same color. For a fixed value of m , relative x positions are irrelevant. For a given metric and each value m , the black point represents the average of the metric over all the observations and bars represent standard deviations.

$J_{E_1,F_1}(m; s_i)$ has smaller average than $J_{E_1,E_2}(m; s_i)$, $J_{F_1,F_2}(m; s_i)$, and for $m > 500$, $J_{E_1,E_2}(m; s_i)$ has larger average than $J_{F_1,F_2}(m; s_i)$ (two-paired t-tests, p-values $< 2.2 \cdot 10^{-16}$). Similarly, $P_{E_1,F_1}(m; s_i)$ and $P_{F_1,E_1}(m; s_i)$ have smaller average than $P_{E_1,E_2}(m; s_i)$, $P_{F_1,F_2}(m; s_i)$, and for $J_{E_1,E_2}(m; s_i)$ has smaller average than $J_{F_1,F_2}(m; s_i)$ (two-paired t-tests, p-values $< 2.2 \cdot 10^{-16}$).

However, these result should be interpreted with caution, as it is not straightforward to rigorously compare the quantities for all values of m . The metrics $J_{E_1,E_2}(m; s_i)$, $J_{F_1,F_2}(m; s_i)$, $P_{E_1,E_2}(m; s_i)$ and $P_{F_1,F_2}(m; s_i)$ can be as large as 1 for all values of m . Instead, for large m , $J_{E_j,F_k}(m; s_i)$ is upper bounded by $u_I/u_U = 0.338$, $P_{E_j,F_k}(m; s_i)$ is upper bounded by $u_I/u_E = 0.629$, and $P_{F_j,E_k}(m; s_i)$ is upper bounded by $u_I/u_F = 0.422$. For general values of m , tight upper bounds for these quantities depend on s_i and therefore on the network structure. Therefore, we proceed by computing empirical upper bounds $J_{\cdot,\cdot}^U(m)$ and $P_{\cdot,\cdot}^U(m)$ over all simulations, and rescaling the quantities above by dividing by the corresponding bounds and obtaining rescaled metrics $\bar{J}_{\cdot,\cdot}(m, s_i)$ and $\bar{P}_{\cdot,\cdot}(m, s_i)$.

Table 5.3 reports the averages of the original and rescaled Jaccard similarity measures. Table 5.4 reports the averages of the original and rescaled precision measures. For all values of m , $\bar{J}_{E_1,F_1}(m; s_i)$ has smaller average than $\bar{J}_{E_1,E_2}(m; s_i)$ and $\bar{J}_{F_1,F_2}(m; s_i)$, and for $m > 500$, $\bar{J}_{E_1,E_2}(m; s_i)$ has larger average than $\bar{J}_{F_1,F_2}(m; s_i)$ (two-sample t-tests, p-values $< 2.2 \cdot 10^{-16}$). For all values of m , $\bar{P}_{E_1,F_1}(m; s_i)$ has smaller average than $\bar{P}_{E_1,E_2}(m; s_i)$ and $\bar{P}_{F_1,F_2}(m; s_i)$, whereas $\bar{P}_{F_1,E_1}(m; s_i)$ has smaller average than $\bar{P}_{F_1,F_2}(m; s_i)$, only for $m \in \{500, 1000, 2000, 5000\}$ and larger for $m \in \{10000, 20000\}$ (two-sample

Table 5.3. Single seed infection on the static networks. Jaccard similarity measures: empirical upper bounds, average of original measures, average of the rescaled measures.

m	\bar{J}_{E_1, F_1}	\bar{J}_{E_1, E_2}	\bar{J}_{F_1, F_2}	J_{E_1, F_1}	J_{E_1, E_2}	J_{F_1, F_2}
500	0.283	0.300	0.402	0.012	0.046	0.055
1000	0.404	0.538	0.517	0.0211	0.0677	0.100
2000	0.553	0.704	0.650	0.0341	0.098	0.163
5000	0.749	0.877	0.823	0.061	0.155	0.251
10000	0.852	0.931	0.911	0.090	0.212	0.309
20000	0.929	0.956	0.952	0.128	0.292	0.365

Table 5.4. Single seed infection on the static networks. Precision measures: average of original and rescaled measures.

m	\bar{P}_{E_1, F_1}	\bar{P}_{F_1, E_1}	\bar{P}_{E_1, E_2}	\bar{P}_{F_1, F_2}	P_{E_1, F_1}	P_{E_1, F_1}	P_{E_1, E_2}	P_{F_1, F_2}
500	0.267	0.250	0.312	0.406	0.034	0.021	0.088	0.111
1000	0.359	0.370	0.559	0.471	0.055	0.034	0.126	0.190
2000	0.503	0.456	0.701	0.558	0.085	0.055	0.179	0.291
5000	0.686	0.692	0.883	0.731	0.136	0.101	0.268	0.409
10000	0.788	0.846	0.916	0.833	0.181	0.154	0.350	0.475
20000	0.885	0.920	0.955	0.908	0.237	0.219	0.452	0.535

t-tests, p-values $< 2.2 \cdot 10^{-16}$).

The rescaled measures suggest that the network structure has a large impact on the spread of the infection between the friendship and encounter networks, at a microscopic level.

5.9 Acknowledgments

Chapter 5, in part, is currently being prepared for submission for publication of the material. L. Coviello, and M. Franceschetti. The dissertation author was the primary investigator and co-author of this paper.

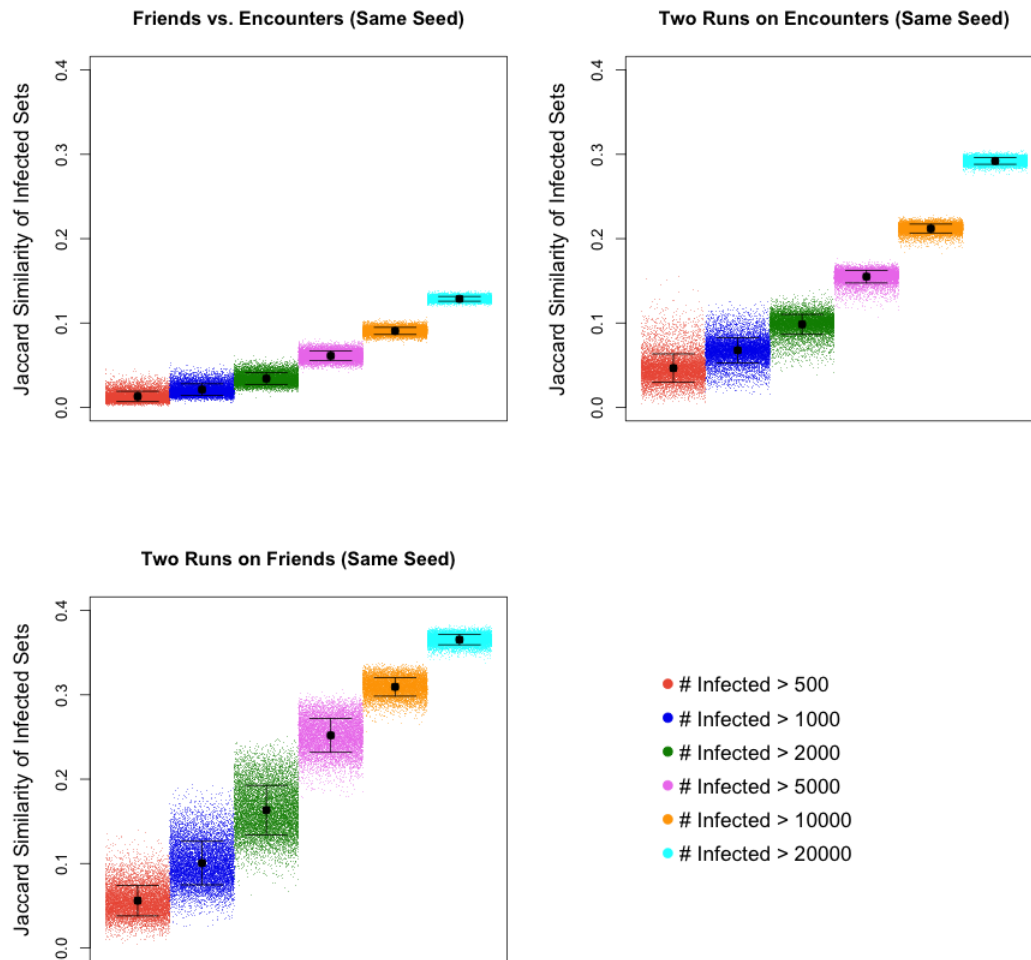


Figure 5.15. Single seed infection on the static networks: Jaccard similarity. The three panels show the metrics $J_{E_1, F_1}(m; s_i)$ (top-left), $J_{E_1, E_2}(m; s_i)$ (top-right) and $J_{F_1, F_2}(m; s_i)$ (bottom), for 10,000 random choices of a single seeds, and different values of the target set size m . For each seed, two simulations on the friendship network and two simulations on the encounter network are run separately. The top-left panel considers, for each of the 10,000 seeds, a pair of simulations on the two networks. The top-right panel considers the 10,000 pairs of simulations ran on the encounter network. The bottom panel considers the 10,000 pairs of simulations ran on the friendship network. On the x -axis, observations for a given value of m form a block with a constant color (within the block, the x position is irrelevant). For a given metric and each value m , the black point represents the average of the metric over all the observations and bars represent standard deviations.

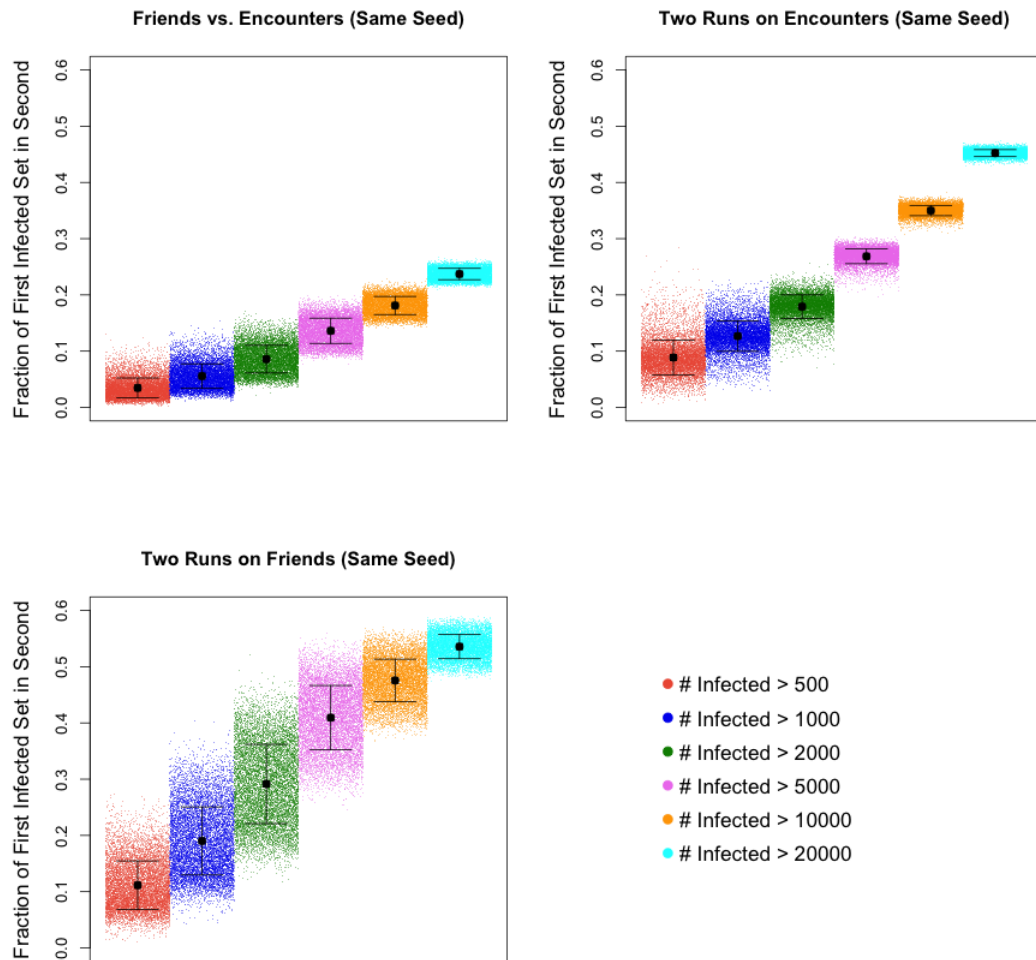


Figure 5.16. Single seed infection on the static networks: precision measures. The three panels show the metrics $P_{E_1, F_1}(m; s_i)$ (top-left), $P_{E_1, E_2}(m; s_i)$ (top-right) and $P_{F_1, F_2}(m; s_i)$ (bottom), for 10,000 random choices of a single seeds, and different values of the target set size m . For each seed, two simulations on the friendship network and two simulations on the encounter network are run separately. The top-left panel considers, for each of the 10,000 seeds, a pair of simulations on the two networks. The top-right panel considers the 10,000 pairs of simulations ran on the encounter network. The bottom panel considers the 10,000 pairs of simulations ran on the friendship network. On the x -axis, observations for a given value of m form a block with a constant color (within the block, the x position is irrelevant). For a given metric and each value m , the black point represents the average of the metric over all the observations and bars represent standard deviations.

Chapter 6

Query incentive networks with split-contracts: robustness to individuals' selfishness

6.1 Introduction

A challenging class of crowdsourcing problems requires an interested party to provide incentives for large groups of people to contribute to the search and retrieval of rare information [205, 147, 94]. The small world problem, i.e. distributed routing of messages to unknown individuals, is the seminal example of this class and has illustrated the difficulty of the approach for almost 50 years [155, 208, 193, 62, 214]. In this class of problems, individuals in the social network act as intermediaries to create a channel between the querier and the answer. Observe that the chief difficulty of this approach is to offer incentives to the individuals to propagate the query further in the network as well as to return the answer all the way back to the querier [62]. The goal is therefore to incentivize participation of the users using some form of (possibly financial) reward. In this way, a node who does not know the answer but is offered a sufficiently high reward can act as intermediary and propagate the query by offering the neighbors a share of its reward. This setting models the social network as a marketplace of information where the users strategically act in order to maximize their utility, and raises several questions

about the system's performance and the incentive propagation, the main one being: can we retrieve the answer to a difficult query when given a limited budget?

The Defense Advanced Research Projects Agency (DARPA), a research organization of the United States Department of Defense, designed a so called "Network Challenge" that conveyed a positive answer to this question.¹ The challenge consisted of locating ten moored red weather balloons placed at ten undisclosed locations in the continental United States. A single \$40,000 cash prize was allocated for the first participant to submit the correct latitude and longitude (within one mile error) of all ten balloons within the contest period. In particular, the competition consisted in recruiting a team to achieve the goal. This task posed varied issues of large-scale, time-critical mobilization. In particular, in order to guarantee the participation and coordination of a large team, an adequate structure of economic incentives had to be built.

The MIT Media Laboratory team won the competition in less than 9 hours, adopting a recruitment scheme based on recursive incentives.² Specifically, using the \$40,000 they could possibly win, they allocated an amount of \$4,000 for finding each balloon. For each balloon, they would distribute the \$4,000 up the chain of participants leading to successful balloon spotting, as described in their website: "[In the case we win the competition,] we're giving \$2,000 per balloon to the first person to send us the correct coordinates, but that's not all – we're also giving \$1000 to the person who invited them. Then we're giving \$500 whoever invited the inviter, and \$250 to whoever invited them, and so on...". This is equivalent to say that a node u who does not have the desired answer, can offer its friends a $1/2$ -split contract, stipulating that if the answer is found in the subtree of a child v of u , then u will get back from v a $1/2$ fraction of whatever amount v gets. However, if u is not the querier, the total amount pocketed by u is less, as u has to

¹<https://networkchallenge.darpa.mil/>

²<http://balloon.media.mit.edu/>

give a $1/2$ fraction of its reward to its recruiter.

While the success of this strategy has been hailed as an empirical testimony to the power of incentive structures [206], the theoretical efficiency of the proven scheme has remained an open question, and motivates this work. In particular, we analyze this economic structure in the model for query incentive networks introduced by Kleinberg and Raghavan in [129]. This model considers a competitive environment where every node plays strategically. To fit the *split contracts* to this model, we generalize the splits to any fraction $0 < \rho < 1$, in the sense that any node u can offer a child v a ρ -split contract stipulating the following: if v has the answer, then v would pocket a $(1 - \rho)$ of the whole reward while returning a fraction ρ to u ; if v does not have the answer, then v can in turn offer some ρ' -split to its (still unrecruited) friends, and so on. Given the strategic setting, nodes will choose the splits to offer to their children so to maximize their expected payoffs; observe that contracts between different nodes can have different splits — and this is indeed the case in the Nash equilibrium as our results show. The details of the original model introduced in [129] follow.

6.1.1 Query Incentive Networks

The scenario of interest is that of a node, the root, that is willing to invest some amount r^* to retrieve certain information from a large network in which every node plays strategically. The main goal is to characterize the tradeoff between the investment and the rarity of the information. The model, introduced by Kleinberg and Raghavan [129], is as follows: the querier node is the root of an infinite d -ary tree, where each node possesses independently the desired information with probability $1/n$, where n represents the rarity of the answer. The root offers each child u a “*fixed-payment*” contract of r^* , stipulating that the root will pay u that amount upon u providing the answer. The query propagates down the tree according to the following scheme: every node u has an integer-valued

function f_u encoding its strategies; if u is offered a reward of r by its parent and does not possess the answer, then in turn it offers a reward of $1 \leq f_u(r) \leq r - 1$ to its children. When the answer to the query is found, the root selects for payment one among the answer-holders using a fixed non-strategic rule. The payment is then propagated down through the path to that selected node, with each node along the path pocketing its share. If an intermediate node u on this path was offered r by its parent, then its overall payoff is $r - f_u(r) - 1$, where the *unit cost* is associated with the act of returning the answer³. The game-theoretical aspect of the model is that any node u chooses the function f_u so to maximize its payoff. To break ties, it is assumed that a node who is offered a reward of one (and does not possess the answer) will always forward the query to its children, even if its expected payoff is zero (since the unit reward would be spent when returning the answer up to its parent).

As pointed out in [129], there is a subtle deficiency with a deterministic tree: the Nash equilibria of a game played in a deterministic network tacitly assume that the nodes know the entire network. Indeed, in a Nash equilibrium, each node chooses its best strategy by knowing the strategies of every other node. However, this is unrealistic, as we want to model a setting where nodes are only aware of their neighbors. To deal with this technical issue, Kleinberg and Raghavan consider a network that can be thought as a branching process from the root. In particular, the number of children of each node is chosen independently from a binomial distribution $\text{Bin}(d, q)$, where q is a constant probability of a node being present. The expected number of children of a node — i.e., the branching factor — is then $b = qd$. By classical results in the theory of branching

³As observed in [129], if nodes placed no value on this answering effort then the root could simply invest an arbitrarily small reward $\epsilon > 0$, and it would retrieve an answer because each node would have a positive payoff from participating in the game and returning the answer. To avoid this situation, a unit price is placed on the effort of returning the answer, while the cost of participating to the game is zero. This is motivated by the fact that the cost of forwarding requests to a list of friends is typically considered negligible in peer-to-peer and social-network systems [120, 219, 220] (see [129] for additional details on the motivations).

processes, if $b < 1$ the process dies out almost surely; therefore there is no amount that the root can offer to obtain an answer with constant probability if the rarity n of the answer is large enough. Instead, for any $b > 1$, there is a constant non-zero probability that the process will generate infinitely many nodes, so that the answer is present within the first $O(\log n)$ levels of the tree with high probability. Nevertheless, Kleinberg and Raghavan show that in the Nash equilibrium the investment needed at the root can be much larger than logarithmic in n . Specifically, while an investment $r^* = O(\log n)$ is sufficient to retrieve the answer with constant probability for $b > 2$, an investment of $r^* = n^{\Theta(1)}$ is needed when $1 < b < 2$. That is, in the latter case the root must invest a reward that is exponentially larger than the expected distance from the closest answer.

Arcaute *et al.* [15] generalized the work in [129] showing that this threshold behavior at $b = 2$ still holds for arbitrary Galton-Watson branching process. They also proved that in a *ray* — a deterministic infinite path ($b = 1$, but with zero extinction probability) — the reward needed is super-exponential in the expected depth of the search tree, that is $r^* = \Omega(n!)$. Finally, they observed that this threshold behavior vanishes if the root desires to find the answer with probability tending to 1: if the desired probability is $1 - 1/n$, then for any branching process with $b > 1$ and no extinction, the needed reward is $n^{\Theta(1)}$.

6.1.2 Our results

We present a theoretical study of the multi-level marketing strategies adopted by the winning team of the DARPA Network Challenge. Given the strong affinity between this challenge and the model of query incentive networks introduced in [129, 15], we frame these strategies in this model by considering split contracts as the possible offers between nodes.

Our main result is that split contracts, unlike fixed-payment contracts, are robust

to a strategic environment, where every node selfishly determines the offers to its children based on the offer received from its parent. We show that for any constant $\varepsilon > 0$ and Galton-Watson branching process with $b > 1$, the Nash equilibrium with split contracts uses an investment of $r^* = O(\log n)$ to retrieve the answer with probability at least $1 - \zeta - \varepsilon$, where ζ is the extinction probability of the process. As the expected distance to the closest answer is $\Theta(\log n)$ and nodes pay a unit cost to return the answer, this is a constant approximation with respect to an ideal centralized non-strategic setting. In other words, the price of anarchy of the game with split contracts is constant (ignoring some pathological equilibria, see Section 6.4 and Section 6.7.8).

Unlike previous work that assumed the parameters of the branching process to be held constant, we are also able to characterize the dependence of the investment with respect to the branching process and the success accuracy. This allows us to show additional improvements of split contracts over fixed-payment contracts: for example, for branching processes with no extinction, an investment of $O(n \log n)$ is enough to retrieve the answer with probability at least $1 - 1/n$, improving upon the $n^{\Theta(1)}$ investment provided in [15]. In fact, our result is even stronger since it guarantees a success probability of at least $1 - \zeta - 1/n$ in general branching processes. In the case of a ray (where the expected distance from the closest answer is n), we show that the investment needed to find the answer with constant probability is $O(n^2)$, while $\Omega(n!)$ is needed when using fixed-payment contracts [15].

6.1.3 Additional related work

Pickard et al. [177] described and analyzed the winning strategy of the DARPA Network Challenge. However, we distinguish ourselves from [177] in both aims and methods. The authors of [177] are mainly concerned with the motivation of the exact 1/2-split winning strategy that was implemented by the MIT Media Laboratory, for which

they show that it is in the participants' interest to recruit the highest number of friends and back the theory with an empirical analysis of the diffusion cascades. Our work considers the more general setting of split contracts in the model of query incentive networks introduced in [129] and analyzes the efficiency, in terms of investment, of the Nash equilibria.

In the context of query incentive networks with fixed-payment contracts, Kota and Narahari [133] applied the results of general branching processes from [15] to analyze the reward when the degree distribution follows a power-law and the desired success probability is at least $1 - 1/n$ and show a threshold behavior of the reward with respect to the scaling exponent. Dikshit and Yadati [61] considered the issue of the quality of the answers in query incentive networks. In particular, they define a quality conscious model of incentives and derive the same threshold behavior around the branching factor $b = 2$ found in [15, 129].

It is worth to mention additional related work that is not in the context of query incentive networks. Emek et al. [66] studied strategies of multi-level marketing, in which each individual is rewarded according to direct and indirect referrals, and show that geometric reward schemes are the only guarantee to certain desirable properties. Our setting is substantially different from [66], as the reward is based on referral rather than information retrieval. Douceur and Moscibroda [63] proposed the lottery tree as a mechanism to incentivize the adoption of a distributed systems and the solicitation of new participants. Influence in social networks is also related to our work. Kempe et al. [126] considered the algorithmic question of selecting an influential set of individuals. Jackson and Yariv [109] proposed a game-theoretic framework to model incentives in adoption processes. Hartline et al. [95] studied influence in social networks from a revenue maximization point of view. Singer [195] developed incentive-compatible mechanisms for influence maximization in several models.

6.2 Preliminaries

We model the network as a tree generated via a Galton-Watson branching process with offspring distribution $\{c_k\}_{k=0}^d$, that is, c_k is the probability that any node has exactly k children and $\sum_{k=0}^d c_k = 1$. We adopt the convention that the root of the tree is at level 0, its children at level 1, and so on. The probability generating function of the offspring distribution is given by

$$\Psi(x) = \sum_{k=0}^d c_k x^k, \quad 0 \leq x \leq 1.$$

The branching factor of the process is defined as $b = \Psi'(1) = \sum_{k=0}^d k c_k$. A fundamental result in the theory of Galton-Watson processes states that the extinction probability ζ of a branching process is the smallest non-negative root of the equation $x = \Psi(x)$. It follows that $\zeta = 1$ if and only if $b < 1$, or $b = 1$ with $c_0 > 0$, and that $0 \leq \zeta < 1$ otherwise. For classical theory on Galton-Watson branching processes we refer to [18].

We assume that each node in the network possesses the answer to the query independently of the other nodes with probability $1/n$, where n represents the *rarity* of the answer. Note that n is the expected number of nodes to query before finding the answer. For $i \geq 0$, let ϕ_i be the probability that no node at level $j \leq i$ has the answer and $\lambda_i = \phi_{i-1} - \phi_i$ be the probability that some node at level i and no node at a lower level possesses the answer. (These probabilities are over the randomness of the branching process and of the process assigning the answers.) Moreover, conditional on the event that the branching process with probability generating function Ψ does not die out, let $h_\Psi(\varepsilon, n)$ be the minimum integer i such that $\phi_i < \varepsilon$. For branching factor $b > 1$, we have that $h_\Psi(\varepsilon, n) = O(\log n)$ for any $\varepsilon = n^{-O(1)}$, whereas in the case of $b = 1$ and $c_0 = 0$ (i.e., a ray), $h_\Psi(\varepsilon, n) = n \ln \frac{1}{\varepsilon}$.

Assume that r^* is the investment available at the root, which desires to retrieve

the answer with probability at least $1 - \zeta - \varepsilon$, for a given success accuracy $\varepsilon > 0$. With the notation introduced so far, we will show that for any constants $b > 1$ and $\varepsilon > 0$ an investment of $r^* = O(h_\Psi(\varepsilon, n))$ suffices to propagate the query down to level $h_\Psi(\varepsilon, n)$ of the tree, and hence to retrieve the answer with probability at least $1 - \zeta - \varepsilon$. For ease of analysis, we assume that the root is not willing to explore the tree below level $h_\Psi(\varepsilon, n)$, that is, we *truncate* the tree at that height.

6.2.1 Split contracts

We now formalize the notion of *split contracts*. Every node including the root can offer a ρ -split contract to its children, for some $0 < \rho < 1$, stipulating the following. If the root offers a ρ -split to a child u who possesses the answer, then u receives a payment of r^* but is required to return a ρ fraction to the root, earning a total of $r^*(1 - \rho) - 1$, where we introduced a unit cost for returning the answer to the parent, as in [129, 15]. If instead u does not possess the answer then it might decide to propagate the query to its children, according to its strategy $f_u(\cdot)$, that is, offering a $f_u(\rho)$ -split contract to its children. If one among u 's children possesses the answer, then u receives an $f_u(\rho)$ fraction of the reward but it gives a ρ fraction back to the root and pays the unit cost to return the answer, with an overall earning of $r^*(1 - \rho)f_u(\rho) - 1$. In general, consider a node u_ℓ which is reached by a query and possesses the answer, and let u_0, u_1, \dots, u_ℓ be the path connecting the root to u_ℓ , where u_0 is the root. Then, if the root offered a ρ_{u_0} -split to its children, and $\rho_{u_i} = f_{u_i}(\rho_{u_{i-1}})$ is the split offered by u_i to its children for all $i < \ell$, then the root u_0 (who need not to pay the unit cost) receives a payoff of

$$r^* \cdot \rho_{u_0} \cdot f_{u_1}(\rho_{u_0}) \cdot f_{u_2}(f_{u_1}(\rho_{u_0})) \cdots f_{u_{\ell-1}}(f_{u_{\ell-2}}(\cdots)) = r^* \cdot \prod_{j=0}^{\ell-1} \rho_{u_j}.$$

Similarly, for $1 \leq i \leq \ell$, the payoff of node u_i is $(r^*(1 - \rho_{u_{i-1}}) \cdot \prod_{j=i}^{\ell-1} \rho_{u_j}) - 1$.

Without loss of generality, we assume that nodes never propose *useless* split-

offers to their children, that is, ρ -split where $\rho > \rho_1 := 1 - 1/r^*$, since their children would not have incentive to play even if they possessed the answer themselves. Also, for simplicity we assume discrete domain and range for the strategy f_u of every node u , that is, $f_u : \mathcal{D}_M \rightarrow (\mathcal{D}_M \cup \perp)$, where $f_u(\rho) = \perp$ indicates that u chooses not to propagate the query, and $\mathcal{D}_M = \{\frac{\rho_1}{M}, \frac{2\rho_1}{M}, \dots, \frac{(M-1)\rho_1}{M}, \rho_1\}$ is a discretization of the interval $(0, \rho_1]$.

6.2.2 Propagation of the payment

We remark that the above payoffs for the path u_0, \dots, u_ℓ will turn into concrete payments only if the root selects u_ℓ among the answer-holders. Indeed, among all answer-holders reached by the query the root, will select only one for payment. In the fixed-payment model of [129, 15], this selection is made using a fixed arbitrary procedure that does not affect the strategies of the nodes (e.g., performing a random walk from the root descending down the tree; the first hit answer-holder will be paid along with its ancestors). In their setting, this choice is coherent as the root always spends a fixed investment, no matter how deep in the tree the payment is propagated. In our case this peculiarity is missing as a result of the split contract mechanism. In our model we will assume the root selects for payment one among the answer-holders (reached by the query) at *smallest depth*. This is motivated by different facts. First, if we consider some notion of time related to propagating the query one level down, then our selection mechanism better depicts the strategy adopted in the DARPA Network Challenge, where the payment was given to the first participant reporting the correct location of a balloon. Second, the actual investment of the root is in general smaller if the path to the answer is shorter. Finally, a selection mechanism based on smallest depth alleviates the false-name issue discussed in [177]. In case of multiple answer-holders at smallest depth, we assume that the root breaks ties in a way that does not affect the strategies of the nodes (e.g., performing a random walk from the root to one of the leaves of the subtree formed by all

shortest paths to the answers, and selecting the corresponding answer-holder).

6.2.3 Difference with respect to previous work

We would like to spend a few words highlighting some of the main differences between our analysis and those in [129, 15]. One of these differences, the propagation of payments, has been already discussed above; from the technical point of view, the *smallest depth* selection mechanism introduces the hurdle that the strategy of each node does not only depend on the strategies in its subtree (as in the case of [129, 15]), but potentially on those of all nodes. We remark that the gap in efficiency of the two models is not related to the different propagation of payment. In fact, if the answer-holder were to be selected according to the smallest depth mechanism in the fixed-payment setting, then the investment needed to retrieve the answer would increase. Roughly speaking, this happens as a node further down in the tree requires higher reward to forward the query, in order to compensate for the smaller probability of having a payment candidate in its subtree.

Another salient difference between the two models concerns the *values* of the contracts: while the nature of the fixed-payment contracts of [129, 15] implies that a node being offered a reward of r can only offer an amount $r' < r$ to its children, we do not enjoy this property on the ρ 's in the case of split contracts. This unfortunately precludes the inductive arguments adopted in [129, 15], making a more involved analysis necessary.

We conclude this section discussing about the gap in efficiency between split contracts, for which an investment proportional to the depth of the search tree suffices for any branching factor $b > 1$, and the results in [129, 15], for which the investment becomes exponential in the depth of the search tree when the branching factor drops below 2. In the setting of [129], the additional amount of reward δ_j that the root needs in

order to explore j levels of the tree (rather than stopping at level $j - 1$) can be expressed as

$$\delta_{j+1} = \frac{1 - \phi_{j-1}}{\lambda_j} \delta_j + 1.$$

When the branching factor drops below 2, the ratio $\frac{1 - \phi_{j-1}}{\lambda_j}$ is greater than 1, and the investment needed at the root to propagate the query down to depth $h_\Psi(\varepsilon, n)$ becomes exponential in $\log n$ (hence, $\text{poly}(n)$).

In contrast, the dependency on the ratio $\frac{1 - \phi_{j-1}}{\lambda_j}$ is softer in our setting. In the proof of Theorem 9, we show that the ρ -split a node at level ℓ needs to receive in order to propagate the query i levels down its subtree is

$$\rho_i^{(\ell)} = 1 - \frac{1}{r^* - i(1 + O(\frac{1 - \phi_{i-1}}{\lambda_i}))}.$$

Since we can show that $\frac{1 - \phi_{i-1}}{\lambda_i}$ is bounded by a constant for any branching process with $b > 1$, an investment $r^* = O(h_\Psi(\varepsilon, n)) = O(\log n)$ suffices for the value $\rho_{h_\Psi(\varepsilon, n)}^{(1)}$ offered by the root to its children to be well-defined (i.e., in \mathcal{D}_M), and hence for the answer to be retrieved cheaply.

6.2.4 Roadmap

The rest of the paper is structured as follows. In Section 6.3, we derive properties that hold for any Nash equilibrium. In Section 6.4, we develop a condition that we call h -consistency under which we can show that a set of strategies \mathbf{g} for the nodes propagates the query to the desired level and is substantially the unique Nash equilibrium. In Section 6.5, we derive a bound on the investment r^* , depending on quantities related to the branching process, for which h -consistency is guaranteed to hold. Finally, in Section 6.6, we study such quantities of the branching process to conclude that $r^* = O(h_\Psi(\varepsilon, n)) = O(\log n)$. Due to space constraints, all proofs are deferred to Section 6.7.

6.3 Properties of Nash Equilibria

In this section we present the notion of Nash equilibrium that naturally arises in the context of split-contracts, and we then derive a manageable expression that any Nash equilibrium has to maximize. Let f_v be the function representing the strategy of node v , and \mathbf{f} be the set of strategies of all nodes up to level $h_\Psi(\varepsilon, n)$, as we assumed that nodes in lower levels do not play.

Definition 11 (Nash equilibrium) *Let r^* , Ψ , ε , n be the parameters of the model, and \mathbf{f} be a set of functions for all nodes up to level $h_\Psi(\varepsilon, n)$. For any such node v , let ρ^v be the split contract offered to v by its parent under \mathbf{f} . Then, \mathbf{f} is a Nash equilibrium if, for each node v , v does not increase its expected payoff by deviating from $f_v(\rho^v)$ when all other nodes play according to \mathbf{f} . The expectation is taken over the randomness of the branching process and of the process assigning answers to nodes.*

We now give a few definitions that will be useful to derive properties of any Nash equilibrium. Given a realization of the branching process, we say that a node v at level $\ell \leq h_\Psi(\varepsilon, n)$ is *active* if the branching process reaches v . Moreover, given a set \mathbf{f} of strategies and a realization of the branching process, we say that an active node v is *\mathbf{f} -reachable* if \mathbf{f} forwards the query down to v . Given a realization of the branching process and of the process assigning the answer to nodes, we say that an \mathbf{f} -reachable node v at level ℓ is an *\mathbf{f} -candidate* if v holds the answer and no \mathbf{f} -reachable node at a level $\ell' < \ell$ does. Observe that the root selects for payment one among the \mathbf{f} -candidates. For each node v at level $\ell \leq h_\Psi(\varepsilon, n)$, set \mathbf{f} of strategies, and $j \geq 1$, let $\alpha_v^{\mathbf{f}}(j|\rho)$ be the probability that there is an \mathbf{f} -candidate in v 's subtree at distance j from v , conditional on v being \mathbf{f} -reachable and offering a ρ -split to its children. Similarly, for $j \geq 1$, let $\beta_v^{\mathbf{f}}(j|\rho)$ be the expected payment that v receives from its children given that v offers a ρ -split to its children and there is an \mathbf{f} -candidate in v 's subtree at distance j from v to whom the

root propagates the payment.

The following lemma characterizes an expression that must be maximized by every node up to level $h_\Psi(\varepsilon, n)$ in any Nash Equilibrium.

Lemma 16 *Consider any set \mathbf{f} of strategies, and let ρ^v be the split contract offered to v by its parent under \mathbf{f} . Then, \mathbf{f} is a Nash equilibrium if and only if, for every node v up to level $h_\Psi(\varepsilon, n)$, $f_v(\rho^v)$ is a value of ρ maximizing the function*

$$\chi_v^{\mathbf{f}}(\rho; \rho^v) := \sum_{j \geq 1} \alpha_v^{\mathbf{f}}(j|\rho) \left((1 - \rho^v) \beta_v^{\mathbf{f}}(j|\rho) - 1 \right). \quad (6.1)$$

To break ties in case of multiple maxima for $\chi_v^{\mathbf{f}}(\cdot; \rho^v)$, we make the same assumption as in [129, 15] that nodes favor strategies that forward the query further down in the tree. Using Lemma 16, we can now prove that Nash equilibria are “leveled”.

Lemma 17 *Consider any Nash equilibrium \mathbf{f} . Then for each active node v at level ℓ , v is \mathbf{f} -reachable if and only if every active node at level ℓ is.*

By means of Lemma 17, we will say that a Nash equilibrium \mathbf{f} is k -tall if level k is \mathbf{f} -reachable and level $k + 1$ is not. This notion is useful in decoupling the probabilities $\alpha_v^{\mathbf{f}}(j|f_v(\rho))$ from the particular equilibrium \mathbf{f} and node v . To see how, assume \mathbf{f} is k -tall, $k \leq h_\Psi(\varepsilon, n)$. For any node v at level $\ell \leq k$ and any $j \leq k - \ell$, let $\gamma_j^{\langle \ell \rangle}$ be the probability that there exists an \mathbf{f} -candidate in v 's subtree at distance j from v (and therefore there is no \mathbf{f} -candidate in the first $\ell + j - 1$ levels). Then, as \mathbf{f} is k -tall, we have that for any node at level ℓ , $\gamma_j^{\langle \ell \rangle}$ depends only on ℓ and j (and not on \mathbf{f} or the specific node). This observation directly yields the following result relating the probabilities $\alpha_v^{\mathbf{f}}(j|f_v(\rho^v))$ and $\gamma_j^{\langle \ell \rangle}$.

Lemma 18 *Let \mathbf{f} be a k -tall Nash equilibrium, $k \leq h_\Psi(\varepsilon, n)$. Then, for every $\ell \leq k$ and*

node v at level ℓ ,

$$\alpha_v^{\mathbf{f}}(j|f_v(\rho^v)) = \begin{cases} \gamma_j^{(\ell)}, & \text{for } 1 \leq j \leq k - \ell \\ 0, & \text{for } j > k - \ell \end{cases}$$

In general, if the query is forwarded j levels down v 's subtree when v offers a ρ -split to its children, then we have $\alpha_v^{\mathbf{f}}(j|\rho) = \gamma_j^{(\ell)}$.

6.4 The Nash Equilibrium

In this section, we derive conditions for the existence of a Nash equilibrium that forwards the query down to level $h_{\Psi}(\varepsilon, n)$, or, equivalently, retrieves the answer with the desired probability $1 - \zeta - \varepsilon$. For ease of notation, let $h = h_{\Psi}(\varepsilon, n)$. We proceed as follows. First we define the functions $e_i^{(\ell)}$ and the thresholds $\rho_i^{(\ell)}$, which intuitively represent expected rewards and contracts for a special set of strategies \mathbf{g} . However, to define \mathbf{g} , we will need all $\rho_i^{(\ell)}$ to exist and be decreasing in i , for all $\ell \leq h$, property that we will dub h -consistency. Finally, assuming h -consistency, we will show that \mathbf{g} forwards the query to level h and is a Nash equilibrium (in fact with an extra property, we will say \mathbf{g} is a best-interest Nash Equilibrium).

We begin by defining the aforementioned functions and values. We provide an inductive process which defines, for each $1 \leq \ell \leq h$, a sequence of functions $e_i^{(\ell)} : [0, 1] \rightarrow \mathbb{R}$, $0 \leq i \leq h - \ell$, and values $\rho_i^{(\ell)} \in \mathcal{D}_M$, $1 \leq i \leq h - \ell + 1$. For every $0 \leq \ell \leq h$, set $e_0^{(\ell)}(\rho) = 0$ and $\rho_1^{(\ell)} = \rho_1 = 1 - 1/r^*$. Suppose that all $\rho_i^{(\ell')}$ have been defined for $\ell < \ell' \leq h$ and $1 \leq i \leq h - \ell' + 1$. Then, for all $1 \leq i \leq h - \ell$, the function $e_i^{(\ell)}(\rho)$ is defined as

$$e_i^{(\ell)}(\rho) = \sum_{j=1}^i \gamma_j^{(\ell)} \left[(1 - \rho)r^* \left(\prod_{t=0}^{j-1} \rho_{i-t}^{(\ell+t+1)} \right) - 1 \right].$$

Having defined $e_i^{(\ell)}(\rho)$, we define

$$\rho_{i+1}^{(\ell)} = \max\{\rho \in \mathcal{D}_M : e_i^{(\ell)}(\rho) \geq e_{i-1}^{(\ell)}(\rho)\},$$

if such value exists, and leave $\rho_{i+1}^{(\ell)}$ undefined otherwise.

For a node v at level $\ell \leq h$, $e_i^{(\ell)}$ has the intuitive meaning of the expected reward that v receives from its children when the query is propagated i levels down v 's subtree (assuming the other nodes play accordingly). The value $\rho_{i+1}^{(\ell)}$ represents the ‘‘cheapest’’ split to offer a node v at level $\ell \leq h$ so that v prefers to propagate the query i levels down its subtree rather than $i - 1$ (recall that, to break ties, we assumed that nodes prefer to propagate the query further down the tree). To guarantee the propagation of the query to level h , we will need the values $\rho_{h-\ell}^{(\ell)}$ to be defined.

Definition 12 (h -consistency) *We say that h -consistency holds if, for all $1 \leq \ell \leq h$ and $2 \leq i \leq h - \ell + 1$, the value $\rho_i^{(\ell)}$ is defined and $\rho_i^{(\ell)} < \rho_{i-1}^{(\ell)}$ (note that $\rho_1^{(\ell)}$ is always defined).*

Intuitively, the ordering of the values $\rho_i^{(\ell)}$ in the definition of h -consistency states that if a node v propagates the query i levels down its subtree when offered a ρ -split by its father, then, in order to propagate the query $i + 1$ levels down, it must be that v is offered a split not greater than ρ . This property is at the basis of the following definition of the set of strategies \mathbf{g} , which we will then show to be a Nash equilibrium. Note how, under \mathbf{g} , nodes at the same level play the same strategy.

Definition 13 (Strategy \mathbf{g}) *Assume h -consistency holds. For each $1 \leq \ell \leq h$, consider the function $t^{(\ell)}(\rho) : [0, 1] \rightarrow \mathcal{D}_M \cup \{\perp\}$ defined by $t^{(\ell)}(\rho) = \rho_{i-1}^{(\ell+1)}$ for the unique i such that $\rho_{i+1}^{(\ell)} < \rho \leq \rho_i^{(\ell)}$ (such i exists under h -consistency), where we assume $\rho_{h-\ell+2}^{(\ell)} = 0$*

and $\rho_0^{\langle \ell+1 \rangle} = \perp$. The set of strategies \mathbf{g} is defined by setting $g_v(\rho) = t^{\langle \ell \rangle}(\rho)$ to every node v at level ℓ , for each $1 \leq \ell \leq h$, and letting the root play $\rho_h^{\langle 1 \rangle}$.

It follows that, under \mathbf{g} , the root (at level zero) offers a $\rho_h^{\langle 1 \rangle}$ -split to its children, who in turn offer $\rho_{h-1}^{\langle 2 \rangle}$ -split contracts to their own children, and so on, until the nodes at level h , who do not forward the query (they play $t^{\langle h \rangle}(\rho_1^{\langle h \rangle}) = \perp$). Observe that \mathbf{g} is h -tall, as all nodes up to level h are \mathbf{g} -reachable. The following theorem states that \mathbf{g} is a Nash equilibrium.

Theorem 7 (Nash equilibrium) *Assuming h -consistency, the set of strategies \mathbf{g} is a Nash equilibrium.*

The key fact in the proof is to show that, for any node v at level $1 \leq \ell \leq h$ (which under \mathbf{g} receives a $\rho_{h-\ell+1}^{\langle \ell \rangle}$ -split from its parent and in turn offers a $\rho_{h-\ell}^{\langle \ell+1 \rangle}$ -split to its children), $\chi_v^{\mathbf{g}}(\rho_j^{\langle \ell+1 \rangle}; \rho) = e_j^{\langle \ell \rangle}(\rho)$ for all $j \leq h - \ell$, and that $\rho_{h-\ell}^{\langle \ell+1 \rangle}$ is the only maximizer of (6.1) that propagates the query to level h . Then the theorem follows by Lemma 16.

Even though \mathbf{g} is not the only Nash equilibrium, the proof of Theorem 7 shows that \mathbf{g} enjoys the additional property that, for each node v and $\rho \in \mathcal{D}_M$,

$$g_v(\rho) = \arg \max_{\rho'} \{\chi_v^{\mathbf{g}}(\rho'; \rho)\}.$$

We call any equilibrium enjoying such property a *best-interest* equilibrium, as nodes choose their best option in any scenario. The following theorem shows that \mathbf{g} is substantially the only best-interest equilibrium, meaning that every other best-interest equilibrium \mathbf{f} coincides with \mathbf{g} on all the split-offers that are actually offered to nodes under \mathbf{f} . As a remark, we observe that even the game with fixed-payment contracts in [129, 15] admits multiple equilibria, although the authors claim uniqueness (a counter-example is presented in Section 6.7.8). On the positive side, the equilibrium analyzed in [129, 15] is the unique best-interest Nash equilibrium of their game.

Theorem 8 (Uniqueness) *Assume h -consistency. Let \mathbf{f} be any ℓ -tall best-interest Nash equilibrium, for some $1 \leq \ell \leq h$, and, for each node v up to level ℓ , let ρ^v be the split contract offered to v by its parent under \mathbf{f} . Then, for every node v up to level ℓ , $f_v(\rho^v) = g_v(\rho^v)$.*

Theorem 8 implies that every best-interest equilibrium \mathbf{f} in which the root offers a $\rho_h^{(1)}$ -split to its children has to be h -tall, as \mathbf{f} agrees with \mathbf{g} on all split-offers made in \mathbf{g} . As h -tall equilibria retrieve the answer with the desired probability, the root has incentive to play $\rho_h^{(1)}$ as its strategy and would have incentive to deviate to $\rho_h^{(1)}$ if playing a different strategy. The following result is then implied.

Corollary 4 *Under the assumption of h -consistency, all best-interest equilibria retrieve the answer with the desired probability.*

6.5 Guaranteeing h -consistency

Until now, we assumed h -consistency both in the definition of \mathbf{g} and in the proof that \mathbf{g} is a Nash equilibrium. It therefore remains to derive conditions that ensure h -consistency. In the following theorem, we provide a lower bound on the reward r^* above which h -consistency is guaranteed. The bound reads in terms of the probabilities $\gamma_i^{(\ell)}$ through the quantities $\Gamma_i^{(\ell)} = \frac{1}{\gamma_i^{(\ell)}} \sum_{j=1}^{i-1} \gamma_j^{(\ell)}$, which, for all $1 \leq \ell \leq h$ and $1 \leq i < h - \ell$, intuitively represent the ratio between the probability that a node at level ℓ has a candidate at depth $j < i$ in its subtree versus the probability that it has one at depth i .

Theorem 9 *Suppose the discretization parameter M is large enough, say $M = \Theta(r^{*2})$, and that*

$$r^* \geq 4 \cdot h \cdot \max \left\{ 1, \max_{\substack{1 \leq \ell \leq h \\ 1 \leq i < h - \ell}} \Gamma_i^{(\ell)} \right\}. \quad (6.2)$$

Then h -consistency holds. In particular, for all $1 \leq \ell \leq h$ and $1 \leq i \leq h - \ell$, $\rho_i^{(\ell)}$ is defined and satisfies

$$1 - \frac{1}{r^* - i} < \rho_i^{(\ell)} \leq 1 - \frac{1}{r^* - i + 1}. \quad (6.3)$$

The main idea to prove the theorem is to derive tight upper and lower bounds on $\rho_{i+1}^{(\ell-1)}$ and then proceed by induction on both ℓ and i . It can also be proven that, for a fixed i , $\rho_i^{(\ell)}$ is decreasing in ℓ for $0 \leq \ell \leq h - i$. The intuition for this property is that a node further down in the tree is willing to give a smaller fraction of its reward back to its parent, in order to compensate the smaller probability of having a candidate in its subtree. However, we do not need this property to ensure h -consistency.

Theorem 9 along with Corollary 4 directly yields the following pivotal result, which relates the quantities $\Gamma_j^{(\ell)}$ to the investment that is sufficient at the root to retrieve the answer with the desired probability.

Corollary 5 *Suppose condition (6.2) holds. Then, in any best-interest Nash equilibrium, the query reaches all nodes at level $h = h_{\Psi}(\varepsilon, n)$ of the tree. That is, an answer is retrieved with probability at least $1 - \zeta - \varepsilon$.*

6.6 Efficiency

In the previous section, we derived a lower bound on the investment r^* as a function of the values $\Gamma_i^{(\ell)}$, for $1 \leq \ell \leq h$ and $1 \leq i \leq h - \ell$. In this section, we show our main result by relating these values to the branching process and the desired success probability. The following lemma bounds these quantities in terms of the probabilities λ_i and ϕ_i of the branching process. Recall that, for each $i \geq 0$ we defined ϕ_i as the probability that no node at level $j \leq i$ possesses the answer, and $\lambda_i = \phi_{i-1} - \phi_i$ as the probability that a node at level i possesses the answer and no node at level $j < i$ does.

Lemma 19 For every $1 \leq \ell \leq h$ and $1 \leq i \leq h - \ell$, it holds that $\Gamma_i^{(\ell)} \leq \frac{1}{\phi_{\ell+i-1}} \frac{1 - \phi_{i-1}}{\lambda_i}$.

The key in proving Lemma 19 is to express $\gamma_i^{(\ell)}$ in terms of the probabilities ϕ_j and λ_j defined above, and then to bound $\Gamma_i^{(\ell)}$ exploiting the memory-less property of the branching process and of the process assigning the answer to the nodes.

The following technical lemma provides an upper bound to $\frac{1 - \phi_{i-1}}{\lambda_i}$. In particular, for any fixed branching process with $b > 1$, this ratio is bounded by a constant, as long as ϕ_i is bounded away from the extinction probability ζ . The lemma characterizes the bound with respect to the branching process and the gap $\phi_i - \zeta$, and its proof builds on the mathematical properties of the probability generating function of the branching process. Recall that the desired success probability is $1 - \zeta - \varepsilon$.

Lemma 20 Consider any Galton-Watson branching process with branching factor $b > 1$. Then, for every i such that $\zeta + \varepsilon \leq \phi_i \leq 1$, it holds that

$$\frac{1 - \phi_i}{\lambda_{i+1}} \leq \max \left\{ \frac{1}{b-1}, \frac{1}{\varepsilon} \cdot \frac{1}{1 - \Psi'(\zeta)} \right\}.$$

Our main result directly follows by combining Corollary 5, Lemma 19 and Lemma 20, along with the observation that $\phi_{\ell+i-1} > \varepsilon$ (as $\phi_{\ell+i-1} \geq \phi_{h_{\Psi}(\varepsilon, n)-1} > \varepsilon$). For the case of a ray, the bound can be obtained observing that $\phi_i = (1 - 1/n)^i$ and $\lambda_{i+1} = \frac{\phi_i}{n}$, which implies $\Gamma_i^{(\ell)} \leq \Gamma_h^{(1)} \leq \varepsilon^{-2}n$.

Theorem 10 (Efficiency) Consider any Galton-Watson branching process with $b > 1$. Then, the root retrieves the answer with probability at least $\sigma = 1 - \zeta - \varepsilon$ provided an investment of

$$r^* = \frac{4}{\varepsilon} \cdot \max \left\{ \frac{1}{b-1}, \frac{1}{\varepsilon} \cdot \frac{1}{1 - \Psi'(\zeta)} \right\} \cdot h_{\Psi}(\varepsilon, n).$$

In the case of a ray, with $b = 1$ and $c_0 = \zeta = 0$, an investment of $r^* = 4 \cdot \frac{n}{\varepsilon^2} \cdot h_{\Psi}(\varepsilon, n) = 4 \cdot \frac{n^2}{\varepsilon^2} \ln \frac{1}{\varepsilon}$ suffices.

Observe that an investment of $h_\Psi(\varepsilon, n)$ is necessary even in a centralized (non-strategic) setting, where the root decides the strategies of all nodes while only guaranteeing a non-negative payoff to them (each node pays a unit cost when returning the answer). In line with intuition, the investment grows as b tends to 1 (in the limit, when the tree becomes a ray, the investment is polynomial in n), and when the accuracy ε approaches 0. The term $\frac{1}{1-\Psi'(\zeta)}$ can be crudely bounded by $\frac{1}{c_0}$. However, when c_0 tends to zero, so does the extinction probability ζ , which implies $\frac{1}{1-\Psi'(\zeta)} \approx \frac{1}{1-c_1}$, also suggesting a more expensive investment when the tree tends to a ray (i.e., when c_1 approaches 1).

6.7 Proofs

6.7.1 Proof of Lemma 16

Fix the available investment r^* , a set \mathbf{f} of strategies, and a node v at level $\ell \leq h_\Psi(\varepsilon, n)$. We need to define the following events. Let A denote the event that the root propagates the payment down through v , that is, the root selects for payment an \mathbf{f} -candidate in v 's subtree. For each $0 \leq j \leq h_\Psi(\varepsilon, n) - \ell$, let B_j denote the event that the \mathbf{f} -candidates are at level $\ell + j$. Finally let C denote the event that v is \mathbf{f} -reachable and D denote the event that there is an \mathbf{f} -candidate in v 's subtree. Observe that the co-occurrence of B_0 and D means that v itself is an \mathbf{f} -candidate. Given r^* and \mathbf{f} , let $Y_{\mathbf{f}, r^*}^v$ be the random variable denoting the payment assigned to v .

We have that

$$\begin{aligned} E[Y_{\mathbf{f}, r^*}^v] &= \sum_{j \geq 0} E[Y_{\mathbf{f}, r^*}^v | A, B_j, D, C] \Pr(A, B_j, D, C) \\ &= \Pr(A|D) \Pr(C) \sum_{j \geq 0} E[Y_{\mathbf{f}, r^*}^v | A, B_j, D, C] \Pr(B_j, D|C). \end{aligned}$$

The first equality follows from the law of total probability together with the obser-

vation that $\Pr(A, \overline{C}) = \Pr(A, \overline{D}) = 0$ and $E[Y_{\mathbf{f}, r^*}^v | \overline{A}] = 0$. The second equality follows from the chain rule of probability and the fact that $\Pr(A | B_j, D, C) = \Pr(A | D)$ for all $j \geq 0$. Observe that the term corresponding to $j = 0$ (i.e., v is the \mathbf{f} -candidate selected for payment) does not depend on f_v since $E[Y_{\mathbf{f}, r^*}^v | A, B_0, D, C] = (1 - \rho^v)r^* - 1$ and $\Pr(B_0, D | C)$ only depends on the strategies of the nodes that are ancestors of v . Similarly, f_v affects neither $\Pr(C)$, which depends on the strategies of v 's ancestors only, nor $\Pr(A | D)$, which is only based on the root's choice of whom to propagate the payment to. Finally, note that if v offers a ρ -split to its children, then, for $j \geq 1$, $E[Y_{\mathbf{f}, r^*}^v | A, B_j, S, C] = (1 - \rho^v)\beta_v^{\mathbf{f}}(j | \rho) - 1$ and $\Pr(B_j, D | C) = \alpha_v^{\mathbf{f}}(j | \rho)$. Therefore, \mathbf{f} is a Nash equilibrium if and only if, for every node v up to level $h_{\Psi}(\varepsilon, n)$, $f_v(\rho^v)$ is a value ρ maximizing $\chi_v^{\mathbf{f}}(\rho; \rho^v)$.

6.7.2 Proof of Lemma 17

Let \mathbf{f} be a Nash equilibrium. Fix a node v and let

$$\rho_2 = \max\{\rho \in \mathcal{D}_M : \chi_v^{\mathbf{f}}(\rho_1; \rho) \geq \chi_v^{\mathbf{f}}(\perp; \rho)\}$$

be the maximum split v 's father can ask v so that v will in turn prefer to offer a ρ_1 -split to their children rather than just participating in the game without propagating the query. We first argue that ρ_2 does not depend on the chosen node v , then show that $f_v(\rho) = \perp$ for every node v and $\rho_2 < \rho \leq \rho_1 = 1 - \frac{1}{r^*}$, and finally use this fact to prove the lemma.

To see that ρ_2 does not depend on v , observe that $\chi_v^{\mathbf{f}}(\perp; \rho) = 0$ as $\alpha_v^{\mathbf{f}}(j | \perp) = 0$ for $j \geq 1$, and that $\chi_v^{\mathbf{f}}(\rho_1; \rho) = \alpha_v^{\mathbf{f}}(1 | \rho_1)((1 - \rho)\beta_v^{\mathbf{f}}(1 | \rho_1) - 1) = \alpha_v^{\mathbf{f}}(1 | \rho_1)((1 - \rho)r^*\rho_1 - 1)$.

We now show that $f_v(\rho) = \perp$, for every node v and $\rho_2 < \rho \leq \rho_1$. By contradiction, suppose $f_v(\rho) = \rho'$, for some v , $\rho_2 < \rho \leq \rho_1$, $\rho' \in \mathcal{D}_M$. On the one hand, as \mathbf{f} is a Nash equilibrium, Lemma 16 implies that ρ' maximizes $\chi_v^{\mathbf{f}}(\rho'; \rho)$, and thus

$$\chi_v^{\mathbf{f}}(\rho'; \rho) \geq \chi_v^{\mathbf{f}}(\perp; \rho) = 0.$$

On the other hand, we have that

$$\begin{aligned}\chi_v^{\mathbf{f}}(\rho'; \rho) &= \sum_{j \geq 1} \alpha_v^{\mathbf{f}}(j|\rho') \left((1-\rho)\beta_v^{\mathbf{f}}(j|\rho') - 1 \right) \\ &\leq \sum_{j \geq 1} \alpha_v^{\mathbf{f}}(j|\rho') ((1-\rho)r^*\rho_1 - 1),\end{aligned}$$

where the last inequality follows from $\beta_v^{\mathbf{f}}(j|\rho') \leq r^*\rho_1$ for all $j \geq 1$, as ρ' must be at most ρ_1 for v 's children to participate to the game. By definition of ρ_2 , it must be that $\chi_v^{\mathbf{f}}(\rho_1; \rho) < \chi_v^{\mathbf{f}}(\perp; \rho) = 0$, which implies $((1-\rho)r^*\rho_1 - 1) < 0$ and thus $\chi_v^{\mathbf{f}}(\rho'; \rho) < 0$, generating a contradiction.

We are now ready to prove the lemma. By contradiction, suppose the statement of the lemma does not hold. Then there must be two sibling nodes u and v (at some level $\ell < h_{\Psi}(\varepsilon, n)$) and a value $\rho = \rho^u = \rho^v$ such that $f_u(\rho) = \rho' \neq \perp$ and $f_v(\rho) = \perp$, that is, such that u forwards the query when offered a ρ -split by its parent while v does not. By the claim above, $f_u(\rho) = \rho'$ implies that $\rho \leq \rho_2$ and therefore, by definition of ρ_2 , v would have incentive to deviate from f_v , offering a ρ_1 -split to their children than just participating to the game without propagating the query, contradicting that \mathbf{f} is a Nash equilibrium.

6.7.3 Proof of Theorem 7

Under h -consistency, for all $\ell \leq h$ and $2 \leq i \leq h - \ell + 1$, $\rho_i^{\langle \ell \rangle}$ is defined and $\rho_i^{\langle \ell \rangle} < \rho_{i-1}^{\langle \ell \rangle}$ (recall that $\rho_1^{\langle \ell \rangle}$ is defined for all $\ell \leq h$). In the proof of the theorem, we make use of the following technical lemma, that is a consequence of h -consistency.

Claim 3 *Assume h -consistency. Then, for every $1 \leq \ell \leq h$, $1 \leq i \leq h - \ell$, and $\rho_{i+2}^{\langle \ell \rangle} < \rho \leq \rho_{i+1}^{\langle \ell \rangle}$, we have that*

$$e_i^{\langle \ell \rangle}(\rho) > e_{i+1}^{\langle \ell \rangle}(\rho) > \cdots > e_{h-\ell}^{\langle \ell \rangle}(\rho),$$

and

$$e_i^{(\ell)}(\rho) \geq e_{i-1}^{(\ell)}(\rho) \geq \dots \geq e_0^{(\ell)}(\rho),$$

where we assume $\rho_{h+1}^{(1)} = 0$.

Proof. Consider any $i+1 \leq j \leq h-\ell$, and observe that, by definition,

$$\rho_{j+1}^{(\ell)} = \max\{\rho' \in \mathcal{D}_M : e_j^{(\ell)}(\rho') \geq e_{j-1}^{(\ell)}(\rho')\},$$

and, by h -consistency (as $j+1 \geq i+2$), $\rho_{j+1}^{(\ell)} \leq \rho_{i+2}^{(\ell)} < \rho$. It follows that $e_j^{(\ell)}(\rho) < e_{j-1}^{(\ell)}(\rho)$ for all $i+1 \leq j \leq h-\ell$, which implies that

$$e_{h-\ell}^{(\ell)}(\rho) < e_{h-\ell-1}^{(\ell)}(\rho) < \dots < e_i^{(\ell)}(\rho),$$

proving the first chain of inequalities in the lemma. Now consider any $2 \leq m \leq i+1$, and observe that, by definition of $\rho_m^{(\ell)}$,

$$e_{m-1}^{(\ell)}(\rho_m^{(\ell)}) \geq e_{m-2}^{(\ell)}(\rho_m^{(\ell)})$$

and, by h -consistency (as $i+1 \geq m$), $\rho \leq \rho_m^{(\ell)}$. This implies that $e_{m-1}^{(\ell)}(\rho) \geq e_{m-2}^{(\ell)}(\rho)$ for all $2 \leq m \leq i+1$. It follows that

$$e_i^{(\ell)}(\rho) \geq e_{i-1}^{(\ell)}(\rho) \dots \geq e_0^{(\ell)}(\rho),$$

which proves the second chain of inequalities in the lemma. \square

To show that \mathbf{g} is a Nash equilibrium, by Lemma 16, it suffices to prove that, for every node v at level up to h , $g_v(\rho^v)$ is the value that maximizes $\chi_v^{\mathbf{g}}(\cdot; \rho^v)$, where ρ^v is the split offer v receives from its parent. Let $0 \leq i \leq h-1$, and fix a node v at level $\ell = h-i$. Under \mathbf{g} , v receives a $\rho_{i+1}^{(\ell)}$ -split from its parent and in turn offers a

$t^{(\ell)}(\rho_{i+1}^{(\ell)}) = \rho_i^{(\ell+1)}$ -split to its children. Therefore, it suffices to show that

$$\rho_i^{(\ell+1)} = \arg \max_{\rho'} \{\chi_v^{\mathbf{g}}(\rho'; \rho_{i+1}^{(\ell)})\}.$$

We will in fact prove something stronger, that is, for all $\rho \in \mathcal{D}_M$,

$$g_v(\rho) = t^{(\ell)}(\rho) = \arg \max_{\rho'} \{\chi_v^{\mathbf{g}}(\rho'; \rho)\}. \quad (6.4)$$

Fix any $\rho \in \mathcal{D}_M$. A few observations allow to prove condition (6.4) for the chosen ρ . First, by h -consistency, there exists unique k such that $\rho_{k+2}^{(\ell)} < \rho \leq \rho_{k+1}^{(\ell)}$, where we assume $\rho_{h+1}^{(\ell)} = 0$. Second, by definition of \mathbf{g} and $\chi_v^{\mathbf{g}}(\cdot; \cdot)$, node v has an incentive to play a given $\rho' \in \mathcal{D}_M$ only if there is no $\hat{\rho} > \rho'$ such that v 's children would play exactly the same split contract if either offered a $\hat{\rho}$ -split or a ρ' -split (otherwise, the query would propagate the same number of levels down the tree, but with v earning more if offering a $\hat{\rho}$ -split to its children). This implies that if ρ' maximizes $\chi_v^{\mathbf{g}}(\cdot, \rho)$, then $\rho' = \rho_j^{(\ell+1)}$ for some $0 \leq j \leq i$ (recall that node v is at level $h - i$). Third, by definition of \mathbf{g} , $e_j^{(\ell)}(\cdot)$ and $\chi_v^{\mathbf{g}}(\cdot; \cdot)$, and by Claim 18, we have that $\chi_v^{\mathbf{g}}(\rho_j^{(\ell+1)}; \rho) = e_j^{(\ell)}(\rho)$ for all $0 \leq j \leq h - \ell$. Finally, by Lemma 3, as $\rho_{k+2}^{(\ell)} < \rho \leq \rho_{k+1}^{(\ell)}$, we have that $e_k^{(\ell)}(\rho) > e_j^{(\ell)}(\rho)$ for all $k < j < h - \ell$ and $e_k^{(\ell)}(\rho) \geq e_j^{(\ell)}(\rho)$ for all $0 \leq j < k$. We have that $\rho_k^{(\ell+1)}$ is the only maximizers of $\chi_v^{\mathbf{g}}(\cdot; \rho)$ which forwards the query to level h , while any other maximizer forwards the query to some level $\ell' < h$. Therefore, as we assumed that nodes break ties preferring to propagate the query further down the tree, $\rho_k^{(\ell+1)} = t^{(\ell)}(\rho)$ is the preferred strategy of node v when offered a ρ -split from its parent. Considering all $\rho \in \mathcal{D}_M$, condition (6.4) follows and the theorem is proven.

6.7.4 Proof of Theorem 8

Under h -consistency, for all $\ell \leq h$, $\rho_i^{(\ell)}$ is defined and $\rho_i^{(\ell)} < \rho_{i-1}^{(\ell)}$ for all $2 \leq i \leq h - \ell + 1$. Let \mathbf{f} be a best-interest Nash equilibrium that is ℓ -tall for some $\ell \leq h$. As \mathbf{f} is best-interest, for every node v up to level ℓ ,

$$f_v(\rho) = \arg \max_{\rho'} \{\chi_v^{\mathbf{f}}(\rho'; \rho)\}, \quad \forall \rho \in \mathcal{D}_M.$$

We prove a stronger claim than the one in the theorem, that is, for every node v up to level ℓ ,

$$f_v(\rho) = g_v(\rho), \quad \forall \rho^v \leq \rho \leq \rho_1, \quad (6.5)$$

where ρ^v is the split offered to v by its parent under \mathbf{f} , $\rho_1 = 1 - 1/r^*$ and \mathbf{g} is the best-interest Nash equilibrium from Definition 13.

We proceed by induction on the levels of the tree, starting from level ℓ and going backwards. In particular we prove by induction that (6.5) holds for every node at level ℓ , for every level $\ell' \leq \ell$. Consider any node v at level ℓ . As \mathbf{f} is ℓ -tall (i.e., level ℓ is \mathbf{f} -reachable, while level $\ell + 1$ is not), node v plays \perp . Therefore, v 's parent (at level $\ell - 1$) has incentive to offer v a ρ_1 -split (the maximum split such that v has incentive to forward the answer to its parent). It follows that $\rho^v = \rho_1$ and $f_v(\rho_1) = \perp = g_v(\rho_1)$, and (6.5) holds for level ℓ .

Fix $0 \leq i < \ell$, and suppose (6.5) holds for every node at level $\ell - i$. Let $\ell' = \ell - i - 1$, and consider any node v at level ℓ' . In the proof of Theorem 7, we showed that, for every $\rho \in \mathcal{D}_M$ and $1 \leq j \leq i$,

$$\chi_v^{\mathbf{g}}(\rho_j^{(\ell'+1)}; \rho) = e_j^{(\ell')}(\rho).$$

By the inductive hypothesis on level $\ell' + 1 = \ell - i$ and the fact that both \mathbf{f} and \mathbf{g} are

best-interest, we have that, for every $\rho \in \mathcal{D}_M$ and $1 \leq j \leq i$,

$$\chi_v^{\mathbf{f}}(\rho_j^{\langle \ell'+1 \rangle}; \rho) = \chi_v^{\mathbf{g}}(\rho_j^{\langle \ell'+1 \rangle}; \rho).$$

The last two observations imply that, for every $\rho \in \mathcal{D}_M$ and $1 \leq j \leq i$,

$$\chi_v^{\mathbf{f}}(\rho_j^{\langle \ell'+1 \rangle}; \rho) = e_j^{\langle \ell' \rangle}(\rho). \quad (6.6)$$

Lemma 3, together with (6.6), implies that

- (i) for every $j < i$ and $\rho_{j+2}^{\langle \ell' \rangle} < \rho' \leq \rho_{j+1}^{\langle \ell' \rangle}$, node v has incentive to play $\rho_j^{\langle \ell'+1 \rangle}$ among all $\rho_i^{\langle \ell'+1 \rangle} \leq \rho \leq \rho_1$, and
- (ii) for $\rho' = \rho_{i+1}^{\langle \ell' \rangle}$, node v has incentive to play $\rho_i^{\langle \ell'+1 \rangle}$ among all $\rho_i^{\langle \ell'+1 \rangle} \leq \rho \leq \rho_1$.

We need the following technical result in order to proceed with the proof.

Claim 4 *Let v be a node at level $\ell' = \ell - i - 1$. Suppose that v receives a ρ' -split from its parent, with $\rho_{j+2}^{\langle \ell' \rangle} < \rho' \leq \rho_{j+1}^{\langle \ell' \rangle}$ for some $j \leq i$, and that v forwards the query exactly to level $\hat{\ell} \leq \ell$. Moreover, assume that (6.5) holds for every node below v . Then, $\hat{\ell} = \ell' + j + 1$ and $f_v(\rho') = \rho_j^{\langle \ell'+1 \rangle}$.*

Proof Let $m = \hat{\ell} - \ell' - 1$. First we show that $f_v(\rho') \leq \rho_m^{\langle \ell'+1 \rangle}$, and then we argue that equality must hold. To show that $f_v(\rho') \leq \rho_m^{\langle \ell'+1 \rangle}$, suppose by contradiction that $f_v(\rho') > \rho_m^{\langle \ell'+1 \rangle}$, that is, there exists $k < m$ such that $\rho_{k+1}^{\langle \ell'+1 \rangle} < f_v(\rho') \leq \rho_k^{\langle \ell'+1 \rangle}$. Then, the query would only be forwarded to level $\ell' + 1 + j < \ell' + 1 + m = \hat{\ell}$, as we assumed that (6.5) holds for all nodes below v . This generates a contradiction, and, therefore, it must be $f_v(\rho') \leq \rho_m^{\langle \ell'+1 \rangle}$.

We now show that $f_v(\rho') = \rho_m^{\langle \ell'+1 \rangle}$. As $f_v(\rho') \leq \rho_m^{\langle \ell'+1 \rangle}$, we have that

$$\beta_v^{\mathbf{f}}(k | f_v(\rho')) \leq \beta_v^{\mathbf{f}}(k | \rho_m^{\langle \ell'+1 \rangle})$$

for all $1 \leq k \leq m$, with equality if and only if $f_v(\rho') = \rho_m^{\langle \ell'+1 \rangle}$. This yields

$$\chi_v^{\mathbf{f}}(f_v(\rho'); \rho') < \chi_v^{\mathbf{f}}(\rho_m^{\langle \ell'+1 \rangle}; \rho')$$

for $f_v(\rho') < \rho_m^{\langle \ell'+1 \rangle}$, which implies $f_v(\rho') = \rho_m^{\langle \ell'+1 \rangle}$. By (i), it must be $m = j$, which gives $\hat{\ell} = \ell' + m + 1 = \ell' + j + 1$. *square*

We now proceed with the proof. As \mathbf{f} is ℓ -tall, $f_v(\rho^v)$ must forward the query exactly to level ℓ . We first show that $\rho^v = \rho_{i+1}^{\langle \ell' \rangle}$ and $f_v(\rho_{i+1}^{\langle \ell' \rangle}) = \rho_i^{\langle \ell'+1 \rangle}$, and then we show that (6.5) holds for v . Note that the claim above implies that if v receives a ρ' -split from its parent with $\rho' > \rho_{i+1}^{\langle \ell' \rangle}$, then $f_v(\rho')$ does not forward the query to level exactly ℓ . Therefore, it suffices to show that $f_v(\rho_{i+1}^{\langle \ell' \rangle}) = \rho_i^{\langle \ell'+1 \rangle}$. Indeed, this would imply that $\rho^v = \rho_{i+1}^{\langle \ell' \rangle}$, as no better (larger) split forwards the query to level exactly ℓ . By contradiction, suppose v plays $f_v(\rho_{i+1}^{\langle \ell' \rangle}) = \hat{\rho} \neq \rho_i^{\langle \ell'+1 \rangle}$. As we are assuming v is offered a $\rho_{i+1}^{\langle \ell' \rangle}$ -split, and (ii) implies that v prefers to play $\rho_i^{\langle \ell'+1 \rangle}$ among all $\rho > \rho_i^{\langle \ell'+1 \rangle}$, it must be $\hat{\rho} < \rho_i^{\langle \ell'+1 \rangle}$. Moreover, it must be the case that $\hat{\rho}$ forwards the query below level ℓ , otherwise v would prefer to play $\rho_i^{\langle \ell'+1 \rangle}$. However, if it was the case, v would prefer to play $\hat{\rho}$ over $\rho_i^{\langle \ell'+1 \rangle}$ when offered any ρ' -split with $\rho' < \rho_{i+1}^{\langle \ell' \rangle}$. This contradicts the assumption that \mathbf{f} is ℓ -tall, for which there exists $\rho' = \rho^v$ such that $f_v(\rho')$ forwards the query exactly to level ℓ .

We showed that $\rho_v = \rho_{i+1}^{\langle \ell' \rangle}$ and $f_v(\rho_v) = g_v(\rho_v)$. To complete the inductive step, we need to prove that $f_v(\rho') = g_v(\rho')$ for all $\rho^v \leq \rho' \leq \rho_1$. Fix any $\rho^v \leq \rho' \leq \rho_1$. We already proved that $f_v(\rho')$ does not forward the query exactly to level ℓ . Moreover, $f_v(\rho')$ cannot forward the the query below level $\ell' > \ell$, as otherwise v would prefer this strategy even when offered a ρ_v -split. Thus, $f_v(\rho')$ must forward the query to some level $\hat{\ell} < \ell$, and Claim 4 concludes the proof.

6.7.5 Proof of Theorem 9

Suppose condition (6.2) holds, that is,

$$r^* \geq 4 \cdot h \cdot \max \left\{ 1, \max_{\substack{1 \leq \ell \leq h \\ 1 \leq i < h-\ell}} \Gamma_i^{(\ell)} \right\}.$$

We show by induction that, if the discretization parameter M is large enough, for all $1 \leq \ell \leq h$ and $1 \leq i \leq h - \ell + 1$, $\rho_i^{(\ell)}$ is defined and satisfies

$$1 - \frac{1}{r^* - i} < \rho_i^{(\ell)} \leq 1 - \frac{1}{r^* - i + 1}, \quad (6.7)$$

that is, h -consistency holds.

By definition we have $\rho_1^{(\ell)} = \rho_1 = 1 - 1/r^*$, for all $1 \leq \ell \leq h$. Therefore (6.7) holds for all $1 \leq \ell \leq h$ and $i = 1$. Fix $\ell \leq h$ and suppose the claim holds for all $\ell \leq \ell' \leq h$ and $1 \leq i \leq h - \ell'$. We recall that $\rho_{i+1}^{(\ell-1)}$ is defined as

$$\rho_{i+1}^{(\ell-1)} = \max \{ \rho \in \mathcal{D}_M : e_i^{(\ell-1)}(\rho) \geq e_{i-1}^{(\ell-1)}(\rho) \},$$

where

$$e_i^{(\ell-1)}(\rho) = \sum_{j=1}^i \gamma_j^{(\ell-1)} \left[(1 - \rho) r^* \left(\prod_{t=0}^{j-1} \rho_{i-t}^{(\ell-1)+t+1} \right) - 1 \right].$$

By definition of $\rho_{i+1}^{(\ell-1)}$, it must be that

$$1 - \frac{1}{r^* \Delta_i} - \frac{1}{M} \leq \rho_{i+1}^{(\ell-1)} \leq 1 - \frac{1}{r^* \Delta_i},$$

where

$$\Delta_i = \prod_{j=0}^{i-1} \rho_{i-j}^{(\ell+j)} - \sum_{j=1}^{i-1} \frac{\gamma_j^{(\ell)}}{\gamma_i^{(\ell)}} \left[\prod_{t=0}^{j-1} \rho_{i-t-1}^{(\ell+t)} - \prod_{t=0}^{j-1} \rho_{i-t}^{(\ell+t)} \right],$$

and M is the discretization parameter of the domain \mathcal{D}_M . To see this, compute the difference $e_i^{\langle \ell-1 \rangle}(\rho_{i+1}^{\langle \ell-1 \rangle}) - e_{i-1}^{\langle \ell-1 \rangle}(\rho_{i+1}^{\langle \ell-1 \rangle})$, and argue that $1 \leq (1 - \rho_{i+1}^{\langle \ell-1 \rangle})r^* \Delta_i \leq 1 + r^* \Delta_i / M$.

We find lower and upper bounds to the term between brackets in the expression for Δ_i . First, by the inductive hypothesis, $\rho_{i-t-1}^{\langle \ell+t \rangle} > \rho_{i-t}^{\langle \ell+t \rangle}$ for all $0 \leq t \leq h - \ell$ and $0 \leq t \leq j - 1$ (with $j < i$). Therefore, we have

$$\prod_{t=0}^{j-1} \rho_{i-t-1}^{\langle \ell+t \rangle} - \prod_{t=0}^{j-1} \rho_{i-t}^{\langle \ell+t \rangle} > 0.$$

Also by induction, we have

$$\begin{aligned} \prod_{t=0}^{j-1} \rho_{i-t-1}^{\langle \ell+t \rangle} - \prod_{t=0}^{j-1} \rho_{i-t}^{\langle \ell+t \rangle} &< \prod_{t=0}^{j-1} \frac{r^* - i + t + 1}{r^* - i + t + 2} - \prod_{t=0}^{j-1} \frac{r^* - i + t - 1}{r^* - i + t} \\ &= \frac{r^* - i + 1}{r^* - i + j + 1} - \frac{r^* - i - 1}{r^* - i + j - 1} \\ &= \frac{2j}{(r^* - i + j + 1)(r^* - i + j - 1)} < \frac{2i}{(r^*)^2}, \end{aligned}$$

as $j < i$ in the expression of Δ_i . Therefore, again by induction, we have

$$\frac{r^* - i - 1}{r^* - 1} - \frac{2i}{(r^*)^2} \Gamma_i^{\langle \ell \rangle} < \Delta_i < \frac{r^* - i}{r^*}.$$

The upper bound on $\rho_{i+1}^{\langle \ell-1 \rangle}$ follows immediately. For the lower bound, it suffices to show that $r^* \cdot \Delta_i > (r^* - i - 1)(1 + r^*/M)$. Also, note that this would imply that $\Delta_i > 0$, and so that $\rho_{i+1}^{\langle \ell-1 \rangle}$ is defined. Rearranging the terms, it suffices to show that

$$\frac{2i}{r^*(r^* - i - 1)} \Gamma_i^{\langle \ell \rangle} < \frac{1}{r^* - 1} - \frac{r^*}{M}. \quad (6.8)$$

By (6.2), we have that

$$i \leq \frac{r^*}{4} \min\{1, 1/\Gamma_i^{(\ell)}\} - 1.$$

Then, (6.8) holds if

$$\frac{1/2}{1 - \frac{1}{4} \min\{1, 1/\Gamma_i\}} < 1 - \frac{(r^*)^2}{M},$$

which is satisfied for M large enough.

6.7.6 Proof of Lemma 19

Recall that, for each $i \geq 0$ we defined ϕ_i as the probability that no node at level $j \leq i$ possesses the answer, and $\lambda_i = \phi_{i-1} - \phi_i$ as the probability that a node at level i possesses the answer while no node at level $j < i$ does. Also, for every $0 \leq \ell \leq h$ and $0 \leq i \leq \ell$, we defined

$$\Gamma_i^{(\ell)} = \frac{\sum_{j=1}^{i-1} \gamma_j^{(\ell)}}{\gamma_i^{(\ell)}},$$

where $\gamma_i^{(\ell)}$ is the probability that, fixed any node v at level ℓ , there is a \mathbf{g} -candidate u in v 's subtree at distance i from v , given that v is active. We recall that a node u at level ℓ' is a \mathbf{g} -candidate if, under strategy \mathbf{g} , u is an active answer-holder and there is no active answer-holder in the first $\ell' - 1$ levels. Let L_j be the event that there is an answer holder at level j of the tree, and F_j be the event that no event L_k happens for all $k \leq j$. Observe that $\Pr(L_j, F_{j-1}) = \lambda_j$ and $\Pr(F_j) = \phi_j$. Fix a node v at level $\ell < h$. Let L_j^v be the event that there is an answer holder in v 's subtree at distance j from v , and F_j^v be the event that

no L_k^v happens for all $k \leq j$. Also, let A_v be the event that v is active. We have

$$\begin{aligned}
\gamma_j^{(\ell)} &= \Pr(L_j^v, F_{\ell+j-1} | A^v) \\
&= \Pr(L_j^v | A^v, F_{\ell+j-1}) \Pr(F_{\ell+j-1} | A^v) \\
&= \Pr(L_j | F_{j-1}) \Pr(F_{\ell+j-1} | A^v) \\
&= \frac{\Pr(L_j, F_{j-1})}{\Pr(F_{j-1})} \Pr(F_{\ell+j-1} | A^v) \\
&= \frac{\Pr(L_j, F_{j-1})}{\Pr(F_{j-1})} \frac{\Pr(A^v | F_{\ell+j-1}) \Pr(F_{\ell+j-1})}{\Pr(A^v)},
\end{aligned}$$

where the third equality follows by the fact that the branching process is memory-less, and the last equality follows by Bayes' rule. Observe that the probability that v is active only depends on the existence of answer-holders on the path from the root to v or in the subtree rooted at v . Therefore, letting P^v be the event that there is no answer-holder in the path from the root to v , we can write

$$\begin{aligned}
\Pr(A^v | F_{\ell+j-1}) &= \Pr(A^v | P^v, F_{j-1}^v) \\
&= \frac{\Pr(F_{j-1}^v | A^v, P^v) \Pr(A^v | P^v)}{\Pr(F_{j-1}^v | P^v)} \\
&= \frac{\Pr(F_{j-1}) \Pr(A^v | P^v)}{\Pr(F_{j-1}^v | P^v)},
\end{aligned}$$

where the second equality follows by Bayes' rule, and the third equality by the memory-less property of the branching factor. It follows that, for all $\ell \leq h$ and $0 \leq j \leq h - \ell$,

$$\gamma_j^{(\ell)} = \frac{\Pr(L_j, F_{j-1}) \Pr(F_{\ell+j-1}) \Pr(A^v | P^v)}{\Pr(F_{j-1}^v | P^v) \Pr(A^v)}.$$

Plugging the last expression into the definition of $\Gamma_i^{(\ell)}$, we get

$$\begin{aligned}\Gamma_i^{(\ell)} &= \frac{1}{\Pr(L_i, F_{i-1}) \Pr(F_{\ell+i-1})} \sum_{j=1}^{i-1} \Pr(L_j, F_{j-1}) \Pr(F_{\ell+j-1}) \frac{\Pr(F_{i-1}^v | P^v)}{\Pr(F_{j-1}^v | P^v)} \\ &= \frac{1}{\lambda_i \phi_{\ell+i-1}} \sum_{j=1}^{i-1} \lambda_j \phi_{\ell+j-1} \frac{\Pr(F_{i-1}^v | P^v)}{\Pr(F_{j-1}^v | P^v)}.\end{aligned}$$

As $\Pr(F_{i-1}^v | P^v) \leq \Pr(F_{j-1}^v | P^v)$ for $j \leq i$, and $\phi_{\ell+j-1} \leq 1$, we have that

$$\Gamma_i^{(\ell)} \leq \frac{1}{\lambda_i \phi_{\ell+i-1}} \sum_{j=1}^{i-1} \lambda_j < \frac{1}{\phi_{\ell+i-1}} \frac{1 - \phi_{i-1}}{\lambda_i}.$$

6.7.7 Proof of Lemma 20

For all $i \geq 0$, let $\hat{\phi}_i = \phi_i/p$ be the probability that for all levels up to i no node has the answer given that the root (at level zero) does not. Observe that no node up to level $i+1$ has the answer given that the root does not if and only if the root's children and their subtrees up to depth i do not have the answer. Therefore, we have that $\hat{\phi}_{i+1} = \Psi(p \cdot \hat{\phi}_i)$, where $\Psi(x)$, $0 \leq x \leq 1$ is the probability generating function of the branching process. It follows that

$$\begin{aligned}\lambda_{i+1} &= \phi_i - \phi_{i+1} = \phi_i p \cdot \hat{\phi}_{i+1} = \phi_i - p \cdot \sum_{k=0}^d c_k \hat{\phi}_i^k p^k \\ &= \phi_i - p \sum_{k=0}^d c_k \phi_i^k > \phi_i - \sum_{k=0}^d c_k \phi_i^k = \phi_i - \Psi(\phi_i).\end{aligned}\tag{6.9}$$

For $0 < \varepsilon \leq 1 - \zeta$ and $0 \leq z < 1 - \zeta$, let

$$a(\varepsilon) = \max \left\{ \frac{1}{b-1}, \frac{1}{\varepsilon} \frac{1}{1 - \Psi'(\zeta)} \right\},$$

and

$$t(z, \varepsilon) = a(\varepsilon) \cdot (1 - z - \Psi(1 - z)) - z.$$

We need to show that, for any $0 < \varepsilon \leq 1 - \zeta$,

$$\frac{1 - \phi_i}{\lambda_{i+1}} \leq a(\varepsilon).$$

Observe that, by inequality (6.9),

$$\frac{1 - \phi_i}{\lambda_{i+1}} \leq \frac{1 - \phi_i}{\phi_i - \Psi(\phi_i)}$$

and therefore it suffices to prove that, for every $\varepsilon > 0$,

$$t(1 - \phi_i, \varepsilon) = a(\varepsilon) (\phi_i - \Psi(\phi_i)) - (1 - \phi_i) \geq 0.$$

First, observe that, for every $\varepsilon > 0$, we have $t(0, \varepsilon) = 0$, since $\Psi(1) = 1$ (see [18]). Also note that

$$\left. \frac{\partial}{\partial z} t(z, \varepsilon) \right|_{z=0} = a(\varepsilon) \cdot (\Psi'(1) - 1) - 1 = a(\varepsilon) \cdot (b - 1) - 1,$$

which is non-negative since $a(\varepsilon) \geq 1/(b - 1)$. Also, observe that $\frac{\partial^2}{\partial z^2} t(z, \varepsilon) < 0$ and $\frac{\partial}{\partial \varepsilon} t(z, \varepsilon) > 0$, for all z and ε in their respective domains. Therefore, since the function $t(z, \varepsilon)$ is continuous, it suffices to check that $\lim_{\varepsilon \rightarrow 0} t(1 - \zeta - \varepsilon, \varepsilon) \geq 0$. As $(1 - b)^{-1} \leq \varepsilon^{-1} (1 - \Psi'(\zeta))^{-1}$ for ε small enough, we have that

$$\lim_{\varepsilon \rightarrow 0} t(1 - \zeta - \varepsilon, \varepsilon) > \lim_{\varepsilon \rightarrow 0} \left[\frac{1}{1 - \Psi'(\zeta)} \frac{1}{\varepsilon} (\zeta + \varepsilon - \Psi(\zeta + \varepsilon)) \right] - 1.$$

Since $\zeta = \Psi(\zeta)$, by l'Hôpital's rule, we conclude that $\lim_{\varepsilon \rightarrow 0} t(1 - \zeta - \varepsilon, \varepsilon) > 0$.

6.7.8 Non-uniqueness of the Nash equilibrium

In this section, we discuss the existence of multiple Nash equilibria both in the game with fixed-payment contract of [129, 15] and in the game with split contracts presented in this work.

First we recall the setting of [129, 15]. Each node has an *integer-valued* function f_v ; if v is offered a reward of $r \geq 1$ by its parent, and v does not possess the answer to the query, then v offers in turn a reward of $f_v(r) < r$ to its children. Also, by definition, $f_v(1) = 0$. Kleinberg and Raghavan [129] show that a set of strategies \mathbf{f} is a Nash equilibrium if and only if, for every node v , $f_v(r^v)$ is the value x maximizing the function

$$h_v(x; r^v) = (r^v - x - 1)p_v(\mathbf{f}, x).$$

Here r^v is the reward offered to v by its parent under \mathbf{f} , and $p_v(\mathbf{f}, x)$ is the probability that the subtree below v yields the answer, given that v does not possess the answer and offers reward x to its children. This characterization of the Nash equilibria for the game with fixed-payment contract is analogous to our result of Lemma 16 for split contracts, where the optimization is with respect to the function $\chi_v^{\mathbf{f}}(\cdot, \rho^v)$.

Using the functions $h_v(x; r^v)$, it is possible to construct a set of strategies $\mathbf{g}^{\text{fixed}}$ which optimizes $h_v(x; r^v)$ for every node v and is therefore a Nash equilibrium of the game with fixed-payment contracts. Theorem 2.2 in [129] claims that $\mathbf{g}^{\text{fixed}}$ is the unique equilibrium, in the sense that any other Nash equilibrium \mathbf{f} in which $f_v(2) = 1$ is such that for all nodes v and rewards r that are reachable at v with respect to f , $f_v(r) = g_v^{\text{fixed}}(r)$. Note that this claim would imply that all equilibria have the same efficiency, in that the query is forwarded to the same levels in every equilibrium.

Unfortunately, this claim can be showed to hold true only when restricted to *best-interest* equilibria (as in our setting, see Theorem 8), that is, when considering only

equilibria where $f_v(r')$ is the value x maximizing $h_v(x; r')$, for every r' . Note that in a *best-interest* equilibrium, nodes choose their strategies to optimize their payoff for any possible offer they *may* receive. This suggests that equilibria that are *not* best-interest are somewhat pathological, as contain nodes who do not consider their payoff globally. It is possible to show that both games admit (non-best-interest) equilibria that can be very inefficient in the sense that the query is only forwarded to a constant number of levels in the tree no matter how large the available investment r^* is. We present one of these equilibria for the case of fixed-payment contracts (the case with split contracts is similar). Consider the set of strategies \mathbf{f} in which all nodes at level 1 play $f_1(r)$, all nodes at level 2 play $f_2(r)$, and all nodes below play $f_3(r)$ (recall that the root is at level zero). For a parameter $r' \geq 4$, the functions are defined as follows.

$$f_1(r) = \begin{cases} 0, & \text{if } r = 1 \\ 1, & \text{if } r = 2 \\ 2, & \text{if } r \geq 3 \text{ and } (r - r' - 1)(\lambda_1 + \lambda_2 + \lambda_3) < (r - 2 - 1)(\lambda_1 + \lambda_2) \\ r', & \text{if } r \geq 3 \text{ and } (r - r' - 1)(\lambda_1 + \lambda_2 + \lambda_3) \geq (r - 2 - 1)(\lambda_1 + \lambda_2) \end{cases}$$

$$f_2(r) = \begin{cases} 0, & \text{if } r = 1 \\ 1, & \text{if } 2 \leq r < r' \\ 2, & \text{if } r \geq r' \end{cases}$$

$$f_3(r) = \begin{cases} 0, & \text{if } r = 1 \\ 1, & \text{if } r \geq 2 \end{cases}$$

It can be verified that \mathbf{f} is a Nash equilibrium, which thus forwards the query to level at most 3, regardless of the reward r^* offered by the root to the nodes at level 1. The bottleneck in the equilibrium is created by the nodes at level 3 or more, who cannot forward the query more than a single level as they never offer their children more than 1;

in light of this, the nodes at level 2 are not going to offer their children more than 2 (and they do so when receiving at least r'), and in turn the nodes at level 1 do not offer more than r' . This causes the query not to be forwarded efficiently. This phenomenon cannot happen in a best-interest equilibrium as, roughly speaking, the nodes at level 3 (or more) would consider the scenario in which they get offered an amount larger than 2 and realize that it is more convenient to offer their children an amount larger than 1 (assuming the nodes below reason similarly), therefore forwarding the query deeper down the tree.

6.8 Acknowledgments

Chapter 7, in full, has been submitted for publication of the material as it might appear in the ACM Transaction on Economics and Computation, under submission. L. Coviello, and M. Franceschetti.

Chapter 7

Matching markets with bundle discounts: computing efficient, stable and fair solutions

7.1 Introduction

We model a market in which vendors offer items of different types, and each buyer is interested in purchasing a unit of each type, possibly from different vendors.

Vendors are nonstrategic. Supplies are unlimited and each vendor has a fixed price for each item. Moreover, each vendor has a discount schedule according to which the bundle of all items is offered at discounted price if her demands exceed given thresholds. This can be seen as an incentive to *loyal* customers who buy from a single vendor who can sustain lower sale prices only in an economy of scale. Buyers play strategically and each selfishly tries to maximize her utility, given by the difference between the perceived value of the products and the price paid. In order to maximize their utility, buyers might be willing to cooperate to induce vendors to activate their bundle discounts. Buyers who do not purchase any bundle (i.e., who buy from several vendors) also contribute to the activation of discounts by increasing the total demands.

The externalities present in this scenario (for which a buyer's utility depends on the choices of others) make the core of the game empty in general [167]. That is, given a

market, there might not exist a configuration such that no coalition of buyers can increase their total utility by defecting. Therefore, we consider a notion of stability that looks at deviations by single buyers rather than by coalitions of buyers. Despite it is certainly a weak notion of stability in this multi-player setup, it is suitable to model a scenario where communication and coordination between buyers is mediated by a central entity (e.g., consider an online setting in which buyers might only be able to set reserve prices).

We consider the case of *transferable utility*, in which utility can be transferred between buyers in the form of side-payments, to induce cooperation. To illustrate the potential benefit of side-payments, consider a buyer who desires the bundle from a certain vendor at a discounted price. In order to trigger the discount, she might be willing to pay a subsidy to other buyers to induce them to purchase from the same vendor (as they would otherwise switch to other vendors). Given a market configuration, or *matching*, we ask whether there exist transfers that stabilize it (i.e., side-payments such that no buyer wants to deviate).

Our pricing model combines ideas from auctions with reserve prices, typical of on-line shopping websites such as Ebay (where a buyer specifies the maximum amount she is willing to pay for a product), and “deal-of-the-day” on-line purchasing, made popular by Groupon and Living Social (which offer discounted gift certificates that become valid if enough people sign up to the deal). In particular, it could model an on-line market in which buyers indicate reserve prices for products from different vendors, who in turn offer discounts if enough people sign up. In the context of our model, reserve prices might represent buyers’ valuations, each buyer prefers the pair of products with the higher difference between her reserve price and the selling price, and buyers might be willing to pay prices that are slightly different between each other in order to trigger deals. Even if the selling price of a product is higher than her reserve price, a buyer might be willing to purchase it if somebody else bears part of her cost. Similarly, if a

buyer's reserve price for a product choice is high enough with respect to a discounted price, then she might be willing to pay a price higher than the selling price to decrease the effective price of other buyers and induce them to buy – contributing to the activation of the discount. In such a scenario, a stable assignment of buyers to vendors must be computed in a centralized fashion.

Given a matching that maximizes social welfare, it is easy to prove the existence of transfers that stabilize it. However, arbitrary transfers might be undesirable for buyers, and we look for transfers that enjoy additional properties of *rationality* and *fairness*. Rationality dictates that buyers who benefit from bundle discounts are the only who pay transfers, and each only subsidizes buyers who purchase (at least one item) from her same vendor (as they might be necessary to trigger the discount). This is motivated by the willingness of each buyer to subsidize only buyers she benefits from. Fairness dictates that buyers pay transfers that are proportional to their *surplus*, that is, the difference between her current utility and the utility of their best alternative. In order to motivate this notion of fairness, observe that it might be undesirable for a buyer to pay a disproportionately large amount of the transfer needed by the buyers she benefits from if there are other buyers willing to contribute to the payment (although, from the strict point of view of stability, a buyer might be willing to transfer an amount up to her entire surplus, independently of the transfers paid by others).

Summary of results Our results show that if cooperation is allowed then social efficiency and stability can coexist in a market presenting complex externalities, and determining the right amount of cooperation is computationally tractable.

In Section 7.3, we show that, given any matching that maximizes the social welfare (or SWM matching), there exist rational transfers that stabilize it (Theorem 11 in Section 7.3). This means that *efficient* matchings are also *stable* up to suitable transfers

(the *price of stability* is one, a property that is not always observed in games [108, 188]). To prove this, we partition buyers according to their choices and surplus: on the one side, groups of “rich” buyers getting the same discounted bundle and with a positive surplus (i.e., willing to pay transfers); on the other side, groups of “poor” buyers with the same product choice and negative surplus (i.e., in need of subsidy). Then, we show that there are “rational” transfers *between groups of buyers* such that: each rich group subsidize poor groups with at least one vendor in common; each rich group transfer at most their available surplus; and each poor group receive the necessary subsidy. Group transfers can be translated into rational and stabilizing transfers. This existence result constitutes the main contribution of this work, and its proof is based on the construction of a graph which encodes the transfers between groups of buyers and has no edges if and only if the transfers are rational.

In Section 7.4, we show how transfers that are rational and fair and stabilize the market can be efficiently computed given a SWM matching. First, group transfers are computed via the Ford-Fulkerson algorithm for the maximum flow on a network such that rational group transfers and feasible flows are in one-to-one correspondence (Section 7.4.1). Then, rational and fair transfers who stabilize the SWM matching are computed (Section 7.4.2). For a market with N buyers, M vendors and c product types, group transfers are computed in time $\mathcal{O}(M^c T)$,¹ where T is the total subsidy needed. If prices and buyers’ valuations do not depend on N and M , this is $\mathcal{O}(NM^c)$. Transfers between buyers are computed in additional $\mathcal{O}(N^2 + NM^{c-1})$, for a cumulative time of $\mathcal{O}(N^2 + NM^c)$. If the the number of vendors is constant (or grows as $M^c = \mathcal{O}(N)$) then the overall complexity is dominated by the N^2 term.

¹Consider two functions $f(x)$ and $g(x)$ of a vector $x = (x_1, \dots, x_n)$. We say that $f(x) = \mathcal{O}(g(x))$ if there exist constants C, M such that $f(x) \leq Cg(x)$ for all x such that $\min_{1 \leq i \leq n} x_i \geq M$. We say that $f(x) = \Omega(g(x))$ if there exist constants C, M such that $f(x) \geq Cg(x)$ for all x such that $\min_{1 \leq i \leq n} x_i \geq M$. We say that $f(x) = \Theta(g(x))$ if both $f(x) = \mathcal{O}(g(x))$ and $f(x) = \Omega(g(x))$.

Section 7.5 deals with the computation of SWM matchings. A natural approach consists in a mixed integer program, see [185], requiring time exponential in N and M , and whose relaxation is not guaranteed to have integral solutions (i.e., corresponding to valid matchings). Conditional on the number of buyers assigned to each pair of vendors, we compute a SWM matching in time $\Theta(N^2M^c)$ via the Ford-Fulkerson algorithm for the maximum flow with minimum cost on a network such that maximum flows and feasible matchings are in one-to-one correspondence. Computing a SWM matching requires to consider a number of cases of the order of N^{M^c} , and this term dominates the computational complexity. Getting rid of the exponential dependency on M does not seem possible, due to the theoretical hardness of the problem. However, the overall time complexity is polynomial in N , and usually M can be assumed much smaller than N or even constant.

We conclude with a discussion in Section 7.6.

Related work Matching models have always received considerable attention by computer scientists [185, 175, 107, 117] and economists [186, 78, 184, 98, 97, 17], as they constitute the abstraction of many real world strategic scenarios in which choice requires mutual agreement. Examples are retail markets, the labor market, college admissions, and the assignment of residents to hospitals. When externalities are present in the market, stability often becomes problematic [98, 191]. In our model, the externalities are the numbers of buyers purchasing each product from each vendor, as they determine who benefits from discounts.

Online retailing has seen a continuous growth during the last two decades [137], and recently “deal of the day” websites such as Groupon and Living Social have introduced a new form of buying, in which enough buyers must sign up for a deal to be valid. An overview of the literature on group-buying in the web is given by [7] and by [119].

Due to the interdependencies between buyers' choices and utility, the core of the market is nonempty only under specific assumptions about preferences or discounts [100, 47].

This paper is also related to and motivated by [149], that considers a more basic model with a single product on the market and where each vendor activates multiple discounts at increasing demand volumes. The novelty of our contribution with respect to [149] is twofold. On the one hand, we extend their model to the more general case of multiple products on the market and to the possibility for vendors to activate discount on bundles of items rather than single items. As discounts can be triggered by buyers who do not necessarily benefit from them, proving the existence of (rational) transfers that stabilize the market is nontrivial and necessitates an inductive argument on the market size (see Section 7.3). We remark that our model, results and algorithms can be extended to the case of more than two products on the market and of more complex discount schedules (see Section 7.6), therefore including [149] as a special case. On the other hand, we also consider the computational side of stability, by proposing a simple and efficient algorithm to compute transfers that stabilize the market and enjoy desirable properties of *rationality* and *fairness*.

7.2 The model

We consider a market \mathcal{M} consisting in a set of N buyers \mathcal{B} and a set of M vendors \mathcal{S} . Each vendor sells items (or products) of c types denoted by $1, \dots, c$, and we assume supplies are unlimited. Let $C = \{1, \dots, c\}$. Each buyer is willing to purchase a single unit of each item type, possibly from two or more different vendors². As a remark, a vendor $s_{\perp} \in \mathcal{S}$, called the *null vendor*, might represent the choice not to buy (i.e., buyer b choosing item k from s_{\perp} means that b does not buy item k). In what follows, prices,

²We do not make any assumption about the nature of the products, which are not assumed to be complements or substitutes. We only assume that buyers assign zero or negative valuation to sets of products involving multiple units of any single item.

discounts and valuations corresponding to such vendor s_{\perp} will be pointed out. Let \mathcal{S}^c denote the cartesian product of c copies of \mathcal{S} .

A matching is a set of tuples $\mu \subset \mathcal{B} \times \mathcal{S}^c$ such that each $b \in \mathcal{B}$ appears in *exactly* a single tuple. A matching represents buyers' choices and, for $\bar{s} = (s_1, s_2, \dots, s_c) \in \mathcal{S}^c$, $(b, \bar{s}) \in \mu$ denotes that $b \in \mathcal{B}$ purchases item k from vendor s_k for $k = 1, \dots, c$. We write $\mu(b) = \bar{s}$, and $\mu^k(b) = s_k$ for $k \in C$.

Given a matching μ , for each $\bar{s} \in \mathcal{S}^c$, let $\hat{\mu}(\bar{s}) = \{b \in \mathcal{B} : \mu(b) = \bar{s}\}$ be the set of buyers who purchase item k from vendor s_k for all $k \in C$, and let $n(\bar{s}) = |\hat{\mu}(\bar{s})|$ be its cardinality. Given a matching μ , for each $s \in \mathcal{S}$ and $k \in C$, let $\hat{\mu}^k(s) = \{b \in \mathcal{B} : \mu^k(b) = s\}$ be the set of buyers who purchase item k from vendor s and $n^k(s) = |\hat{\mu}^k(s)|$. We refer to $n(s) = (n^1(s), \dots, n^c(s)) \in \mathbb{N}^c$ as the *demand vector* of vendor s (where \mathbb{N} is the set of nonnegative integers).

Prices Vendors are nonstrategic. The prices offered by a vendor are determined by her demand vector, according to a *price schedule* defined as follows. Each vendor $s \in \mathcal{S}$ has a base price p_s^k for each item $k \in C$ (we let $p_{s_{\perp}}^k = 0$ for the null vendor s_{\perp}). Moreover, s activates discounted prices on the bundle of all items C when certain thresholds are met, as we explain next. Let $p_s^{(0)} = \sum_{k \in C} p_s^k$ the *base price* of all items offered by s . We assume that s has h vectors $\tau_i(s) = (\tau_i^1(s), \dots, \tau_i^c(s))$ for $i = 1 \dots h$, called the demand thresholds vectors of s , such that $\tau_{i+1}^k(s) \geq \tau_i^k(s)$ and $\sum_{k \in C} \tau_{i+1}^k(s) > \sum_{k \in C} \tau_i^k(s)$ for all $k \in C$ and $i = 0 \dots h - 1$.³ Let $\tau_0(s) = (0, \dots, 0)$. Let $\tau_1(s_{\perp}) = (\infty, \dots, \infty)$ for the null vendor s_{\perp} . We also assume that s has h prices $p_s^{(i)}$ for $i = 1 \dots h$ such that $p_s^{(i+1)} < p_s^{(i)}$ for all $i = 0 \dots h - 1$.

Given a matching μ , with corresponding demand vector $n(s) = (n^1(s), \dots, n^c(s))$

³The constraints on the demand thresholds can be written as $\tau_{i+1}^k(s) \geq \tau_i^k(s)$ for all $k \in C$ and $\tau_{i+1}^k(s) > \tau_i^k(s)$ for some $k \in C$.

for vendor s , s offers the bundle of all items C at a cumulative price $p_s^{(i^*)}$ where

$$i^* = \max_{0 \leq i \leq h} \{i : n^k(s) \geq \tau_i^k(s), \forall k \in C\}.$$

That is, s offers the bundle of all products at the price corresponding to the largest demand threshold vector that is met component-wise.

If s activates one of her discounts, then a buyer b such that $\mu^k(b) = s$ for all $k \in C$ (i.e., b buys all items from s) pays a price $p_s^{(i^*)}$ instead of $p_s^{(0)}$.

Let

$$\mathcal{T}(\mu) = \{s \in \mathcal{S} : \exists i > 0 \text{ s.t. } n^k(s) \geq \tau_i^k(s), \forall k \in C\} \subseteq \mathcal{S}$$

be the set of vendors who activate a discount under the matching μ .

Utility Each buyer $b \in \mathcal{B}$ has a valuation for each possible product choice $\bar{s} \in \mathcal{S}^c$. The valuation v_b of buyer $b \in \mathcal{B}$ is a mapping from \mathcal{S}^c to \mathbb{R}^+ , such that, for $\bar{s} = (s_1, \dots, s_c) \in \mathcal{S}^c$, $v_b(\bar{s})$ is the valuation b assigns to purchasing product k from vendor s_k for each $k \in C$. For each $b \in \mathcal{B}$, let $v_b(\bar{s}) = 0$ for $\bar{s} = (s_\perp, \dots, s_\perp)$ (i.e., the choice not to buy any item has zero valuation).

Given a matching μ , each $b \in \mathcal{B}$ has quasi-linear utility function given by

$$u_b(\mu) = v_b(\mu(b)) - p_b(\mu),$$

where $p_b(\mu)$ is the price paid by b under the matching μ . Given a matching μ , the price $p_b(\mu)$ is computed as follows. If $\mu^k(b) = s$ for some $s \in \mathcal{T}(\mu)$ and all $k \in C$ then b pays the price $p_s^{(i^*)}$ corresponding to the largest threshold that is met. Otherwise b pays $\sum_{k \in C} p_s^{\mu^k(b)}$, that is, the sum of the base price for each single item.

Buyers play strategically, and each tries to maximize her utility. The *social welfare* $SW(\mu)$ of a matching μ is the sum of all buyers' utilities.

$$SW(\mu) = \sum_{b \in \mathcal{S}} u_b(\mu).$$

Definition 14 *Matching μ is social welfare maximizing (SWM) if $SW(\mu) \geq SW(\mu')$ for every matching μ' .*

Transferable utility We consider markets with transferable utility. That is, utility can be transferred between buyers in the form of side-payments, made in order to induce cooperation. To illustrate the potential benefit of side-payments, consider a scenario in which the best option for buyer b is to buy all items from vendor s at a discounted price. In order to purchase the desired products at a low price, b might be willing to bear some of the cost incurred by other users purchasing one or multiple items from s , which would otherwise choose other vendors.

For each $b, b' \in \mathcal{B}$, let $t_{b \rightarrow b'} \geq 0$ denote the transfer from b to b' . Let t denote the vector of transfers between all pairs of buyers. Given matching μ and transfer vector t , let (μ, t) be the market configuration in which buyers choose items according to μ and transfers t are exchanged. The utility of $b \in \mathcal{B}$ under (μ, t) is given by $u_b(\mu, t) = u_b(\mu) + \sum_{b' \in \mathcal{B}} (t_{b' \rightarrow b} + t_{b \rightarrow b'})$, where the sum is the net amount of transfer received by b . Under the assumption of transferable utility, given a matching μ , we ask whether there exist transfers t such that (μ, t) is stable, according to a suitable notion of stability.

As a remark, transfers are not equivalent to buyers becoming intermediaries. In fact, a buyer might in general subsidize only a fraction of the transfer needed by another buyer, and a buyer might receive transfer from multiple other buyers.

Stability The strongest notion of stability for a market configuration is to exhibit the *core* property [167]. A matching-transfer pair (μ, t) , has the core property if no coalition of buyers can increase their total utility by deviating from μ . Maximizing the social welfare is necessary condition for the core property (otherwise all buyers can increase their social welfare by deviating to a SWM matching). However, the core of a market \mathcal{M} (the set of matching-transfer pairs with the core property) can be empty (refer to the example in Section 7.7). We therefore turn our attention to a notion of stability which looks at deviations by single users rather than groups of users. Given a matching-transfer pair (μ, t) , there are two ways a buyer b can deviate from it. First, b might deviate by changing her product choice (resulting in a matching μ' such that $\mu'(b) \neq \mu(b)$ and $\mu'(b') = \mu(b')$ for each $b' \neq b$). In this case b 's utility would be given by the difference between her valuation of the newly chosen product pair and the price paid. We assume that, after defection, b will not be involved in any transfer (as this would not constitute an unilateral action by b), and that she cannot enjoy any discount (as other buyers might not allow b to enjoy discounts without paying transfers). Therefore, we assume that after deviation, b pays the base prices of the chosen products. Second, b might deviate by dropping her transfers in full or in part. A buyer b who enjoys a discount from $s \in \mathcal{T}(\mu)$ can benefit from other buyers purchasing from vendor s as they can trigger a lower price for b . In this case, b 's payoff after defection assumes that buyers loose incentive to buy from vendor in s , resulting in a price increase. That is, we assume that b dropping her transfers results in the deviation by both subsidized and non-subsidized buyers purchasing from s . This assumption does not affect the validity of our results, as we will look at stabilizing SWM matchings: any SWM matching minimizes the number of buyers that need to be subsidized in order to trigger a given price, and the deviation of each of these buyers would result in a price increase. Letting μ and μ' be respectively the matching before and after defection by b , we have that $u_b(\mu') = v_b(\mu'(b)) - \sum_{k \in C} p_{\mu'(b)}^k$, in both

cases of $\mu'(b) \neq \mu(b)$ and $\mu'(b) = \mu(b)$. The following definition formalizes the notion of stability just presented.

Definition 15 *A matching-transfer pair (μ, t) is stable if no buyer can unilaterally and profitably deviate from it. That is, for all $b \in \mathcal{B}$, $u_b(\mu) + \sum_{b' \in \mathcal{B}} (t_{b' \rightarrow b} + t_{b \rightarrow b'}) \geq u_b(\mu')$ for each μ' such that $\mu'(b') = \mu(b')$ for each $b' \neq b$.*

Given a matching μ , let $u_b^*(\mu)$ be the maximum utility b can achieve by deviating from μ , and let $\sigma_b(\mu) = u_b(\mu) - u_b^*(\mu)$ be the *surplus* of b under μ . If $\sigma_b(\mu) < 0$ then b needs to receive a positive transfer in order not to deviate from $\mu(b)$ to her best alternative. If $\sigma_b(\mu) > 0$ then b might be willing to pay a subsidy to certain buyers to induce them not to deviate from μ .

Definition 16 *Given a market \mathcal{M} and a matching μ , a transfer vector t is stabilizing if the matching-transfer pair (μ, t) is stable.*

Given a SWM matching μ , the existence of a stabilizing transfer is trivial to prove.

Observation 1 *For any market \mathcal{M} and any SWM matching μ , there exist stabilizing transfers t .*

We prove Observation 1 by contradiction. Let $x = \sum_{b: \sigma_b(\mu) > 0} \sigma_b(\mu)$ be the total transfer available under μ , and let $y = \sum_{b: \sigma_b(\mu) < 0} -\sigma_b(\mu)$ be the total transfer needed. Assume there are no stabilizing transfers, that is, $x < y$. A matching in which each b such that $\sigma_b(\mu) < 0$ switches to her best alternatives has utility at least $SW(\mu) + y - x > SW(\mu)$, generating a contradiction. Observe that, maximizing the social welfare is sufficient but not necessary for the existence of stabilizing transfers (see counterexample in Section 7.8).

Rational and fair transfers We are not interested in arbitrary transfers, as they could be undesirable for certain buyers. Observe that not all buyers are willing to pay transfers.

Under a matching μ , a buyer b is willing to pay transfers to other buyers only if the price paid by b under μ is strictly smaller than the sum of the base prices of the chosen items (i.e., b buys all products from a single $s \in \mathcal{T}(\mu)$) and b has positive surplus.

Consider buyers b and b' such that b buys all products from a single $s \in \mathcal{T}(\mu)$, $\sigma_b(\mu) > 0$ and $\sigma_{b'}(\mu) < 0$. If $s \notin \mu(b') = \emptyset$ then b is not willing to pay any transfer to b' as her choice does not affect the price $p_b(\mu)$. If $s \in \mu(b') \neq \emptyset$ then b is willing to pay a transfer to buyer b' . In particular, b is willing to pay a cumulative transfer of at most $\sigma_b(\mu)$ to all such buyers b' . The reason is that these buyers might be necessary to trigger the discount b currently benefits of, and they might defect if they do not receive any transfer. Moreover, if two buyers purchase the same items from the same vendors and have the same surplus, it would be undesirable for one of them to pay a higher transfer than the other. We consider the following definitions of rationality and fairness.

Definition 17 *A transfer vector t is rational if, for all $b, b' \in \mathcal{B}$, $t_{b \rightarrow b'} > 0$ only if $\sigma_b(\mu) > 0$, $\sigma_{b'}(\mu) < 0$, $\mu^k(b) = s$ for all $k \in C$ and $s \in \mu(b')$ for some $s \in \mathcal{T}(\mu)$.*

Definition 18 *A rational transfer vector t is fair if for each $b, b' \in \mathcal{B}$ such that $\mu(b) = \mu(b')$, $\sigma_b(\mu) > 0$ and $\sigma_{b'}(\mu) > 0$, the transfers paid by b and b' are proportional to $\sigma_b(\mu)$ and $\sigma_{b'}(\mu)$.*

7.3 Existence of rational and stabilizing transfers

Our main result states that, maximizing the social welfare is sufficient condition for the existence of rational and stabilizing transfers.

Theorem 11 *For any market \mathcal{M} and SWM matching μ , there exist rational and stabilizing transfers.*

For each $s \in \mathcal{T}(\mu)$, let

$$\mathcal{P}(s) = \left\{ b \in \mathcal{B} : \mu^k(b) = s \ \forall k \in \mathcal{C}, \sigma_b(\mu) > 0 \right\}$$

be the set of buyers who purchase all items from vendor s (at a discounted price) and have positive surplus. For $s \notin \mathcal{T}(\mu)$ let $\mathcal{P}(s) = \emptyset$. Each $b \in \mathcal{P}(s)$ is willing to pay transfers up to $\sigma_b(\mu)$ to buyers who have negative surplus and purchase at least a product $k \in \mathcal{C}$ from s , for a total of

$$P(s) = \sum_{b \in \mathcal{P}(s)} \sigma_b(\mu).$$

For $s \notin \mathcal{T}(\mu)$, let $P(s) = 0$.

For each subset of vendors $x \subseteq \mathcal{S}$, let

$$\mathcal{N}(x) = \left\{ b \in \mathcal{B} : \mu^k(b) \in x \ \forall k \in \mathcal{C}, \sigma_b(\mu) < 0 \right\}$$

be the set of buyers who purchase items from all and only the vendors in x and have negative surplus. Observe that $\mathcal{N}(x) = \emptyset$ for all $|x| > c$, so we will implicitly assume $|x| \leq c$. In order not to deviate from μ by switching to her best alternative, each $b \in \mathcal{N}(x)$ must receive a transfer of $-\sigma_b(\mu)$, for a total of

$$N(x) = - \sum_{b \in \mathcal{N}(x)} \sigma_b(\mu).$$

According to Definition 17, given rational transfers t , if $b \in \mathcal{P}(s)$ and $b' \in \mathcal{N}(x)$ for some $x \subseteq \mathcal{S}$ such that $s \notin x$ then $t_{b \rightarrow b'} = 0$.

As a remark, given a SWM matching μ , if $s \notin \mathcal{T}(\mu)$ for all $s \in x \subseteq \mathcal{S}$ then $\mathcal{N}(x) = \emptyset$, otherwise, a matching with higher social welfare is obtained if buyers $\mathcal{N}(x)$

switch to their best alternatives.⁴

Group transfers In the proof of Theorem 11, we will consider transfers between groups of buyers rather than transfers between single buyers. This is enough as transfers between single buyers can be computed from group transfers in arbitrary ways (we provide a computationally efficient way which also guarantees fairness in Section 7.4.2). In particular, for $s \in \mathcal{S}$ and $x \subseteq \mathcal{S}$, let

$$\bar{t}_{s \rightarrow x} = \sum_{b \in \mathcal{P}(s)} \sum_{b' \in \mathcal{N}(x)} t_{b \rightarrow b'}$$

be the total transfer from buyers $\mathcal{P}(s)$ to buyers $\mathcal{N}(\{s\})$. If transfers t are rational then $\bar{t}_{s \rightarrow x} = 0$ whenever $s \notin x$ (and the group transfers are said to be rational). To prove Theorem 11, we need to show that there exist group transfers \bar{t} such that

$$\begin{cases} P(s) \geq \sum_{x: s \in x} \bar{t}_{s \rightarrow x} & \forall s \in \mathcal{S} \\ N(x) = \sum_{s \in x} \bar{t}_{s \rightarrow x} & \forall x \subseteq \mathcal{S} \\ \bar{t}_{s \rightarrow x} = 0 & s \notin x. \end{cases} \quad (7.1)$$

The first two constraints require that the matching μ can be stabilized by group transfers \bar{t} , while the third constraint requires \bar{t} to be rational. Group transfers \bar{t} satisfying (7.1) are said *rational and stabilizing*. We consider the following definition of *cross-transfer*.

Definition 19 For $s \in \mathcal{S}$ and $x \subseteq \mathcal{S}$, group transfers $\bar{t}_{s \rightarrow x}$ is a *cross-transfer* if $s \notin x$.

Group transfer \bar{t} are rational if all cross-transfers are zero. Transfers t (between buyers) are rational if and only if all cross-transfers (between groups) are zero.

⁴In particular, if $s \notin \mathcal{S}(\mu)$, then $\mathcal{N}(\{s\}) = \emptyset$. Buyers $\mathcal{N}(\{s\})$ are the ones who purchase all items from s and have negative surplus. When we restrict our attention to rational transfers, buyers $\mathcal{N}(\{s\})$ can only receive transfer from buyers $\mathcal{P}(s)$.

Group transfers \bar{t} and \bar{t}' are *equivalent* if buyers $\mathcal{P}(s)$ pay the same transfer and buyers $\mathcal{N}(x)$ receive the same transfer under \bar{t} and \bar{t}' .

Definition 20 *Group transfers \bar{t} and \bar{t}' are equivalent if*

$$\begin{cases} \sum_{x \subseteq \mathcal{S}} \bar{t}_{s \rightarrow x} = \sum_{x \subseteq \mathcal{S}} \bar{t}'_{s \rightarrow x} & \forall s \in \mathcal{S} \\ \sum_{s \in \mathcal{S}} \bar{t}_{s \rightarrow x} = \sum_{s \in \mathcal{S}} \bar{t}'_{s \rightarrow x} & \forall x \subseteq \mathcal{S}. \end{cases}$$

Proof of Theorem 11 We assume μ is a SWM matching. We proceed by contradiction, making the following assumption.

Assumption 2 *There are no stabilizing and rational group transfers \bar{t} . That is, for any stabilizing group transfer \bar{t} , there are no equivalent and rational group transfers \bar{t}' .*

Given a SWM matching μ and stabilizing group transfers \bar{t} , we construct a graph $G(\bar{t})$ (called the *cross-transfer graph*) which encodes all cross-transfers in \bar{t} and has no edges if and only if group transfers \bar{t} are rational. We then show that, given group transfers \bar{t} , there exist equivalent group transfers \bar{t}' such that the corresponding graph $G(\bar{t}')$ is directed and acyclic. Assumption 2 implies that any such $G(\bar{t}')$ has edges, and we complete the proof by showing that a matching μ' with $SW(\mu') > SW(\mu)$ can be obtained, generating a contradiction with the assumption that μ is a SWM matching.

Definition 21 *Given group transfers \bar{t} , the cross-transfer graph $G(\bar{t})$ is the directed graph with node set equal to \mathcal{S} , and directed edge (s, s') if and only if there exist $x \subseteq \mathcal{S}$ such that $s \notin x, s' \in x, \bar{t}_{s \rightarrow x} > 0$.*

In words, in $G(\bar{t})$ there is an edge from $s \in \mathcal{S}$ to $s' \in \mathcal{S}$ if buyers $\mathcal{P}(s)$ pay a cross-transfer to buyers $\mathcal{N}(x)$ for some $x \subseteq \mathcal{S}$ such that $s \notin x, s' \in x$. An example of cross-transfer graph is given in Figure 7.1.




Figure 7.1. Example of cross-transfer graph. Let that $\mathcal{S} = \{s_1, s_2, s_3, s_4\}$, and that the only nonzero cross-transfers are $\bar{t}_{s_1 \rightarrow x'} > 0$ for $x' = \{s_3, s_4\}$ and $\bar{t}_{s_2 \rightarrow x''} > 0$ for $x'' = \{s_4\}$. According to Definition 21, $G(\bar{t})$ has nodes $\{s_1, s_2, s_3, s_4\}$ and directed edges $\{(s_1, s_4), (s_1, s_3), (s_2, s_4)\}$.

The following results state that rational group transfers correspond to cross-transfer graphs with no edges, and that we can restrict our attention to directed acyclic graphs.

Lemma 21 *Group transfers \bar{t} are rational if and only if $G(\bar{t})$ has no edge.*

Lemma 22 *Given group transfers \bar{t} , there exist equivalent group transfers \bar{t}' such that the corresponding cross-transfer graph $G(\bar{t}')$ is acyclic.*

The proof of Lemma 21 follows by the definition of cross-transfer graph and is therefore omitted. The proof of Lemma 22 is given in Section 7.9.

Without loss of generality, consider stabilizing group transfers \bar{t} and assume that $G(\bar{t})$ is a directed acyclic graph. By Assumption 2, there are no equivalent group transfers \bar{t}' such that $G(\bar{t}')$ has no edge. A vendor $s \in \mathcal{S}$ is called a *source* node if there is no edge (s', s) in $G(\bar{t}')$, and an *internal* node otherwise. Let $\mathcal{S}^{SRC} \subseteq \mathcal{S}$ be the set of vendors corresponding to source nodes in $G(\bar{t}')$. Let $\mathcal{S}^{IN} \subseteq \mathcal{S}$ be the set of vendors corresponding to internal nodes in $G(\bar{t}')$. Let

$$\mathcal{N}^{IN} = \bigcup \{ \mathcal{N}(x) \text{ s.t. } x \subseteq \mathcal{S}^{IN} \}$$

be the set of buyers who purchase products only from vendors \mathcal{S}^{IN} (and possibly some product from the null vendor s_{\perp}) and have negative surplus. By Lemma 22, we can assume without loss of generality that all buyers who receive transfer purchase products only from vendors \mathcal{S}^{IN} . Let

$$\mathcal{P}^{IN} = \bigcup \{ \mathcal{P}(s) \text{ s.t. } s \in \mathcal{S}^{IN} \}$$

be the set of buyers who buy all items C from a single vendor in \mathcal{S}^{IN} and have positive surplus. Similarly, let

$$\mathcal{P}^{SRC} = \bigcup \{ \mathcal{P}(s) \text{ s.t. } s \in \mathcal{S}^{SRC} \}.$$

According to $G(\bar{r}')$, buyers \mathcal{P}^{IN} are not able to pay the total amount of transfer needed by buyers \mathcal{N}^{IN} , and additional transfer from \mathcal{P}^{SRC} is needed (observe that the latter buyers get no benefit from the product choice of buyers \mathcal{N}^{IN}). Under Assumption 2, letting

$$X = \sum_{b \in \mathcal{P}^{IN}} \sigma_b(\mu) \quad \text{and} \quad Y = - \sum_{b \in \mathcal{N}^{IN}} \sigma_b(\mu),$$

be the amounts of transfer made available by \mathcal{P}^{IN} and needed by \mathcal{N}^{IN} respectively, we have that $X < Y$. Consider the matching μ' in which *all* buyers \mathcal{N}^{IN} and \mathcal{P}^{IN} deviate to their best alternatives⁵. Buyers \mathcal{N}^{IN} incur a cumulative gain of at least Y (the gain would be strictly greater than Y if some new threshold is activated for these buyers, after deviation⁶). Buyers \mathcal{P}^{IN} can either gain or loose utility after deviation, but each cannot

⁵Matching μ' is the results of a deviation from μ by multiple buyers. We do not directly use this deviation to proof the stability of a matching-transfer pair (whose definition looks at unilateral deviations). We use μ' to derive a contradiction on the assumption that μ is a SWM matching.

⁶Even if for the sake of stability buyers cannot enjoy discounts after deviation, here we consider that discount thresholds might be triggered as we are interested in computing the social welfare of μ' .

incur a loss larger than $\sigma_b(\mu)$, resulting in an upper bound of X on the cumulative loss.⁷ Buyers \mathcal{P}^{SRC} cannot lose utility, as no buyer deviates from sellers \mathcal{S}^{SRC} as we consider deviations by buyers \mathcal{N}^{IN} . All remaining buyers are the ones in $\mathcal{N}(x)$ for $x \subseteq \mathcal{S}$ such that $x \cap \mathcal{S}^{SRC} \neq \emptyset$ (i.e., buyers with negative surplus who do not buy all items from \mathcal{S}^{IN}) and all buyers with nonnegative surplus who are not enjoying any discount. Since these buyers do not enjoy the discounts by vendors \mathcal{S}^{IN} , they cannot lose utility in μ' with respect to μ . We have that $SW(\mu') \geq SW(\mu) + Y - X > SW(\mu)$, generating a contradiction with the assumption that μ is a SWM matching.

7.4 Computation of fair transfers

Given a market \mathcal{M} and a SWM matching μ , Theorem 11 guarantees the existence of rational and stabilizing transfers. In this section we present an efficient procedure to compute rational and stabilizing transfers that are also fair according to Definition 18. Recall that N , M and c are the numbers of buyers, sellers and product types, respectively. We assume that a SWM matching μ is given, and we proceed as follows. In Section 7.4.1 we show how to compute rational and stabilizing group transfers $\bar{t}_{s \rightarrow x}$ from $\mathcal{P}(s)$ to $\mathcal{N}(x)$ for all $s \in \mathcal{S}$, $x \subseteq \mathcal{S}$ ($|x| \leq c$), via the max-flow Ford-Fulkerson algorithm on a flow network such that feasible flows are in one-to-one correspondence with rational group transfers. Let $T = \sum_x N(x)$ be the total transfer needed by buyers with negative surplus (that is, all b such that $\sigma_b(\mu) < 0$) and who are not purchasing both items from the same vendor. Assuming that prices and valuations are constant in N and M , and observing that $|\cup_x \mathcal{N}(x)| \leq N$, we have that $T = \mathcal{O}(N)$, and rational and stabilizing group transfers can be computed in time $\mathcal{O}(TM^c) = \mathcal{O}(NM^c)$.

Given rational and stabilizing group transfers, in Section 7.4.2 we show how

⁷It is necessary to assume that also buyers \mathcal{P}^{IN} deviate to their best alternatives, as their surplus $\sigma_b(\mu)$ depends on their best alternatives given the matching-transfer pair (μ, t) .

to compute rational and stabilizing transfers between buyers that are fair according to Definition 18. This requires time $\mathcal{O}(N^2 + NM^{c-1})$, for an overall time $\mathcal{O}(N^2 + NM^c)$.

7.4.1 Step 1: rational and stabilizing group transfers

We consider the following flow network \mathcal{G} (refer to Figure 7.2). Nodes are as follows.

- A single source node r , and a single sink node t .
- A node v_x for each $x \subseteq \mathcal{S}$, $|x| \leq c$, corresponding to $\mathcal{N}(x)$. There are $\mathcal{O}(M^c)$ such nodes.
- A node u_s for each $s \in \mathcal{S}$, corresponding $\mathcal{P}(s)$. There are M such nodes.

Edges and capacities are as follows.

- For each node v_x , an edge from r to v_x with capacity $N(x)$. Flow from r to v_x represents the total transfer to $\mathcal{N}(x)$. There are $\mathcal{O}(M^c)$ such edges.
- For each node v_x , and edge from v_x to u_s for all $s \in x$, each with capacity $N(x)$. Flow from v_x to u_s represents the group transfer from $\mathcal{P}(s)$ to $\mathcal{N}(x)$. There are $\mathcal{O}(M^c)$ such edges (as each node v_x has at most a constant number c of outgoing edges).
- For each node u_s , an edge from u_s to t with capacity $P(s)$. Flow from u_s to t represents the total transfer given by $\mathcal{P}(s)$. There are M such edges.

Given a flow f on the network \mathcal{G} , $f(x, y)$ represents the flow from node x to node y . Let $\mathcal{F}(\mathcal{N})$ be the set of all feasible flows on \mathcal{G} and $\mathcal{T}(\mathcal{M})$ be the set of all rational group transfers in the market \mathcal{M} (given the SWM matching μ). Consider the mapping $\omega : \mathcal{F}(\mathcal{N}) \rightarrow \mathcal{T}(\mathcal{M})$ such that a feasible flow $f \in \mathcal{F}(\mathcal{N})$ is mapped to group transfers $\bar{t} = \omega(f)$ such that:

$$\begin{cases} \bar{t}_{s \rightarrow x} = f(v_x, u_s) & x \subseteq \mathcal{S}, |x| \leq c, s \in \mathcal{S} \text{ such that there is edge } (v_x, u_s) \text{ in } \mathcal{G} \\ \bar{t}_{s \rightarrow x} = 0 & \text{otherwise.} \end{cases}$$




Figure 7.2. Scheme of the flow network \mathcal{G} . A single node v_x for a set $x = \{s_1, s_2, s_3\} \subseteq \mathcal{S}$ is represented. There is an edge from the source r to v_x with capacity $N(x)$, to accommodate the total transfer needed by $\mathcal{N}(x)$. For $i = 1, 2, 3$, there is an edge from v_x to u_{s_i} with capacity $N(x)$, to accommodate the transfer from $\mathcal{P}(s_i)$ to $\mathcal{N}(x)$. For $i = 1, 2, 3$, there is an edge from u_{s_i} to the sink t with capacity $P(s_i)$, to accommodate the total transfer from $\mathcal{P}(s_i)$ (transfer not only to $\mathcal{N}(x)$).

Observe that the capacity constraints on edges $(u_s, t), s \in \mathcal{S}$ imply that $\sum_x \bar{t}_{s \rightarrow x} \leq P(s)$ for all $s \in \mathcal{S}$. \bar{t} is rational as in \mathcal{G} there is no edge (v_x, u_s) for $s \notin x$.

Proposition 1 *The mapping $\omega : \mathcal{F}(\mathcal{N}) \rightarrow \mathcal{T}(\mathcal{M})$ is a bijection. Let f^* be a maximum flow on \mathcal{G} . Then, $\omega(f^*)$ defines rational and stabilizing group transfers.*

The proof is given in Section 7.10. We can therefore compute rational and stabilizing group transfers via the Ford-Fulkerson algorithm for the maximum flow (see for example [130]). To bound the running time of the algorithm, we assume that the capacities of all edges in \mathcal{G} are integer, that is, all terms $P(s)$ and $N(x)$ are integer. This is the case if valuations and prices are multiples of the same unit (e.g., dollars or cents). For a network with n nodes, e edges, integer capacities, and the total capacity of the edges exiting the source equal to T , the running time of the algorithm is $\mathcal{O}((m+n)T)$. In \mathcal{G} , we have that $n = \mathcal{O}(M^c)$, $e = \mathcal{O}(M^c)$, and $T = \sum_x N(x)$. Therefore, stabilizing group transfers can be computed in time $\mathcal{O}(TM^c)$. If we assume that prices and valuations are $\mathcal{O}(1)$ (that is, constant in the market size N, M), we have that $T = \mathcal{O}(N)$ (as $|\cup_x \mathcal{N}(x)| \leq N$) and that $\mathcal{O}(TM^c) = \mathcal{O}(NM^c)$.

7.4.2 Step 2: transfer between buyers

In this section we show how rational, fair and stabilizing transfers between buyers can be computed from rational and stabilizing group transfers. Observe that each buyer $b \in \mathcal{B}$ belongs at most to a single set $\mathcal{P}(s)$ for some $s \in \mathcal{S}$ or to a single set $\mathcal{N}(x)$ for some $x \subseteq \mathcal{S}$, $|x| \leq c$. We consider the following definition of fairness, equivalent to Definition 18 when we restrict our attention to stabilizing transfers.

Definition 22 *Given a market \mathcal{M} and a SWM matching μ , rational and stabilizing transfers t (with corresponding group transfers \bar{t}) are fair if, for each $s \in \mathcal{S}$ such that $\mathcal{P}(s) \neq \emptyset$ and each $b \in \mathcal{P}(s)$, the total transfer paid by b is*

$$\sum_{b' \in \mathcal{B}} t_{b \rightarrow b'} = \sigma_b(\mu) \sum_{k \in \mathcal{S}} \bar{t}_{s \rightarrow x} / P(s).$$

Observe that all buyers $\mathcal{P}(s)$ are required to pay a cumulative transfer of $\sum_x \bar{t}_{s \rightarrow x}$ to buyers $\bigcup_x \mathcal{N}(x)$, out of an available cumulative surplus of $P(s) = \sum_{b \in \mathcal{P}(s)} \sigma_b(\mu)$. Under rational, fair and stabilizing group transfers \bar{t} , Condition (7.1) guarantees that no buyer with $\sigma_b(\mu) > 0$ pays more than $\sigma_b(\mu)$, and that each buyer with $\sigma_b(\mu) < 0$ can receive the required side-payment.

We now present our algorithm to compute rational and fair stabilizing transfers from rational and stabilizing group transfers. First, $t_{b \rightarrow b'}$ is initialized at zero for each $b, b' \in \mathcal{B}$. Fair transfers from buyers $\mathcal{P}(s)$ (for a fixed $s \in \mathcal{S}$ such that $\mathcal{P}(s) \neq \emptyset$) are computed by algorithm \mathcal{A}_1 (in Table 4), as follows.

Assume that $\bar{t}_{s \rightarrow x} > 0$ for $x = x_1, \dots, x_h$ (with $s \in x_k$ for all $k = 1, \dots, h$), as output by the algorithm in Section 7.4.1. Observe that $h = \mathcal{O}(M^{c-1})$ as we are considering sets x such that $|x| \leq c$ and $s \in x$.

For each $b \in \mathcal{P}(s)$, at any given point in the execution of the algorithm, $\tilde{\sigma}_b$ denotes b 's residual surplus, that is, the amount b has still available to make side-payments. At

ALGORITHM 4: Algorithm \mathcal{A}_1 , transfers from buyers in $\mathcal{P}(s)$

Input: $\bar{t}_{j \rightarrow jk}$ for all $k = 0, \dots, M$, $\sigma_b(\mu)$ for all $b \in \mathcal{B}$.
Initialize: $\tilde{\sigma}_b = \sigma_b(\mu)$ for each $b \in \mathcal{P}(s)$;
for $k = 0, \dots, M$ **do**
 if $\bar{t}_{j \rightarrow jk} > 0$;
 then
 $s \leftarrow \sum_{b \in \mathcal{P}(s)} \tilde{\sigma}_b$;
 $\alpha \leftarrow \bar{t}_{j \rightarrow jk} / s$;
 $\beta \leftarrow \bar{t}_{j \rightarrow jk} / (\bar{t}_{j \rightarrow jk} + \bar{t}_{k \rightarrow jk})$;
 Algorithm \mathcal{A}_2 with input $\{\alpha \tilde{\sigma}_b : b \in \mathcal{P}(s)\}, \{-\beta \tilde{\sigma}_{b'} : b' \in \mathcal{N}(j, k)\}$;
 for $b \in \mathcal{P}(s)$ **do**
 $\tilde{\sigma}_b \leftarrow (1 - \alpha) \tilde{\sigma}_b$;
 end
 end
end

initialization, let $\tilde{\sigma}_b = \sigma_b(\mu) > 0$. Transfers to buyers $\mathcal{N}(x_k)$ are computed in phases, in increasing order of k . At each phase $k = 0, \dots, h$, let $\alpha = \bar{t}_{s \rightarrow x_k} / \sum_{b \in \mathcal{P}(s)} \tilde{\sigma}_b(\mu)$ be the ratio between the group transfer from $\mathcal{P}(s)$ to $\mathcal{N}(x_k)$ and the residual surplus of $\mathcal{P}(s)$, and let $\beta = \bar{t}_{s \rightarrow x_k} / N(x_k)$ be the fraction of transfer that $\mathcal{N}(x_k)$ receives from $\mathcal{P}(s)$, out of the total transfer from $\cup_{s' \in x_k} \mathcal{P}(s')$. Algorithm \mathcal{A}_2 in Table 5 computes transfers between buyers $\mathcal{P}(s)$ to buyers $\mathcal{N}(x_k)$ such that each $b \in \mathcal{P}(s)$ transfers $\alpha \tilde{\sigma}_b(\mu)$ and each $b' \in \mathcal{N}(x_k)$ receives $-\beta \sigma_{b'}(\mu)$. Before increasing the value of k , each $b \in \mathcal{P}(s)$ updates her residual surplus to $(1 - \alpha) \tilde{\sigma}_b(\mu)$.

The correctness of algorithm \mathcal{A}_2 is straightforward. Given this, the correctness of algorithm \mathcal{A}_1 follows by observing that, for each $s \in \mathcal{S}$ and $b \in \mathcal{P}(s)$, b 's transfer in each instance of algorithm \mathcal{A}_2 never exceed $\tilde{\sigma}_b$, and that for each $x \subseteq \mathcal{S}$, $|x| \leq c$ and $b' \in \mathcal{N}(x)$, b' receives a total of $-\sigma_{b'}(\mu)$ in the (at most) c instances of algorithm \mathcal{A}_2 she is involved in.

Time complexity Let $T_{\mathcal{A}_1}(s)$ and $T_{\mathcal{A}_2}(s, x)$ be the number of operations required, respectively, by algorithm \mathcal{A}_1 for buyers in $\mathcal{P}(s)$, and by algorithm \mathcal{A}_2 to compute transfers

ALGORITHM 5: Algorithm \mathcal{A}_2 , transfers from buyers in $\mathcal{P}(s)$ to buyers in $\mathcal{N}(j, k)$

Input: amounts offered $\{x_1, \dots, x_n\}$ by $\{b_{h_1}, \dots, b_{h_n}\}$; requested $\{y_1, \dots, y_m\}$ by $\{b_{k_1}, \dots, b_{k_m}\}$.

Output: transfers $t_{h_i \rightarrow k_\ell}$ for $i = 1, \dots, n$ and $\ell = 1, \dots, m$

Initialize: $i = 1, \ell = 1$;

while ($\ell \leq m$) and ($y_\ell > 0$) **do**

if $x_i \geq y_\ell$ **then**

$t_{h_i \rightarrow k_\ell} \leftarrow y_\ell$;

$x_i \leftarrow x_i - y_\ell$;

$\ell \leftarrow \ell + 1$;

else

$t_{h_i \rightarrow k_\ell} \leftarrow x_i$;

$y_\ell \leftarrow y_\ell - x_i$;

$i \leftarrow i + 1$;

end

end

from $\mathcal{P}(s)$ to $\mathcal{N}(x)$. The total time to compute fair, rational and stabilizing transfers is $T(M, N) = \mathcal{O}(N^2) + \sum_{s \in \mathcal{S}} T_{\mathcal{A}_1}(s)$, where the first terms accounts for the initialization of t .

We have that $T_{\mathcal{A}_2}(s, x) = \mathcal{O}(|\mathcal{P}(s)| + |\mathcal{N}(x)|)$, as during each iteration of the *while* loop, one of the indexes i and ℓ increases by one, and each iteration requires a constant number of operations.

To upper bound $T_{\mathcal{A}_1}(s)$, each iteration of the *for* loop requires $\mathcal{O}(|\mathcal{P}(s)|)$ operations to compute s , and $T_{\mathcal{A}_2}(s, x)$ operations for the execution of algorithm \mathcal{A}_2 . Therefore,

the cumulative running time is upper bounded by

$$\begin{aligned}
T(M, N) &= \mathcal{O}(N^2) + \sum_{s \in \mathcal{S}} T_{\mathcal{A}_1}(s) \\
&= \mathcal{O}(N^2) + \sum_{s \in \mathcal{S}} \sum_{|x| \leq c: s \in x} (T_{\mathcal{A}_2}(s, x) + \mathcal{O}(|\mathcal{P}(s)|)) \\
&= \mathcal{O}(N^2) + \sum_{s \in \mathcal{S}} \sum_{|x| \leq c: s \in x} \mathcal{O}(|\mathcal{P}(s)| + |\mathcal{N}(x)|) \\
&= \mathcal{O}(N^2) + \sum_{s \in \mathcal{S}} \mathcal{O}(M^{c-1}) \mathcal{O}(|\mathcal{P}(s)|) + \sum_{s \in \mathcal{S}} \sum_{|x| \leq c: s \in x} \mathcal{O}(|\mathcal{N}(x)|) \\
&= \mathcal{O}(N^2) + \mathcal{O}(M^{c-1}N) + \mathcal{O}(N) = \mathcal{O}(N^2 + M^{c-1}N)
\end{aligned}$$

as $\sum_{s \in \mathcal{S}} |\mathcal{P}(s)| \leq N$, $\sum_{s \in \mathcal{S}} \sum_{|x| \leq c: s \in x} |\mathcal{N}(x)| \leq cN$, and $|\{x \subseteq \mathcal{S} : s \in x\}| = \mathcal{O}(M^{c-1})$.

Combining with the result in Section 7.4.1, fair, rational and stabilizing transfers between buyers can be computed in time $\mathcal{O}(N^2 + NM^c)$ given a SWM matching.

7.5 Computation of social welfare maximizing matching

A natural approach to compute a SWM matching is to formulate a mixed integer program (see [185]) in which, for each $b \in \mathcal{B}$ and $\bar{s} \in \mathcal{S}^c$, a binary assignment variable $x_{b, \bar{s}}$ indicates whether $\mu(b) = \bar{s}$, and, for each $s \in \mathcal{S}$, a binary variable $z_{s, i}$ indicates whether the demand of vendor s meets the threshold $\tau_i(s)$ (and the corresponding discount is triggered). These would account to $NM^c + H$ integer variables, where $H \geq M$ is the total number of discount thresholds of all vendors, and a running time exponential in this quantity. A relaxation of the problem by letting each assignment variable lay in the interval $[0, 1]$ would leave only H integer variables (and the computational complexity exponential in M). However, the existence of an integral solution (corresponding to a valid matching) is an open question.

Instead, we follow a different approach, similar to [149]. Conditional on the number of buyers $n(\bar{s}) = |\hat{\mu}(\bar{s})|$ for each $\bar{s} \in \mathcal{S}^c$ (which we refer to as a *partition* of the buyers), we compute a SWM matching via the Ford-Fulkerson algorithm for the max-flow with min-cost in time $\mathcal{O}(N^2M^c)$. Then by considering all feasible allocations $\{n(\bar{s}) : \bar{s} \in \mathcal{S}^c\}$ (that are however exponential in M^c), we determine the SWM matching.

Let $\Pi = \{\{n(\bar{s}) : \bar{s} \in \mathcal{S}^c\} : \sum_{\bar{s} \in \mathcal{S}^c} n(\bar{s}) = N\}$ be the set of all *feasible* partitions, that is, partitions such that each of the N buyers can be assigned to a single pair of vendors.

Fix $\pi \in \Pi$, and define a flow network $\mathcal{G}(\pi)$ as follows (see Figure 7.3). Nodes are the following.

- A single source node r , and a single sink node t .
- For each $b \in \mathcal{B}$, a node b . There are N such nodes.
- For each $\bar{s} \in \mathcal{S}^c$, a node \bar{s} . There are M^c such nodes.

The edges, with corresponding capacities and costs, are as follows.

- For each $b \in \mathcal{B}$, an edge from r to b , with capacity 1 and cost 0. A unit of flow on this edge represents b being assigned to some product choice. There are N such edges.
- For each $b \in \mathcal{B}$ and $\bar{s} \in \mathcal{S}^c$, an edge from b to \bar{s} with capacity 1 and cost $-v_b(\bar{s})$, that is, the opposite of the valuation buyer b gives to product choice \bar{s} . A unit flow on this edge represents buyer $\mu(b) = \mu$. There are NM^c such edges.
- For each $\bar{s} \in \mathcal{S}^c$, an edge from \bar{s} to t with capacity $n(\bar{s})$ and cost 0. An integral flow on this edge represents the total number of buyers choosing \bar{s} . There are $(M+1)^c$ such edges.

Feasible integral flows on $\mathcal{G}(\pi)$ are in one-to-one correspondence with matchings conditional on π . Let $\mu(f)$ be the matching corresponding to flow f . Given an integral flow f on $\mathcal{G}(\pi)$, its value equals the number of buyers that are matched to vendor pairs in $\mu(f)$, and its cost equals the negative of the total valuation by buyers under $\mu(f)$.




Figure 7.3. Scheme for the flow network $\mathcal{G}(\pi)$. Nodes for a single buyer $b \in \mathcal{B}$ and a single set $\bar{s} \in \mathcal{S}^c$ are represented. The cost of $-v_b(\bar{s})$ of the edge from b to \bar{s} is the opposite of the valuation buyer b gives to product choice \bar{s} .

Every max-flow f on $\mathcal{G}(\pi)$ has value N , that is, each buyer is matched to a vendor pair under $\mu(f)$. Given $\pi \in \Pi$, the total price paid by buyers is constant for each max-flow f on $\mathcal{G}(\pi)$. Therefore, maximizing the social welfare of a matching conditional on π corresponds to minimizing the cost of an integral max-flow on $\mathcal{G}(\pi)$. The total numbers of nodes and edges in $\mathcal{G}(\pi)$ are respectively $n = \Theta(N + M^c)$ and $e = \Theta(NM^c)$, and the total capacity of the edges exiting the source is $T = N$. Since all capacities are integer, the Ford-Fulkerson algorithm finds an integral max-flow with minimum cost in time $\Theta(T(n + m)) = \Theta(N^2M^c)$.

To determine the SWM matching of \mathcal{M} , for each $\pi \in \Pi$ we need to determine, a SWM matching conditional on π , for an overall time $\Theta(N^2M^c|\Pi|)$. However, this is dominated by a term N^{M^c} (see Section 7.11).

Getting rid of the exponential dependency in M does not seem possible, due to the theoretical hardness of the problem. In fact, fixed $x > 0$, deciding whether there exists a matching μ with $SW(\mu) \geq x$ is NP-hard, (by a reduction from the Knapsack problem, as noted by [149]). Even if computationally demanding even for small M , the proposed solution requires time polynomial in the number of buyers N . Our solution is significantly more efficient than both the exhaustive maximization of social welfare over all M^{2N} matchings, and solving the integer problem above (both exponential in N). Moreover, M could in general be considered much smaller than N , or even constant.

7.6 Discussion

It is an open question whether Theorem 11 holds in the case of arbitrary price schedules, where a vendor might have several discounted prices on sets of products, as described next. Let $\mathcal{C} = \{x \subseteq C\}$ be the partition of C (i.e., the set of all 2^c subsets of C). The price schedule p_s of vendor $s \in \mathcal{S}$ is a mapping from $\mathbb{N}^c \times \mathcal{C}$ to \mathbb{R}^+ (the set of nonnegative real numbers), such that, for $n \in \mathbb{N}^c$ and $x \in \mathcal{C}$, $p_s(n, x)$ is the price for the bundle of products x offered by s under demand n . Let $p_s(n, \emptyset) = 0$ for each s and n . We require that $p_s(m, x) \leq p_s(n, x)$ for all $x \in \mathcal{C}$ if $m \geq n$ component-wise. Letting e_k be the unit vector with the k -th component equal to one and all other components equal to zero, we refer to $p_s^k = p_s(e_k, \{k\})$ as the *base price* of item k offered by s . The price paid by b under matching μ is determined as follows. For each $s \in \mathcal{S}$, let $x_b(s) = \{k \in C : \mu^k(b) = s\}$ be the set of items b purchases from s . Recalling that $n(s)$ denotes the demand vector of s under the matching μ ,

$$p_b(\mu) = \sum_{s \in \mathcal{S}} p_s(n(s), x_b(s)).$$

Buyers b such that $p_b(\mu) < \sum_{k \in C} p_{\mu^k(b)}^k$ might be willing to pay transfers.

Finally, we observe that buyers might benefit from misreporting their product valuations. For example, consider a SWM matching μ and a buyer b with negative surplus $\sigma_b(\mu)$. Let v be b 's *true* valuation of the products she is matched to. If b reports a valuation of $v' = v - x$, for $x > 0$ such that μ remains a SWM matching under the untruthful reporting, then she can receive a higher subsidy of $-\sigma_b(\mu) + x$ (Theorem 11 guarantees the existence of rational and stabilizing transfers). We leave this issue to future research.

7.7 Emptiness of the core

Consider a market with two product types A and B , vendors $\mathcal{S} = \{s_1, s_2, s_3\}$, buyers $\mathcal{B} = \{b_1, b_2, b_3\}$, and valuations

$$\begin{aligned} b_1 : \quad & v_{b_1}(s_1, s_1) = 8, \quad v_{b_1}(s_2, s_2) = 1, \quad v_{b_1}(s_3, s_3) = 1, \\ b_2 : \quad & v_{b_2}(s_2, s_2) = 8, \quad v_{b_2}(s_3, s_3) = 1, \quad v_{b_2}(s_1, s_1) = 1, \\ b_3 : \quad & v_{b_3}(s_3, s_3) = 8, \quad v_{b_3}(s_1, s_1) = 1, \quad v_{b_3}(s_2, s_2) = 1, \end{aligned}$$

and $v_b(s, s') = 0$ for each b and $s \neq s'$. Assume that $p_s^A = p_s^B = 3$ for each $s \in \mathcal{S}$. Each $s \in \mathcal{S}$ activates a discounted price of $p_s^{AB} = 2$ when thresholds $\tau_s^A = \tau_s^B = 2$ on the demand of A and B are met. There are three SWM matchings, symmetric with respect to each other. Consider one of them, for example $\mu(b_1) = \mu(b_2) = (s_1, s_1)$ and $\mu(b_3) = (s_3, s_3)$, with individual utilities $u_{b_1}(\mu) = 6$, $u_{b_2}(\mu) = -1$, $u_{b_3}(\mu) = 2$. Since $u_{b_1}(\mu) + u_{b_2} = 5$, any transfer between them would leave one with a net utility of at most 2.5. If $u_{b_1}(\mu) + t_{b_2 \rightarrow b_1} \leq 2.5$, then b_1 and b_3 can profitably deviate by agreeing on purchasing both items from s_3 , receiving a cumulative utility of 5, and allocating for example 2.9 units to b_1 and 2.1 units to b_3 . If instead $u_{b_1}(\mu) + t_{b_2 \rightarrow b_1} > 2.5$, then b_2 and b_3 can profitably deviate by agreeing on purchasing from s_2 . The analysis for all other SWM matchings is similar.

7.8 Maximizing the social welfare is not necessary for stability

Consider a variation of the example in Section 7.7 above, in which b_2 's valuations are given by

$$b_2 : \quad v_{b_2}(s_2, s_2) = 8, \quad v_{b_2}(s_3, s_3) = 1, \quad v_{b_2}(s_1, s_1) = 0.5,$$

The matching μ such that $\mu(b_1) = \mu(b_2) = (s_1, s_1)$ and $\mu(b_3) = (s_3, s_3)$ has $SW(\mu) = 13/2$ and is not SWM (the matching μ' such that $\mu'(b_2) = \mu'(b_3) = (s_2, s_2)$ and $\mu'(b_1) = (s_1, s_1)$ has $SW(\mu) = 7$). However, a transfer of $15/4$ from b_1 to b_2 makes μ stable.

7.9 Proof of Lemma 22

First we assume that $G(\bar{t})$ contains a cycle of length two, that is edges (s_1, s_2) and (s_2, s_1) for $s_1, s_2 \in \mathcal{S}$. We show that there exist equivalent group transfers \bar{t}' such that either $G(\bar{t}) = G(\bar{t}') - \{(s_1, s_2)\}$ or $G(\bar{t}) = G(\bar{t}') - \{(s_2, s_1)\}$ or $G(\bar{t}) = G(\bar{t}') - \{(s_1, s_2), (s_2, s_1)\}$.

Then, we assume that the shortest cycles in $G(\bar{t})$ have length $K > 2$ and let $\mathcal{K} = s_1, \dots, s_K, s_{K+1}$ (with s_K, s_{K+1}) be such a cycle. We show that there exist equivalent group transfers \bar{t}' such that $G(\bar{t}')$ has a cycle of length $K - 1$ obtained by replacing two adjacent edges of \mathcal{K} with a single edge. This completes the proof as each cycle can be reduced to a length-two cycle by iterating the argument and finally to a single edge.

Assume $G(\bar{t})$ contains edges (s_1, s_2) and (s_2, s_1) . Let

$$\mathcal{X}_1 = \{x \subseteq \mathcal{S} : s_1 \notin x, s_2 \in x, \bar{t}_{s_1 \rightarrow x} > 0\},$$

$$\mathcal{X}_2 = \{x \subseteq \mathcal{S} : s_2 \notin x, s_1 \in x, \bar{t}_{s_2 \rightarrow x} > 0\}.$$

Let

$$t_{s_1} = \sum_{x \in \mathcal{X}_1} \bar{t}_{s_1 \rightarrow x},$$

$$t_{s_2} = \sum_{x \in \mathcal{X}_2} \bar{t}_{s_2 \rightarrow x}$$

be respectively the total amount of cross-transfer that buyers $\mathcal{P}(s_1)$ pay to all buyers $\mathcal{N}(x), x \in \mathcal{X}_1$ and that buyers $\mathcal{P}(s_2)$ pay to buyers $\mathcal{N}(x), x \in \mathcal{X}_2$.

Suppose that $t_{s_1} \leq t_{s_2}$. We define equivalent group transfers \bar{t}' such that

$$\bar{t}'_{s_1 \rightarrow x} = 0 \quad \text{for each } x \in \mathcal{X}_1, \quad (7.2)$$

where buyers $\mathcal{P}(s_1)$ switch a cumulative amount of transfer t_{s_1} from buyers $\mathcal{N}(x), x \in \mathcal{X}_1$ to buyers $\mathcal{N}(x), x \in \mathcal{X}_2$,

$$\sum_{x \in \mathcal{X}_2} \bar{t}'_{s_1 \rightarrow x} = t_{s_1} + \sum_{x \in \mathcal{X}_2} \bar{t}_{s_1 \rightarrow x}. \quad (7.3)$$

Each group $\mathcal{N}(x), x \in \mathcal{X}_1$ receives the missing amount of transfer from buyers $\mathcal{P}(s_2)$,

$$\bar{t}'_{s_2 \rightarrow x} = \bar{t}_{s_2 \rightarrow x} + \bar{t}_{s_1 \rightarrow x} \quad \text{for each } x \in \mathcal{X}_1, \quad (7.4)$$

for a total of t_{s_1} . Buyers $\mathcal{P}(s_2)$ decrease the cross-transfer to buyers $\mathcal{N}(x), x \in \mathcal{X}_2$ by total amount t_{s_1} ,

$$\sum_{x \in \mathcal{X}_2} \bar{t}'_{s_2 \rightarrow x} = t_{s_2} - t_{s_1}. \quad (7.5)$$

The existence of equivalent group transfers \bar{t}' such that (7.2)-(7.5) hold is straightforward. Observe that $\bar{t}'_{s_1 \rightarrow x} = 0$ for all $x \in \mathcal{X}_1$, and therefore $(s_1, s_2) \notin G(\bar{t}')$. If $t_{s_2} - t_{s_1} > 0$ then $\bar{t}'_{s_2 \rightarrow x} > 0$ for some $x \in \mathcal{X}_2$ and $(s_1, s_2) \in G(\bar{t}')$, otherwise $(s_1, s_2) \notin G(\bar{t}')$. The proof in the case of $t_{s_1} > t_{s_2}$ similarly follows.

Assume now that the shortest cycles in $G(\bar{t})$ have length $K > 2$, and let \mathcal{K} be a shortest cycle. That is, \mathcal{K} is formed by edges (s_i, s_{i+1}) for $i = 1, \dots, K$, with $s_{K+1} = s_1$.

For each $k = 1, \dots, K$ let

$$\mathcal{X}_k = \{x \subseteq \mathcal{S} : s_k \notin x, s_{k+1} \in x, \bar{t}_{s_k \rightarrow x} > 0\},$$

$$t_{s_k} = \sum_{x \in \mathcal{X}_k} \bar{t}_{s_k \rightarrow x}.$$

Without loss of generality, assume that $s_1 \in \arg \min_{s_k \in \mathcal{K}} t_{s_k}$, that is $t_{s_1} \leq t_{s_k}$ for all $k = 2, \dots, K$ (which is always true up to node relabeling). By the assumption that \mathcal{K} is a cycle of minimum length, there is no chord in $G(\bar{t}')$, that is $(s_k, s_j) \notin G(\bar{t}')$ if $s_k, s_j \in \mathcal{K}, s_j \neq s_{k+1}$. We build group transfers \bar{t}' which are equivalent to \bar{t} and such that $(s_1, s_2) \notin G(\bar{t}')$ and (s_i, s_{i+1}) for $i = 2, \dots, K$ with $s_{K+1} = s_2$ is a cycle of length $K - 1$ in $G(\bar{t}')$.

Group transfers \bar{t}' are defined such that

$$\bar{t}'_{s_1 \rightarrow x} = 0 \quad \text{for each } x \in \mathcal{X}_1, \quad (7.6)$$

and buyers $\mathcal{P}(s_1)$ switch a cumulative amount of transfer t_{s_1} from buyers $\mathcal{N}(x), x \in \mathcal{X}_1$ to buyers $\mathcal{N}(x), x \in \mathcal{X}_K$,

$$\sum_{x \in \mathcal{X}_K} \bar{t}'_{s_1 \rightarrow x} = t_{s_1} + \sum_{x \in \mathcal{X}_K} \bar{t}_{s_1 \rightarrow x}. \quad (7.7)$$

Each group $\mathcal{N}(x), x \in \mathcal{X}_1$ receives the missing amount of transfer from buyers $\mathcal{P}(s_K)$,

$$\bar{t}'_{s_K \rightarrow x} = \bar{t}_{s_K \rightarrow x} + \bar{t}_{s_1 \rightarrow x} \quad \text{for each } x \in \mathcal{X}_1, \quad (7.8)$$

for a total of t_{s_1} . Buyers $\mathcal{P}(s_K)$ decrease the cross-transfer to buyers $\mathcal{N}(x), x \in \mathcal{X}_K$ by

total amount t_{s_1} ,

$$\sum_{x \in \mathcal{X}_K} \bar{t}'_{s_K \rightarrow x} = t_{s_K} - t_{s_1}. \quad (7.9)$$

The existence of equivalent group transfers \bar{t}' such that (7.6)-(7.9) hold is straightforward. Observe that $\bar{t}'_{s_1 \rightarrow x} = 0$ for all $x \in \mathcal{X}_1$, and therefore $(s_1, s_2) \notin G(\bar{t}')$. Buyers $\mathcal{P}(s_K)$ pay a transfer of t_{s_1} to groups $\mathcal{N}(x), x \in \mathcal{X}_1$. This last contribution is a cross-transfer as $s_K \notin x, s_2 \in x$ for each $x \in \mathcal{X}_1$ because $(s_1, s_K) \notin G(\bar{t})$. Therefore $(s_K, s_2) \in G(\bar{t}')$. Moreover, if $t_{s_K} - t_{s_1} > 0$ then $\bar{t}'_{s_K \rightarrow x} > 0$ for some $x \in \mathcal{X}_K$ and $(s_K, s_1) \in G(\bar{t}')$, otherwise $(s_K, s_1) \notin G(\bar{t}')$. This completes the proof.

7.10 Proof of Proposition 1

It is straightforward to see that ω is a bijection, so we only prove the second part of the claim. Let $\bar{t} = \omega(f^*)$. \bar{t} are rational group transfers (as ω is a bijection). Suppose by contradiction that \bar{t} is not stabilizing, that is, condition (7.1) does not hold for \bar{t} . Recall that condition (7.1) reads as

$$\begin{cases} P(s) \geq \sum_{x: s \in x} \bar{t}_{s \rightarrow x} & \forall s \in \mathcal{S} \\ N(x) = \sum_{s \in x} \bar{t}_{s \rightarrow x} & \forall x \subseteq \mathcal{S} \\ \bar{t}_{s \rightarrow x} = 0 & s \notin x. \end{cases}$$

First, suppose that $P(s) < \sum_{x: s \in x} \bar{t}_{s \rightarrow x}$ for some $s \in \mathcal{S}$. This would imply that the flow entering node u_s is larger than the capacity of the edge (u_s, t) , generating a contradiction with the feasibility of the maximum flow f^* . Second, suppose that $N(x) > \sum_{s \in x} \bar{t}_{s \rightarrow x}$ for some $x \subseteq \mathcal{S}, |x| \leq c$. This would imply that every flow f' on \mathcal{G} is smaller than $\sum_x N(x)$, and therefore there exist no group transfers \bar{t}' such that $N(x) = \sum_{s \in x} \bar{t}'_{s \rightarrow x} \forall x \subseteq \mathcal{S}$ for all $x \subseteq \mathcal{S}$, generating a contradiction with Theorem 11 (as feasible flows and rational group

transfers are in one-to-one correspondence). Rationality of \bar{t} implies that $\bar{t}_{s \rightarrow x} = 0$ if $s \notin x$.

7.11 Computational complexity for determining SWM matchings

We have that $|\Pi| = \binom{N+M^c-1}{M^c-1}$. To prove this, observe that computing $|\Pi|$ is equivalent to counting the number of ways in which N (indistinguishable) balls can be distributed among a sorted list of M^c set. Consider a line with $N + M^c - 1$ empty positions. There are $\binom{N+M^c-1}{M^c-1}$ ways to place $M^c - 1$ stones on the available positions. The occupied positions (in ascending order) represent the boundaries between the M^c sets, and the cardinality of each set is the number of empty positions between two successive stones (if the first position is occupied by a stone, then the first set is empty; if the ℓ -th and $(\ell + 1)$ -th positions are both occupied, then the $(\ell + 1)$ -th set is empty).

Using Stirling's approximation $n! \sim (n/e)^n (2\pi n)^{1/2}$, considering M constant, we have that

$$|\Pi| \sim \left(\frac{N}{M^c+1} + 1 \right)^{M^c+1} \left(\frac{M^c+1}{N} + 1 \right)^N \left(\frac{1}{2\pi N} + \frac{1}{2\pi(M^c-1)} \right)^{1/2}.$$

Considering M as a constant, we need time $\Theta(N^2 M^c |\Pi|)$, that is,

$$\Theta \left(N^2 M^c \left(\frac{N}{M^c-1} \right)^{M^c-1} \right),$$

By the upper bound $\binom{n}{k} \leq (en/k)^k$, the time to compute a SWM matching is

$$\Theta \left(N^2 M^c \left(\frac{eN}{M^c-1} \right)^{M^c-1} \right),$$

dominated by a term N^{M^c} .

7.12 Acknowledgments

Chapter 7, in full, has been submitted for publication of the material as it might appear in the ACM Transaction on Economics and Computation, under submission. L. Coviello, and M. Franceschetti.

Bibliography

- [1] David J Abraham, Ariel Levavi, David F Manlove, and Gregg OMalley. The stable roommates problem with globally-ranked pairs. In *Internet and Network Economics*, pages 431–444. Springer, 2007.
- [2] Daron Acemoglu, Asuman Ozdaglar, and Ercan Yildiz. Diffusion of innovations in social networks. In *IEEE Conference on Decision and Control*, pages 2329–2334, 2011.
- [3] Heiner Ackermann, Paul W Goldberg, Vahab S Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. *SIAM Journal on Computing*, 40(1):92–106, 2011.
- [4] Nitin Agarwal, Huan Liu, Lei Tang, and Philip S Yu. Identifying the influential bloggers in a community. In *Proceedings of the 2008 international conference on web search and data mining*, pages 207–218. ACM, 2008.
- [5] Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47, 2002.
- [6] Réka Albert, Hawoong Jeong, and Albert-László Barabási. Error and attack tolerance of complex networks. *Nature*, 406(6794):378–382, 2000.
- [7] Krishnan S Anand and Ravi Aron. Group buying on the web: A comparison of price-discovery mechanisms. *Management Science*, 49(11):1546–1562, 2003.
- [8] Roy M Anderson and Robert McCredie May. *Infectious diseases of humans*, volume 1. Oxford university press Oxford, 1991.
- [9] Joshua D Angrist, Guido W Imbens, and Donald B Rubin. Identification of causal effects using instrumental variables. *Journal of the American statistical Association*, 91(434):444–455, 1996.
- [10] Joshua D Angrist and Alan B Krueger. Empirical strategies in labor economics. *Handbook of labor economics*, 3:1277–1366, 1999.

- [11] Sinan Aral. Commentary-identifying social influence: A comment on opinion leadership and social contagion in new product diffusion. *Marketing Science*, 30(2):217–223, 2011.
- [12] Sinan Aral, Lev Muchnik, and Arun Sundararajan. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *Proceedings of the National Academy of Sciences*, 106(51):21544–21549, 2009.
- [13] Sinan Aral and Dylan Walker. Creating social contagion through viral product design: A randomized trial of peer influence in networks. *Management Science*, 57(9):1623–1639, 2011.
- [14] Sinan Aral and Dylan Walker. Identifying influential and susceptible members of social networks. *Science*, 337(6092):337–341, 2012.
- [15] E. Arcaute, A. Kirsch, R. Kumar, D. Liben-Nowell, and S. Vassilvitskii. On threshold behavior in query incentive networks. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 66–74, 2007.
- [16] Esther M Arkin, Sang Won Bae, Alon Efrat, Kazuya Okamoto, Joseph SB Mitchell, and Valentin Polishchuk. Geometric stable roommates. *Information Processing Letters*, 109(4):219–224, 2009.
- [17] Itai Ashlagi, Mark Braverman, and Avinatan Hassidim. Matching with couples revisited. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 335–336. ACM, 2011.
- [18] Krishna B Athreya and Samuel Karlin. *Branching processes*. Springer-Verlag, New York, 2004.
- [19] Lars Backstrom, Eric Sun, and Cameron Marlow. Find me if you can: improving geographical prediction with social and spatial proximity. In *Proceedings of the 19th international conference on World wide web*, pages 61–70. ACM, 2010.
- [20] Eytan Bakshy, Jake M Hofman, Winter A Mason, and Duncan J Watts. Everyone’s an influencer: quantifying influence on twitter. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 65–74. ACM, 2011.
- [21] Eytan Bakshy, Itamar Rosenn, Cameron Marlow, and Lada Adamic. The role of social networks in information diffusion. In *Proceedings of the 21st international conference on World Wide Web*, pages 519–528. ACM, 2012.
- [22] Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.

- [23] Duygu Balcan, Vittoria Colizza, Bruno Gonçalves, Hao Hu, José J Ramasco, and Alessandro Vespignani. Multiscale mobility networks and the spatial spreading of infectious diseases. *Proceedings of the National Academy of Sciences*, 106(51):21484–21489, 2009.
- [24] Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who’s who in networks. wanted: the key player. *Econometrica*, 74(5):1403–1417, 2006.
- [25] Abhijit Banerjee, Arun G Chandrasekhar, Esther Duflo, and Matthew O Jackson. The diffusion of microfinance. *Science*, 341(6144):1236498, 2013.
- [26] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [27] Louise Barkhuus. The mismeasurement of privacy: using contextual integrity to reconsider privacy in hci. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 367–376. ACM, 2012.
- [28] Louise Barkhuus and Anind K Dey. Location-based services for mobile telephony: a study of users’ privacy concerns. In *INTERACT*, volume 3, pages 702–712. Citeseer, 2003.
- [29] Dimitri P Bertsekas. A distributed asynchronous relaxation algorithm for the assignment problem. In *Decision and Control, 1985 24th IEEE Conference on*, volume 24, pages 1703–1704. IEEE, 1985.
- [30] Dimitri P Bertsekas. The auction algorithm: A distributed relaxation method for the assignment problem. *Annals of operations research*, 14(1):105–123, 1988.
- [31] Dimitri P Bertsekas and David A Castanon. A forward/reverse auction algorithm for asymmetric assignment problems. *Computational Optimization and Applications*, 1(3):277–297, 1992.
- [32] Stefano Boccaletti, Vito Latora, Yamir Moreno, Martin Chavez, and D-U Hwang. Complex networks: Structure and dynamics. *Physics reports*, 424(4):175–308, 2006.
- [33] Béla Bollobás, Oliver Riordan, Joel Spencer, Gábor Tusnády, et al. The degree sequence of a scale-free random graph process. *Random Structures & Algorithms*, 18(3):279–290, 2001.
- [34] Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298, 2012.

- [35] John Bound, David A Jaeger, and Regina M Baker. Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American statistical association*, 90(430):443–450, 1995.
- [36] Anton Bovier. *Statistical mechanics of disordered systems: a mathematical perspective*, volume 18. Cambridge University Press, 2006.
- [37] Christina Brandt, Nicole Immorlica, Gautam Kamath, and Robert Kleinberg. An analysis of one-dimensional schelling segregation. In *STOC*, pages 789–804, 2012.
- [38] Michael Brautbar and Michael Kearns. A clustering coefficient network formation game. *Algorithmic Game Theory*, pages 224–235, 2011.
- [39] Dirk Brockmann, Lars Hufnagel, and Theo Geisel. The scaling laws of human travel. *Nature*, 439(7075):462–465, 2006.
- [40] David Card. The causal effect of education on earnings. *Handbook of labor economics*, 3:1801–1863, 1999.
- [41] Shai Carmi, Shlomo Havlin, Scott Kirkpatrick, Yuval Shavitt, and Eran Shir. A model of internet topology using k-shell decomposition. *Proceedings of the National Academy of Sciences*, 104(27):11150–11154, 2007.
- [42] Ciro Cattuto, Wouter Van den Broeck, Alain Barrat, Vittoria Colizza, Jean-François Pinton, and Alessandro Vespignani. Dynamics of person-to-person interactions from distributed rfid sensor networks. *PloS one*, 5(7):e11596, 2010.
- [43] Damon Centola. The spread of behavior in an online social network experiment. *science*, 329(5996):1194–1197, 2010.
- [44] Damon Centola. An experimental study of homophily in the adoption of health behavior. *Science*, 334(6060):1269–1272, 2011.
- [45] Meeyoung Cha, Hamed Haddadi, Fabricio Benevenuto, and P Krishna Gummadi. Measuring user influence in twitter: The million follower fallacy. *ICWSM*, 10(10-17):30, 2010.
- [46] Tanmoy Chakraborty, Stephen Judd, Michael Kearns, and Jinsong Tan. A behavioral study of bargaining in social networks. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 243–252. ACM, 2010.
- [47] Xin Chen. Inventory centralization games with price-dependent demand and quantity discount. *Operations Research*, 57(6):1394–1406, 2009.

- [48] Zhiyuan Cheng, James Caverlee, and Kyumin Lee. You are where you tweet: a content-based approach to geo-locating twitter users. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 759–768. ACM, 2010.
- [49] Zhiyuan Cheng, James Caverlee, Kyumin Lee, and Daniel Z Sui. Exploring millions of footprints in location sharing services. *ICWSM*, 2011:81–88, 2011.
- [50] Eunjoon Cho, Seth A Myers, and Jure Leskovec. Friendship and mobility: user movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1082–1090. ACM, 2011.
- [51] Nicholas A Christakis and James H Fowler. The spread of obesity in a large social network over 32 years. *New England journal of medicine*, 357(4):370–379, 2007.
- [52] Nicholas A Christakis and James H Fowler. The collective dynamics of smoking in a large social network. *New England journal of medicine*, 358(21):2249–2258, 2008.
- [53] Nicholas A Christakis and James H Fowler. Social network sensors for early detection of contagious outbreaks. *PloS one*, 5(9):e12948, 2010.
- [54] Reuven Cohen, Keren Erez, Daniel Ben-Avraham, and Shlomo Havlin. Breakdown of the internet under intentional attack. *Physical review letters*, 86(16):3682, 2001.
- [55] Vittoria Colizza, Alain Barrat, Marc Barthélemy, and Alessandro Vespignani. Predictability and epidemic pathways in global outbreaks of infectious diseases: the sars case study. *BMC medicine*, 5(1):34, 2007.
- [56] Auguste Comte. The course of positive philosophy. *London: George Bell and Sons*, 1876.
- [57] Matthew Cook. Universality in elementary cellular automata. *Complex Systems*, 15(1):1–40, 2004.
- [58] Malcolm Coulthard. *Advances in written text analysis*. Routledge, 2002.
- [59] Robert Craigie Cross. Plato’s republic. 1964.
- [60] Yves-Alexandre de Montjoye, César A Hidalgo, Michel Verleysen, and Vincent D Blondel. Unique in the crowd: The privacy bounds of human mobility. *Scientific reports*, 3, 2013.
- [61] Devansh Dikshit and Narahari Yadati. Truthful and quality conscious query incentive networks. In *Internet and Network Economics*, pages 386–397. Springer, 2009.

- [62] Peter Sheridan Dodds, Roby Muhamad, and Duncan J Watts. An experimental study of search in global social networks. *Science*, 301(5634):827, 2003.
- [63] John R. Douceur and Thomas Moscibroda. Lottery trees: motivational deployment of networked systems. In *SIGCOMM*, pages 121–132, 2007.
- [64] Sylvie Dubois and Barbara M Horvath. Let’s tink about dat: Interdental fricatives in cajun english. *Language Variation and Change*, 10(03):245–261, 1998.
- [65] Nathan Eagle, Alex Sandy Pentland, and David Lazer. Inferring friendship network structure by using mobile phone data. *Proceedings of the National Academy of Sciences*, 106(36):15274–15278, 2009.
- [66] Yuval Emek, Ron Karidi, Moshe Tennenholtz, and Aviv Zohar. Mechanisms for multi-level marketing. In *ACM Conference on Electronic Commerce*, pages 209–218, 2011.
- [67] Daniel P Enemark, Mathew D McCubbins, Ramamohan Paturi, and Nicholas Weller. Does more connectivity help groups to solve social problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 21–26, 2011.
- [68] Norman Fairclough. *Analysing discourse: Textual analysis for social research*. Psychology Press, 2003.
- [69] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM Computer Communication Review*, volume 29, pages 251–262. ACM, 1999.
- [70] Scott L Feld. Why your friends have more friends than you do. *American Journal of Sociology*, pages 1464–1477, 1991.
- [71] Neil M Ferguson, Derek AT Cummings, Simon Cauchemez, Christophe Fraser, Steven Riley, Aronrag Meeyai, Sophon Iamsirithaworn, and Donald S Burke. Strategies for containing an emerging influenza pandemic in southeast asia. *Nature*, 437(7056):209–214, 2005.
- [72] Neil M Ferguson, Derek AT Cummings, Christophe Fraser, James C Cajka, Philip C Cooley, and Donald S Burke. Strategies for mitigating an influenza pandemic. *Nature*, 442(7101):448–452, 2006.
- [73] Neil M Ferguson, Matt J Keeling, W John Edmunds, Raymond Gani, Bryan T Grenfell, Roy M Anderson, and Steve Leach. Planning for smallpox outbreaks. *Nature*, 425(6959):681–685, 2003.

- [74] James H Fowler and Nicholas A Christakis. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the framingham heart study. *Bmj*, 337:a2338, 2008.
- [75] James H Fowler and Nicholas A Christakis. Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences*, 107(12):5334–5338, 2010.
- [76] Linton C Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1979.
- [77] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [78] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [79] Manuel Garcia-Herranz, Esteban Moro, Manuel Cebrian, Nicholas A Christakis, and James H Fowler. Using friends as sensors to detect global-scale contagious outbreaks. *PloS one*, 9(4):e92413, 2014.
- [80] Eric Gilbert and Karrie Karahalios. Predicting tie strength with social media. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 211–220. ACM, 2009.
- [81] Michelle Girvan and Mark EJ Newman. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences*, 99(12):7821–7826, 2002.
- [82] Sharad Goel and Daniel G Goldstein. Predicting individual behavior with social networks. *Marketing Science*, 33(1):82–93, 2013.
- [83] Sharad Goel, Duncan J Watts, and Daniel G Goldstein. The structure of online diffusion networks. In *Proceedings of the 13th ACM conference on electronic commerce*, pages 623–638. ACM, 2012.
- [84] Vindu Goel. facebook tinkers with users emotions in news feed experiment, stirring outcry. *New York Times*, Jun, 29, 2014.
- [85] William Goffman et al. Generalization of epidemic theory. an application to the transmission of ideas. Technical report, DTIC Document, 1964.
- [86] Marcelo FC Gomes, Ana Pastore y Piontti, Luca Rossi, Dennis Chao, Ira Longini, M Elizabeth Halloran, and Alessandro Vespignani. Assessing the international spreading risk associated with the 2014 west african ebola outbreak. *PLOS Currents Outbreaks*, 1, 2014.

- [87] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.
- [88] Sanjeev Goyal, Hoda Heidari, and Michael Kearns. Competitive contagion in networks. In *Proceedings of the ACM Symposium on Theory Of Computing*, pages 759–774, 2012.
- [89] Jesse Graham, Jonathan Haidt, and Brian A Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029, 2009.
- [90] Thilo Gross and Bernd Blasius. Adaptive coevolutionary networks: a review. *Journal of The Royal Society Interface*, 5(20):259–271, 2008.
- [91] Jamie Guillory, Jason Spiegel, Molly Drislane, Benjamin Weiss, Walter Donner, and Jeffrey Hancock. Upset now?: emotion contagion in distributed groups. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 745–748. ACM, 2011.
- [92] M Elizabeth Halloran, Ira M Longini, Azhar Nizam, and Yang Yang. Containing bioterrorist smallpox. *Science*, 298(5597):1428–1432, 2002.
- [93] M Elizabeth Halloran, Alessandro Vespignani, Nita Bharti, Leora R Feldstein, KA Alexander, Matthew Ferrari, Jeffrey Shaman, John M Drake, Travis Porco, JN Eisenberg, SY Del Valle, E Lofgren, SV Scarpino, MC Eisenberg, D Gao, JM Hyman, S Eubank, and Longini IM Jr. Ebola: mobility data. *Science*, 346(6208):433, 2014.
- [94] Eric Hand. Citizen science: People power. *Nature*, 466(7307):685, 2010.
- [95] Jason D. Hartline, Vahab S. Mirrokni, and Mukund Sundararajan. Optimal marketing strategies over social networks. In *WWW*, pages 189–198, 2008.
- [96] John William Hatfield and Scott Duke Kominers. Matching in networks with bilateral contracts. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 119–120. ACM, 2010.
- [97] John William Hatfield and Scott Duke Kominers. Matching in networks with bilateral contracts. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 119–120. ACM, 2010.
- [98] John William Hatfield and Paul R Milgrom. Matching with contracts. *American Economic Review*, pages 913–935, 2005.
- [99] Hans JAP Heesterbeek. *Mathematical epidemiology of infectious diseases: model building, analysis and interpretation*, volume 5. John Wiley & Sons, 2000.

- [100] Wilco van den Heuvel, Peter Borm, and Herbert Hamers. Economic lot-sizing games. *European Journal of Operational Research*, 176(2):1117–1130, 2007.
- [101] Cesar A Hidalgo and C Rodriguez-Sickert. The dynamics of a mobile phone network. *Physica A: Statistical Mechanics and its Applications*, 387(12):3017–3024, 2008.
- [102] Till Hoffmann, Mason A Porter, and Renaud Lambiotte. Generalized master equations for non-poisson dynamics on networks. *Physical Review E*, 86(4):046102, 2012.
- [103] John E Hopcroft and Richard M Karp. A $n^{5/2}$ algorithm for maximum matchings in bipartite. In *Switching and Automata Theory, 1971., 12th Annual Symposium on*, pages 122–125. IEEE, 1971.
- [104] Lars Hufnagel, Dirk Brockmann, and Theo Geisel. Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences of the United States of America*, 101(42):15124–15129, 2004.
- [105] Guido M Imbens and Jeffrey M Wooldridge. Recent developments in the econometrics of program evaluation. Technical report, National Bureau of Economic Research, 2008.
- [106] Amos Israeli and Alon Itai. A fast and simple randomized parallel algorithm for maximal matching. *Information Processing Letters*, 22(2):77–80, 1986.
- [107] Amos Israeli, Mathew D McCubbins, Ramamohan Paturi, and Andrea Vattani. Low memory distributed protocols for 2-coloring. In *Stabilization, Safety, and Security of Distributed Systems*, pages 303–318. Springer, 2010.
- [108] Matthew O Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of economic theory*, 71(1):44–74, 1996.
- [109] Matthew O Jackson and Leat Yariv. Diffusion of behavior and equilibrium properties in network games. *American Economic Review*, 97(2):92–98, 2007.
- [110] Hawoong Jeong, Bálint Tombor, Réka Albert, Zoltan N Oltvai, and A-L Barabási. The large-scale organization of metabolic networks. *Nature*, 407(6804):651–654, 2000.
- [111] Long Jiang, Mo Yu, Ming Zhou, Xiaohua Liu, and Tiejun Zhao. Target-dependent twitter sentiment classification. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 151–160. Association for Computational Linguistics, 2011.
- [112] Thorsten Joachims. *Learning to classify text using support vector machines: Methods, theory and algorithms*. Kluwer Academic Publishers, 2002.

- [113] J Stephen Judd and Michael Kearns. Behavioral experiments in networked trade. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 150–159. ACM, 2008.
- [114] Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral dynamics and influence in networked coloring and consensus. *Proceedings of the National Academy of Sciences*, 107(34):14978, 2010.
- [115] Stephen Judd, Michael Kearns, and Yevgeniy Vorobeychik. Behavioral conflict and fairness in social networks. *Internet and Network Economics*, pages 242–253, 2011.
- [116] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica: Journal of the Econometric Society*, pages 263–291, 1979.
- [117] Yashodhan Kanoria, Mohsen Bayati, Christian Borgs, Jennifer Chayes, and Andrea Montanari. Fast convergence of natural bargaining dynamics in exchange networks. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1518–1537. SIAM, 2011.
- [118] Richard M Karp, Eli Upfal, and Avi Wigderson. Constructing a perfect matching is in random nc. *Combinatorica*, 6(1):35–48, 1986.
- [119] Robert J Kauffman and Eric A Walden. Economics and electronic commerce: Survey and directions for research. *International Journal of Electronic Commerce*, 5:5–116, 2001.
- [120] Henry Kautz, Bart Selman, and Mehul Shah. Referral web: combining social networks and collaborative filtering. *Communications of the ACM*, 40:63–65, March 1997.
- [121] Michael Kearns. Experiments in social computation. *Communications of the ACM*, 55(10):56–67, 2012.
- [122] Michael Kearns, Stephen Judd, Jinsong Tan, and Jennifer Wortman. Behavioral experiments on biased voting in networks. *Proceedings of the National Academy of Sciences*, 106(5):1347–1352, 2009.
- [123] Michael Kearns, Stephen Judd, and Yevgeniy Vorobeychik. Behavioral experiments on a network formation game. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 690–704. ACM, 2012.
- [124] Michael Kearns, Siddharth Suri, and Nick Montfort. An experimental study of the coloring problem on human subject networks. *Science*, 313(5788):824–827, 2006.

- [125] Matt J Keeling and Pejman Rohani. *Modeling infectious diseases in humans and animals*. Princeton University Press, 2008.
- [126] David Kempe, Jon M. Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.
- [127] Maksim Kitsak, Lazaros K Gallos, Shlomo Havlin, Fredrik Liljeros, Lev Muchnik, H Eugene Stanley, and Hernán A Makse. Identification of influential spreaders in complex networks. *Nature Physics*, 6(11):888–893, 2010.
- [128] Frank Kleibergen and Richard Paap. Generalized reduced rank tests using the singular value decomposition. *Journal of econometrics*, 133(1):97–126, 2006.
- [129] Jon Kleinberg and Prabhakar Raghavan. Query incentive networks. In *IEEE Symposium on Foundations of Computer Science*, pages 132–141, 2005.
- [130] Jon Kleinberg and Eva Tardos. *Algorithm design*. Pearson Education India, 2006.
- [131] Jeffrey R Kling, Jeffrey B Liebman, and Lawrence F Katz. Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119, 2007.
- [132] Jim Koopman. Modeling infection transmission. *Annu. Rev. Public Health*, 25:303–326, 2004.
- [133] Nagaraj Kota and Y. Narahari. Threshold behavior of incentives in social networks. In *ACM international conference on Information and knowledge management*, pages 1461–1464, 2010.
- [134] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.
- [135] Adam DI Kramer, Jamie E Guillory, and Jeffrey T Hancock. Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24):8788–8790, 2014.
- [136] Kai Kupferschmidt. Estimating the ebola epidemic. *Science*, 345(6201):1108–1108, 2014.
- [137] Kenneth C Laudon and Carol Guercio Traver. *E-commerce*. Pearson Prentice Hall, 2007.
- [138] David Lazer, Alex Sandy Pentland, Lada Adamic, Sinan Aral, Albert Laszlo Barabasi, Devon Brewer, Nicholas Christakis, Noshir Contractor, James Fowler, Myron Gutmann, Tony Jebara, Gary King, Michael Macy, Deb Roy, and Marshall Van Alstyne. Life in the network: the coming age of computational social science. *Science (New York, NY)*, 323(5915):721, 2009.

- [139] Scott Lederer, Jennifer Mankoff, and Anind K Dey. Who wants to know what when? privacy preference determinants in ubiquitous computing. In *CHI'03 extended abstracts on Human factors in computing systems*, pages 724–725. ACM, 2003.
- [140] Jure Leskovec, Lars Backstrom, and Jon Kleinberg. Meme-tracking and the dynamics of the news cycle. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 497–506. ACM, 2009.
- [141] Thomas M Liggett. *Interacting particle systems*, volume 276. Springer Verlag, 1985.
- [142] Suyu Liu, Nicola Perra, Márton Karsai, and Alessandro Vespignani. Controlling contagion processes in activity driven networks. *Physical review letters*, 112(11):118702, 2014.
- [143] Alun L Lloyd and Robert M May. How viruses spread among computers and people. *Science*, 292(5520):1316–1317, 2001.
- [144] Huma Lodhi, Craig Saunders, John Shawe-Taylor, Nello Cristianini, and Chris Watkins. Text classification using string kernels. *The Journal of Machine Learning Research*, 2:419–444, 2002.
- [145] Eric T Lofgren, M Elizabeth Halloran, Caitlin M Rivers, John M Drake, Travis C Porco, Bryan Lewis, Wan Yang, Alessandro Vespignani, Jeffrey Shaman, Joseph NS Eisenberg, et al. Opinion: Mathematical models: A key tool for outbreak response. *Proceedings of the National Academy of Sciences*, 111(51):18095–18096, 2014.
- [146] Ira M Longini, Azhar Nizam, Shufu Xu, Kumnuan Ungchusak, Wanna Hanshaworakul, Derek AT Cummings, and M Elizabeth Halloran. Containing pandemic influenza at the source. *Science*, 309(5737):1083–1087, 2005.
- [147] Jan Lorenz, Heiko Rauhut, Frank Schweitzer, and Dirk Helbing. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*, 108(22):9020, 2011.
- [148] Zvi Lotker, Boaz Patt-Shamir, and Seth Pettie. Improved distributed approximate matching. In *Proceedings of the ACM symposium on Parallelism in algorithms and architectures*, pages 129–136, 2008.
- [149] Tyler Lu and Craig E Boutilier. Matching models for preference-sensitive group purchasing. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 723–740. ACM, 2012.
- [150] Charles F Manski. Identification of endogenous social effects: The reflection problem. *The review of economic studies*, 60(3):531–542, 1993.

- [151] Mathew D McCubbins, Ramamohan Paturi, and Nicholas Weller. Connected coordination network structure and group coordination. *American Politics Research*, 37(5):899–920, 2009.
- [152] Miller McPherson, Lynn Smith-Lovin, and James M Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pages 415–444, 2001.
- [153] Matthias R Mehl and James W Pennebaker. The sounds of social life: a psychometric analysis of students’ daily social environments and natural conversations. *Journal of personality and social psychology*, 84(4):857, 2003.
- [154] Stefano Merler, Marco Ajelli, Laura Fumanelli, Marcelo FC Gomes, Ana Pastore y Piontti, Luca Rossi, Dennis L Chao, Ira M Longini, M Elizabeth Halloran, and Alessandro Vespignani. Spatiotemporal spread of the 2014 outbreak of ebola virus disease in liberia and the effectiveness of non-pharmaceutical interventions: a computational modelling analysis. *The Lancet Infectious Diseases*, 2015.
- [155] Stanley Milgram. The small world problem. *Psychology today*, 2(1):60–67, 1967.
- [156] Andrea Montanari and Amin Saberi. Convergence to equilibrium in local interaction games. In *IEEE Symposium on Foundations of Computer Science*, pages 303–312. IEEE, 2009.
- [157] Jonathan Morduch. The microfinance promise. *Journal of economic literature*, pages 1569–1614, 1999.
- [158] Lev Muchnik, Sinan Aral, and Sean J Taylor. Social influence bias: A randomized experiment. *Science*, 341(6146):647–651, 2013.
- [159] Heinrich H Nax, Bary SR Pradelski, and H Peyton Young. Decentralized dynamics to optimal and stable states in the assignment game.
- [160] Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *Automatic Control, IEEE Transactions on*, 54(1):48–61, 2009.
- [161] Mark EJ Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- [162] Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in networks. *Physical review E*, 69(2):026113, 2004.
- [163] Anastasios Noulas, Salvatore Scellato, Renaud Lambiotte, Massimiliano Pontil, and Cecilia Mascolo. A tale of many cities: universal patterns in human urban mobility. *PloS one*, 7(5):e37027, 2012.

- [164] Anastasios Noulas, Salvatore Scellato, Cecilia Mascolo, and Massimiliano Pontil. An empirical study of geographic user activity patterns in foursquare. *ICWSM*, 11:70–573, 2011.
- [165] Alex Olshevsky and John N Tsitsiklis. Convergence speed in distributed consensus and averaging. *SIAM Journal on Control and Optimization*, 48(1):33–55, 2009.
- [166] Jukka-Pekka Onnela, Jari Saramäki, Jorkki Hyvönen, György Szabó, David Lazer, Kimmo Kaski, János Kertész, and Albert-László Barabási. Structure and tie strengths in mobile communication networks. *Proceedings of the National Academy of Sciences*, 104(18):7332–7336, 2007.
- [167] Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT Press, 1994.
- [168] Gergely Palla, Imre Derényi, Illés Farkas, and Tamás Vicsek. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.
- [169] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *arXiv preprint arXiv:1408.2701*, 2014.
- [170] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.
- [171] James W Pennebaker, Roger J Booth, and Martha E Francis. Linguistic inquiry and word count: Liwc [computer software]. *Austin, TX: liwc. net*, 2007.
- [172] James W Pennebaker, Cindy K Chung, Molly Ireland, Amy Gonzales, and Roger J Booth. The development and psychological properties of liwc2007, 2014.
- [173] James W Pennebaker, Matthias R Mehl, and Kate G Niederhoffer. Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, 54(1):547–577, 2003.
- [174] Nicola Perra, Andrea Baronchelli, Delia Mocanu, Bruno Gonçalves, Romualdo Pastor-Satorras, and Alessandro Vespignani. Random walks and search in time-varying networks. *Physical review letters*, 109(23):238701, 2012.
- [175] Seth Pettie and Peter Sanders. A simpler linear time $2/3-\epsilon$ approximation for maximum weight matching. *Information Processing Letters*, 91(6):271–276, 2004.
- [176] Xuan-Hieu Phan, Le-Minh Nguyen, and Susumu Horiguchi. Learning to classify short and sparse text & web with hidden topics from large-scale data collections.

- In *Proceedings of the 17th international conference on World Wide Web*, pages 91–100. ACM, 2008.
- [177] Galen Pickard, Wei Pan, Iyad Rahwan, Manuel Cebrián, Riley Crane, Anmol Madan, and Alex Pentland. Time critical social mobilization. *Science*, 334(6055):509–512, 2011.
- [178] Chiara Poletto, MF Gomes, A Pastore y Piontti, Luca Rossi, L Bioglio, Denis L Chao, Ira M Longini, M Elizabeth Halloran, Vittoria Colizza, Alessandro Vespignani, et al. Assessing the impact of travel restrictions on international spread of the 2014 west african ebola epidemic. *Eurosurveillance*, 19(42), 2014.
- [179] John W Pratt. Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, pages 122–136, 1964.
- [180] Anatol Rapoport. Spread of information through a population with socio-structural bias: I. assumption of transitivity. *The bulletin of mathematical biophysics*, 15(4):523–533, 1953.
- [181] Jonathan M Read, Ken TD Eames, and W John Edmunds. Dynamic social networks and the implications for the spread of infectious disease. *Journal of The Royal Society Interface*, 5(26):1001–1007, 2008.
- [182] Daniel M Romero, Brendan Meeder, and Jon Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 695–704. ACM, 2011.
- [183] J Niels Rosenquist, James H Fowler, and Nicholas A Christakis. Social network determinants of depression. *Molecular psychiatry*, 16(3):273–281, 2011.
- [184] Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *The Journal of Political Economy*, pages 991–1016, 1984.
- [185] Alvin E Roth, Uriel G Rothblum, and John H Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993.
- [186] Alvin E Roth and Marilda A Oliveira Sotomayor. *Two-sided matching: A study in game-theoretic modeling and analysis*. Cambridge Univ Press, 1992.
- [187] Alvin E Roth and John H Vande Vate. Random paths to stability in two-sided matching. *Econometrica: Journal of the Econometric Society*, pages 1475–1480, 1990.

- [188] Tim Roughgarden and Eva Tardos. Introduction to the inefficiency of equilibria. *Algorithmic Game Theory*, 17:443–459, 2007.
- [189] Marcel Salathé, Maria Kazandjieva, Jung Woo Lee, Philip Levis, Marcus W Feldman, and James H Jones. A high-resolution human contact network for infectious disease transmission. *Proceedings of the National Academy of Sciences*, 107(51):22020–22025, 2010.
- [190] Anand D Sarwate and Alexandros G Dimakis. The impact of mobility on gossip algorithms. *Information Theory, IEEE Transactions on*, 58(3):1731–1742, 2012.
- [191] Hiroo Sasaki and Manabu Toda. Two-sided matching problems with externalities. *Journal of Economic Theory*, 70(1):93–108, 1996.
- [192] Hiroki Sayama, Irene Pestov, Jeffrey Schoolmid, Benjamin James Bush, Chun Wong, Junichi Yamanoi, and Thilo Gross. Modeling complex systems with adaptive networks. *Computers & Mathematics with Applications*, 65(10):1645–1664, 2013.
- [193] Sebastian Schnettler. A structured overview of 50 years of small-world research. *Social Networks*, 31(3):165–178, 2009.
- [194] Herbert Alexander Simon. *Administrative behavior*, volume 4. Cambridge University Press, 1965.
- [195] Yaron Singer. How to win friends and influence people, truthfully: Influence maximization mechanisms for social networks. In *WSDM*, 2011.
- [196] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.
- [197] Bharath Sriram, Dave Fuhry, Engin Demir, Hakan Ferhatosmanoglu, and Murat Demirbas. Short text classification in twitter to improve information filtering. In *Proceedings of the 33rd international ACM SIGIR conference on Research and development in information retrieval*, pages 841–842. ACM, 2010.
- [198] Douglas O Staiger and James H Stock. *Instrumental variables regression with weak instruments*, 1994.
- [199] Juliette Stehlé, Nicolas Voirin, Alain Barrat, Ciro Cattuto, Vittoria Colizza, Lorenzo Isella, Corinne Régis, Jean-François Pinton, Nagham Khanafer, Wouter Van den Broeck, et al. Simulation of an seir infectious disease model on the dynamic contact network of conference attendees. *BMC medicine*, 9(1):87, 2011.
- [200] James H Stock and Motohiro Yogo. Testing for weak instruments in linear iv regression. *Identification and inference for econometric models: Essays in honor of Thomas Rothenberg*, 2005.

- [201] Michael Stubbs. *Text and corpus analysis: Computer-assisted studies of language and culture*. Blackwell Oxford, 1996.
- [202] Lijun Sun, Kay W Axhausen, Der-Horng Lee, and Manuel Cebrian. Efficient detection of contagious outbreaks in massive metropolitan encounter networks. *Scientific reports*, 4, 2014.
- [203] Lijun Sun, Kay W Axhausen, Der-Horng Lee, and Xianfeng Huang. Understanding metropolitan patterns of daily encounters. *Proceedings of the National Academy of Sciences*, 110(34):13774–13779, 2013.
- [204] Siddharth Suri and Duncan J Watts. Cooperation and contagion in web-based, networked public goods experiments. *PLoS One*, 6(3):e16836, 2011.
- [205] James Surowiecki. *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. Doubleday Books, 2004.
- [206] John C. Tang, Manuel Cebrián, Nicklaus A. Giacobe, Hyun-Woo Kim, Taemie Kim, and Douglas ”Beaker” Wickert. Reflecting on the darpa red balloon challenge. *Commun. ACM*, 54(4):78–85, 2011.
- [207] Yla R Tausczik and James W Pennebaker. The psychological meaning of words: Liwc and computerized text analysis methods. *Journal of language and social psychology*, 29(1):24–54, 2010.
- [208] Jeffrey Travers and Stanley Milgram. An experimental study of the small world problem. *Sociometry*, 32(4):425–443, 1969.
- [209] Thomas W Valente. *Network models of the diffusion of innovations*, volume 2. Hampton Press Cresskill, NJ, 1995.
- [210] Alessandro Vespignani. Modelling dynamical processes in complex socio-technical systems. *Nature Physics*, 8(1):32–39, 2012.
- [211] John Von Neumann, Arthur Walter Burks, et al. *Theory of self-reproducing automata*. 1966.
- [212] Jing Wang, Siddharth Suri, and Duncan J Watts. Cooperation and assortativity with endogenous partner selection. In *Proceedings of the ACM Conference of Electronic Commerce*, 2012.
- [213] Ronald Wardhaugh and Janet M Fuller. *An introduction to sociolinguistics*. John Wiley & Sons, 2014.
- [214] Duncan J Watts, Peter Sheridan Dodds, and Mark EJ Newman. Identity and search in social networks. *Science*, 296(5571):1302, 2002.

- [215] Duncan J Watts and Steven H Strogatz. Collective dynamics of small-world networks. *nature*, 393(6684):440–442, 1998.
- [216] M Wood. okcupid plays with love in user experiments. *New York Times*, Jul, 28, 2014.
- [217] Rongjing Xiang, Jennifer Neville, and Monica Rogati. Modeling relationship strength in online social networks. In *Proceedings of the 19th international conference on World wide web*, pages 981–990. ACM, 2010.
- [218] Feng Xiao and Long Wang. Asynchronous consensus in continuous-time multi-agent systems with switching topology and time-varying delays. *Automatic Control, IEEE Transactions on*, 53(8):1804–1816, 2008.
- [219] Bin Yu and Munindar P. Singh. Searching social networks. In *AAMAS*, pages 65–72, 2003.
- [220] Zhang, Jun and Van Alstyne, Marshall. SWIM: fostering social network based information search. In *Proceedings of ACM CHI 2004 Conference on Human Factors in Computing Systems*, volume 2, page 1568, 2004.