

# **Stony Brook University**



OFFICIAL COPY

**The official electronic file of this thesis or dissertation is maintained by the University Libraries on behalf of The Graduate School at Stony Brook University.**

**© All Rights Reserved by Author.**

# Causal Strategic Inference in Social and Economic Networks

A Dissertation Presented

by

**Mohammad Tanvir Irfan**

to

The Graduate School

in Partial Fulfillment of the Requirements

for the Degree of

**Doctor of Philosophy**

in

**Computer Science**

Stony Brook University

August 2013

**Stony Brook University**

The Graduate School

**Mohammad Tanvir Irfan**

We, the dissertation committee for the above candidate for the  
Doctor of Philosophy degree, hereby recommend  
acceptance of this dissertation.

Luis E. Ortiz – Dissertation Advisor  
Assistant Professor, Department of Computer Science

Joseph S.B. Mitchell – Chairperson of Defense  
Professor, Department of Applied Mathematics & Statistics

Jie Gao  
Associate Professor, Department of Computer Science

Jing Chen  
Assistant Professor, Department of Computer Science

David C. Parkes  
Harvard College Professor and  
George F. Colony Professor of Computer Science  
School of Engineering and Applied Sciences, Harvard University

This dissertation is accepted by the Graduate School.

Charles Taber  
Interim Dean of the Graduate School

Abstract of the Dissertation

# Causal Strategic Inference in Social and Economic Networks

by

Mohammad Tanvir Irfan

Doctor of Philosophy

in

Computer Science

Stony Brook University

2013

Who are the most influential senators in Congress? Is there a small coalition of senators who are influential enough to prevent filibusters? In a different setting of microfinance markets, can we predict the effects of interventions to help policy makers? In order to pursue such diverse questions, we propose *causal strategic inference*, a game-theoretic counterpart of causal probabilistic inference. Using this general framework, we study two different sets of problems, broadly on social networks and networked microfinance economies.

In the first study, we introduce a new approach to the study of influence that captures the strategic aspects of the complex interactions in a network. We design *influence games*, a new class of graphical games, as a model of the behavior of a large but finite networked population. Influence games can deal with positive as well as negative influence without having to consider network dynamics. We characterize the computational complexity of various problems on influence games, propose effective solutions to the hard problems,

and design approximation algorithms, with provable guarantees, for identifying the most influential individuals in a network. Our empirical study is based on the real-world data obtained from congressional voting records and Supreme Court rulings.

Our second study is on microfinance economies. It is motivated by the challenge of formulating economic policies without the privilege of conducting trial-and-error experiments. First, we model a microfinance market as a two-sided economy. We then learn the parameters of the model from real-world data and design algorithms for various computational problems. We show the uniqueness of equilibrium interest rates for a special case and give a constructive proof of equilibrium existence in the general case. Using data from Bangladesh and Bolivia, we show that our model captures various real-world phenomena and can be used to assist policy makers in the microfinance sector.

Despite contrasting application areas, these two studies bear a common signature that is prevalent in many other domains as well: the actions of the entities in a network-structured complex system are strategically inter-dependent. This dissertation presents a computational game-theoretic framework for studying causal questions in such scenarios.

In loving memory of three of my strongest supporters, whom I could only say  
the final goodbye from Stony Brook, thousands of miles away:  
my father-in-law *Muhammad Ayaz*,  
my grandmother *Jahanara Begum*, and  
my uncle *Abdul Mannan Chowdhury*.

# Contents

<b>List of Figures</b>	<b>viii</b>
<b>List of Tables</b>	<b>xi</b>
<b>Acknowledgements</b>	<b>xiii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Causality: A Contested Ground . . . . .	2
1.2 Causal Probabilistic Inference . . . . .	4
1.3 Causal Strategic Inference . . . . .	5
1.4 Causal Strategic Inference: A Comparative Study . . . . .	8
1.4.1 Causal Strategic Inference in Econometrics . . . . .	9
1.4.2 Causal Strategic Inference: Our Approach . . . . .	18
1.5 Organization . . . . .	22
<b>2 Causal Strategic Inference in Social Networks</b>	<b>25</b>
2.1 Influence in Networks . . . . .	26
2.1.1 Connection to Rational Calculus Models of Collective Action . . . . .	28
2.1.2 Connection to Literature on the Most Influential Nodes	34
2.1.3 A Brief Note on Mechanism Design . . . . .	36
2.2 Influence Games . . . . .	37
2.2.1 General Game-Theoretic Model . . . . .	37
2.2.2 Most Influential Nodes: Problem Formulation . . . . .	38
2.2.3 Linear Influence Games . . . . .	39
2.2.4 Connection to Polymatrix Games . . . . .	40
2.2.5 Learning Influence Games . . . . .	42
2.3 Equilibria Computation in Influence Games . . . . .	42
2.3.1 Nonnegative Influence Factors . . . . .	43
2.3.2 Special Influence Structures and Potential Games . . . . .	43
2.3.3 Tree-Structured Influence Games . . . . .	46

2.3.4	Hardness Results . . . . .	49
2.3.5	Complete Proofs of Hardness Results . . . . .	51
2.3.6	Heuristics for Computing and Counting Equilibria . . . . .	62
2.4	Computing the Most Influential Nodes . . . . .	64
2.5	Experimental Results . . . . .	67
2.5.1	Random Influence Games . . . . .	67
2.5.2	Preferential-Attachment LIGs . . . . .	69
2.5.3	Illustration: Supreme Court Rulings . . . . .	69
2.5.4	Illustration: Congressional Voting . . . . .	73
2.5.5	Contrasting with Diffusion Model . . . . .	76
2.5.6	Filibuster . . . . .	77
2.5.7	Influence of the “Gangs” of Senators . . . . .	81
2.6	Conclusion . . . . .	82
2.A	Brief Review of Collective Behavior and Collective Action in Sociology . . . . .	83
2.A.1	Classical Treatment of Collective Behavior: Mass Hysteria . . . . .	84
2.A.2	Emergent Norm Theory . . . . .	85
2.A.3	Collective Action . . . . .	86
2.B	Experimental Results in Tabular Form . . . . .	89
<b>3</b>	<b>Causal Strategic Inference in Economic Networks</b>	<b>93</b>
3.1	Microfinance Systems . . . . .	93
3.2	A Model of Microfinance Markets . . . . .	98
3.2.1	Problem Formulation . . . . .	99
3.2.2	Special Case: No Diversification of Loan Portfolios . . . . .	102
3.2.3	Eisenberg-Gale Formulation . . . . .	104
3.2.4	Equilibrium Properties of General Case . . . . .	109
3.3	Computational Scheme . . . . .	110
3.3.1	Computing an Equilibrium Point . . . . .	110
3.3.2	Learning the Parameters of the Model . . . . .	116
3.4	Empirical Study . . . . .	117
3.4.1	Case Study: Bolivia . . . . .	117
3.4.2	Case Study: Bangladesh . . . . .	123
3.5	Policy Experiments . . . . .	126
3.6	Conclusion . . . . .	129
<b>4</b>	<b>Conclusion</b>	<b>130</b>
	<b>Bibliography</b>	<b>133</b>



# List of Figures

2.1	<i>Illustration of our approach to influence in networks. Each node has a binary choice of behavior, <math>\{-1, +1\}</math>, and wants to behave like the majority of its neighbors (and is indifferent if there is a tie). We adopt pure-strategy Nash equilibrium (to be defined later), abbreviated as PSNE, as the notion of stable outcome. The network is shown in (a) and the enumeration of PSNE (a row for each PSNE, where black denotes node's behavior 1, gray <math>-1</math>) in (b). We want to achieve the objective of every node choosing 1 (desirable outcome). Selecting the set of nodes <math>\{1, 2, 3\}</math> and assigning these nodes behavior prescribed by the desirable outcome (i.e., 1 for each) lead to two consistent stable outcomes of the system, shown in (c) and (d). Thus, <math>\{1, 2, 3\}</math> cannot be a most influential set of nodes. On the other hand, selecting <math>\{1, 6\}</math> and assigning these nodes behavior 1 lead to the desirable outcome as the unique stable outcome remaining. Therefore, <math>\{1, 6\}</math> is a most influential set, even though these two nodes are at the fringes of the network. Furthermore, note that <math>\{1, 6\}</math> is not most influential in the diffusion setting, since it does not maximize the spread of behavior 1. (It should be mentioned that we study a much richer class of games here than the one shown in this example.) . . . . .</i>	28
2.2	<i>Illustration of the structure of an LIG instance from a 3-SAT instance (each undirected edge represents two arcs of opposite directions between the same two nodes). In this example, the 3-SAT instance is <math>(i_1 \vee i_2 \vee i_3) \wedge (\neg i_2 \vee i_3 \vee i_4) \wedge (\neg i_3 \vee i_4 \vee \neg i_5)</math>.</i>	53
2.3	<i>Illustration of the NP-hardness reduction of Theorem 2.3.10. The monotone one-in-three SAT instance is <math>(i_1 \vee i_2 \vee i_3) \wedge (i_2 \vee i_3 \vee i_4) \wedge (i_3 \vee i_4 \vee i_5)</math>. The threshold of each variable player is 0, and that of each clause player is <math>\epsilon</math>.</i>	58
2.4	<i>Illustration of the NP-hardness reduction of Theorem 2.3.11.</i>	59

2.5	<i>Illustration of the NP-hardness reduction (Theorem 2.3.12). For the monotone one-in-three SAT instance of Figure 2.3, we first obtain the same construction as in Theorem 2.3.10. We add two extra players, the all-satisfied-verification player and the none-satisfied-verification player, whose tasks are to verify if all clauses are satisfied and if no clause is satisfied, respectively. These two players are connected to <math>m^2</math> extra players. . . . .</i>	61
2.6	<i>A hypergraph representation of three PSNE in a 9-player game with binary actions. The PSNE shown here are the followings: <math>(1, -1, -1, 1, -1, -1, 1, -1, -1)</math> (triangle), <math>(-1, -1, -1, -1, -1, -1, 1, 1, 1)</math> (rectangle), and <math>(-1, -1, -1, -1, -1, 1, -1, 1, 1)</math> (6-gon). . . . .</i>	65
2.7	<i>PSNE computation on random LIGs. The vertical bars denote 95% confidence intervals. . . . .</i>	69
2.8	<i>PSNE computation on preferential-attachment LIGs (<math>y</math>-axis is in log scale). The vertical bars denote 95% confidence intervals. . . . .</i>	70
2.9	<i>Pictorial representation of LIG learned from data—the non-diagonal elements represent influence factors and the diagonal elements biases. The numbering of the players (from 1 to 9) corresponds to the Justices in this order: Justices A Scalia, C Thomas, WH Rehnquist, SD O’Connor, AM Kennedy, SG Breyer, DH Souter, RB Ginsburg, and JP Stevens. The darker the color of a cell, the more negative is the corresponding number. For example, the most negative number (<math>-0.2634</math>) occurs in cell <math>(5, 5)</math> (i.e., the bias of Justice Kennedy). The most positive number (<math>0.4282</math>) occurs in cell <math>(1, 2)</math> (i.e., the influence factor from Justice Scalia to Justice Thomas) and the number closest to zero is <math>0.001</math> in cell <math>(2, 4)</math>. . . . .</i>	72
2.10	<i>Graphical representation of LIG learned from data. Larger node sizes indicate higher thresholds (more stubborn). Positive influence factors are drawn as black arcs and negative as red. Thicker arcs represent higher value of influence factors. While the learned LIG is a complete graph, we have only drawn approximately half of the arcs (i.e., we are not showing the “weakest” arcs in this graph). . . . .</i>	73

2.11	<i>LIG for the 110th U.S. Congress: darker color of nodes represent higher threshold (more stubborn); thicker arcs denote influence factors of higher magnitude (only half of the original arcs with the highest magnitude of influence factors are shown here); circles denote most influential senators; rectangles denote cut nodes used in the divide-and-conquer algorithm. The shaded part of it has been magnified for better visualization in Figure 2.12.</i>	74
2.12	<i>A part of the LIG for the 110th U.S. Congress: blue nodes represent Democrat senators, red Republican, and white independent; darker color of nodes represent higher threshold (more stubborn); thicker arcs denote influence factors of higher magnitude; circled node (Senator Rockefeller) denotes one of the most influential senators; rectangles at the bottom denote cut nodes used in the divide-and-conquer algorithm.</i>	75
2.13	<i>Most influential nodes in our setting. This directed acyclic graph (dag) illustrates all possible options for node selection that our approximation algorithm considers. A source node represents a node selected in the first iteration and a sink node represents a node selected in the last step. Any directed path from a source to a sink represents a sequence of nodes selected in successive iterations by our algorithm. All nodes in the same level and having the same parent, are tied in an iteration of the algorithm. Also note that the same node can appear in different paths of the dag at different levels.</i>	77
2.14	<i>Most influential nodes in the diffusion setting. Each directed path from a source to a sink represents a sequence of nodes that</i>	78
3.1	<i>Optimal objective function values and “negative entropy” of interest rates as <math>\lambda</math> varies. This shows that the objective function value becomes large (which is undesirable) as <math>\lambda</math> grows.</i>	118
3.2	<i>Learned allocations vs. observed allocations. This shows that the learned allocations closely approximate the allocations in the observed data.</i>	119
3.3	<i>Learning curve for the Bolivia dataset under Gaussian random noise (vertical bars denote 95% CI). The learning curve does not suggest overfitting.</i>	120
3.4	<i>Learning curve for the Bolivia dataset under the Dirichlet noise model. Logarithm of the noise parameter <math>\nu</math> is shown on the x-axis (vertical bars denote 95% CI). The learning curve in this noise model does not indicate overfitting as well.</i>	121

3.5	<i>Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Gaussian noise model (vertical bars denote 95% CI). It shows that the equilibrium point computed is robust with respect to noise. . . . .</i>	122
3.6	<i>Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Dirichlet noise model. Logarithm of the noise parameter <math>\nu</math> is shown on the x-axis. Vertical bars denote 95% CI. It also shows the robustness of the computed equilibrium point although the noise model is different—Dirichlet.</i>	122
3.7	<i>Two best response dynamics of MFI Bancosol with different initialization. Both of these converged to the same solution. . .</i>	123
3.8	<i>Learned allocations vs. equilibrium allocations. The two allocations being almost identical, shows that the learning algorithm was able to capture an equilibrium point using the inner part of the nested optimization program. . . . .</i>	124
3.9	<i>Comparison among observed, learned, and equilibrium interest rates. The equilibrium interest rates and the observed interest rates do not completely match, which is fine as we do not assume the data to be an equilibrium point. . . . .</i>	125
3.10	<i>Learned allocations vs. observed allocations. Although they are not exactly the same, the learned allocations do approximate the observed ones. . . . .</i>	125
3.11	<i>Comparison among observed, learned, and equilibrium interest rates. Again, we do not assume the data to be an equilibrium point. Therefore, the difference between observed interest rates and equilibrium interest rates suggests an interesting aspect. That is, some banks are charging less than they should and some are charging more than they should. The first aspect is well-established in the literature with respect to government-owned banks. There is also practical evidence for the second aspect, which will be elaborated when we study interest rate caps.</i>	126
3.12	<i>Learned allocations vs. equilibrium allocations. The two allocations are very similar. . . . .</i>	127
3.13	<i>Equilibrium interest rates before and after adding subsidies to Eco Futuro (MFI 5). Subsidies help bring down the interest rate of not only Eco Futuro but also the other MFIs. . . . .</i>	127

# List of Tables

2.1	<i>Interpretation of Votes</i> . . . . .	71
2.2	<i>PSNE computation on uniform random directed graphs. Offsets of 95% confidence intervals are shown in parenthesis.</i> . . . . .	90
2.3	<i>Computation of the most influential nodes: Comparison between the solution given by the <code>ApproximateUniqueHyperedge</code> algorithm and the optimal one in random directed graphs. Offsets of the 95% confidence intervals are shown in parenthesis.</i> . . . . .	91
2.4	<i>LIG learned from data. Each non-diagonal element in each row represents the influence factor the row player receives from the column players (for example, Justice Scalia’s influence factor on Justice Thomas is 0.2930 and that in the opposite direction is 0.4282). The diagonal elements represent thresholds of the corresponding row players.</i> . . . . .	92

# Acknowledgements

It's been an amazing journey for me—one that actually hurts to realize that it's over. In retrospect, it's a little strange that it's painful to reach the destination that I once set out for with so much enthusiasm. It's also been a unique journey, mainly because it's been driven by many people. Many. In fact, it wouldn't be possible to acknowledge everyone here in this short space.

First and foremost, I'd like to express my utmost gratitude to my advisor **Professor Luis E. Ortiz**. Looking back at the cold December 2008 night when I wrote to you expressing an interest to work, little did I know that I was going to get not only an advisor with extraordinary knowledge and vision but also a mentor with a *unique* combination of youth and wisdom. Thank you for those countless hours of brainstorming, which gave me the very best experience of my PhD life. I've also learned a lot from your truly enlightening classes. I still preserve all those classnotes and boardshots and sometimes flip through the pages to relive some of the moments. Thank you for contributing not only to my graduate studies but also to my career planning and preparation. I've learned from you how to present my research and most importantly, how to present it to a general audience. I've also learned from you a very precious but often ignored lesson on the philosophy of conducting and publishing research. Thank you!

I'd like to thank the Chair of my dissertation committee **Professor Joseph S.B. Mitchell**. Thank you for supporting me in various ways from day one of my graduate student life at Stony Brook. When I think of a role model for research and teaching, I think of you. When I think of an inspirational leader in many aspects, I also think of you. I've learned *a lot* from all those sessions of our Algorithmic Reading Group that you conducted—how to approach new problems, how to formulate and express ideas, and most importantly, how to identify important problems. I feel fortunate to have worked with you on our guarding polyominoes project. I also feel fortunate to have sat at your Computational Geometry class. What an amazing experience it was! And I still keep that one page cheat-sheet that I prepared for the final exam, as a souvenir. Thank you!

I'm enormously grateful to my committee member **Professor Jie Gao**. Thank you for making the time to attend my Prelim as well as my defense despite being in a very challenging situation. Thank you for all your suggestions. Also, thank you for teaching an Algorithms course that my classmates and I still brag about. I'm also thankful to my committee member **Professor Jing Chen**. Having you in my committee was a great privilege. Thank you for traveling all the way from Princeton, NJ to Stony Brook only to attend my defense in person. Thank you for your suggestions as well. I wish I could stay here longer to work with you on many of the exciting topics of our shared interest. Thank you!

I'm immensely indebted to my external member **Professor David C. Parkes**. Thank you for being in my committee while being engaged in the most demanding task for any academician—writing a book. Thank you for reading the drafts of my dissertation thoroughly, raising many stimulating questions, and suggesting many important revisions. I'm especially grateful to you for pointing out the relevant econometrics literature. It was a very rewarding exercise indeed. Thank you!

I'm grateful to the great **state of New York** for supporting me as a Graduate Assistant and to **NSF** for supporting me as a Research Assistant. Coming from a middle class family in Bangladesh, it would've been impossible for me to attend a stellar institutions here without your support. Thank you!

I'd like to thank all of my labmates, including **Hau Chan** and **Chandra Sekhar Mallarapu**, for their contribution. Chandra, it was a great experience working with you on the survey propagation project. I wish I had the opportunity to work longer with you. I'm also thankful to **Abhishek Murthy** and **Ariful Islam** for giving me a desk in their lab and to our system administrator **Brian Tria** for suggesting the arrangement in the first place. Brian, your never-ending resource of humor and quick-wit made sure that your fans, including myself of course, never had a dull moment. Thank you!

I'm grateful to all my friends, near and far. I'm thankful to my boyhood best friend **Ziauddin Faruk** for exchanging messages with me every single day from Bangladesh. Seldom did I pass a boring day, and I'm really fortunate to have a friend for life *for real!* I feel fortunate to have been part of a vibrant Bangladeshi students community here at Stony Brook. I'd like to thank all of you for making my graduate student life worthwhile. I'm particularly grateful to **M. Sajjad Hossain** and **Jalal Mahmud** for making Stony Brook feel like home when I first came here. I'm also thankful to **Amitav**, **Arif**, **Asif**, **Emran**, **Faisal**, **Humed**, **Mahabub**, **Mahsiul**, **Omer**, **Shebuti**, **Utpal**, and many others for attending many of my presentations and giving me excellent feedback from time to time. Thank you!

Finally, I'm thankful to my family beyond words. My wife **Sumayya Ayaz** has made enormous sacrifices over the years for my graduate studies. Thank you for being there for me every moment of the time with unmatched patience and love. I'm thankful to our son **Rafan** for giving a bigger meaning to life, perhaps unknowingly. I don't have any word to express my gratitude to my parents for supporting me in every possible way. I'm also grateful to my parents-in-law for their love and support. My two little sisters, my sister-in-law, and my brothers-in-law always gave me a positive outlook and kept me going. Thank you all!



# Chapter 1

## Introduction

*December 31, 2012*

The much dreaded “fiscal cliff” was only hours away. The tax increase and spending cuts that would take effect from January 1, 2013 could potentially lead the U.S. economy to another depression. Both the Republicans and the Democrats in Congress agreed that it would be an absolutely undesirable situation. However, they were completely polarized on *how* to avoid this situation. In order to deal with the situation, a bipartisan group of senators had already been formed. It had four Democratic and four Republican senators and was called the “gang-of-eight” senators. Apparently, the key idea behind forming such a bipartisan group was to *intervene* the usual decision-making of the senators in Congress and lead the two diverging groups toward a consensus.

As the Wall Street and markets all over the world anxiously waited, several proposals seeking a middle ground came from different sources over the last few weeks, but none got accepted by the opposing party. Some of these proposals were labeled “joke” and some “magic beans and fairy dust.” As the clock kept ticking, a remarkably concerted effort by Senate Majority Leader Harry Reid (D-NV), Senate Minority Leader Mitch McConnell (R-KY), and Vice President Joe Biden was visible. At the eleventh hour, a deal was reached by Senator McConnell and Vice President Biden. It included a tax rate threshold of \$400,000 for individuals and did not include any significant spending cut.

Although a middle ground was reached, it did not please the members of either party. Democrats were seeking a tax rate threshold of \$250,000 and Republicans were opposing tax increase of any kind while favoring the direction of spending cuts. Nevertheless, Senator McConnell and Vice President Biden met with their respective party colleagues to ease the tension. Many, with reluctance, finally decided to side with the proposal.

*January 1, 2013*

In the small hours of the new year, the bill went to the Senate floor and was passed with a vote of 89–8.

The above story of avoiding the fiscal cliff connotes some very important scientific concepts. First, how *collective* outcomes like the above are reached has been studied by sociologists for decades. Second, it gives an anecdotal evidence of *influence* among the senators. It also suggests that although the underlying system for this influence is *large* and *complex*, there is a *network structure* in it, because clearly, not everyone directly influences everyone else. Third, the above story is all about *strategic interactions* in the sense that one’s action (whether to vote “yes” or “no”) depends on the actions of those that are “close” to him or her. Finally, it gives an evidence of *interventions* in real world—interventions by forming groups that would not naturally arise otherwise. As we will see shortly, this component of interventions is a fundamental element of cause and effect studies.

The goal of this dissertation is to study causal questions in strategic settings where a large number of entities interact with each other in a network-structured way. Marked by interventions, following are a few examples of causal questions in the Senate setting. How influential is a group of senators? Who are the most influential senators? Does there exist a small coalition of senators who can prevent filibusters?

In this dissertation, we will also study strategic settings that arise in networked economies, such as microfinance markets. In that setting, examples of causal questions could be the followings. What would happen if a loss-making bank is shut down? How can the government make loans more affordable by providing subsidies to the banks? What should be a sensible cap on interest rates? Note that many of these policy-level questions cannot be evaluated in practice before being implemented. However, if we could mathematically *model* the real world system, we could then evaluate such questions using this model. Therefore, our goal would be to model real world settings as complex systems and design algorithms for answering causal questions, but first, we will talk a little about causality in general.

## 1.1 Causality: A Contested Ground

Causality is one of the most natural quests of the human mind. Not only that it appears in abundance in our daily life, it also has a long history of scientific expedition, often embroiled in debates among statisticians, philosophers, economists, and computer scientists [56, 73, 97]. Such a level of contention

among researchers of diverse backgrounds—on one single topic—is rare and at the same time indicative of its scientific import and wide applicability. This dissertation presents a comprehensive framework for studying causality in strategic settings that often appear in social and economic networks. We call this framework *causal strategic inference*.

As mentioned above, causality has always been a highly contested ground. One beautiful example of this is a book edited by Daniel Little [73]. The chapters of the book pave the way for an enlightening back and forth debate between philosophers and economists. For example, in Chapter 2, philosopher James Woodward presents a causal interpretation of the structural equation models frequently used by economists. Woodward promotes the *manipulability theory* of causation as opposed to other alternatives, such as Granger’s notion of causation [48]. The manipulability theory resonates with our intuitive perception of causation. That is, if one variable causes another in a relationship, then changing or *intervening* the first variable (or other related variables) would provide a way of manipulating the latter. Now, an important question is: does the relationship remain stable while these interventions are being made? In the case of an *autonomous* relationship, the answer is yes. One example of an autonomous relationship is a law of the physics, such as the law of gravitation, which remains valid under a wide range of interventions. In contrast, non-autonomous relationships would break down easily under slight changes.

However, instead of thinking of relationships as simply autonomous versus non-autonomous, Woodward suggests the notion of the *degree of autonomy*, which corresponds to the range of interventions (perhaps limited) under which a relationship would remain stable. Woodward argues that this notion is particularly well suited for interpreting structural equation models. One significance of this is that it gives these models an explanatory power, as Woodward says, “autonomous relationships are causal in character and can be used to provide explanations.”

Later on, in Chapter 4, economist Kevin Hoover presents his view of causality in econometrics while contesting various points made by the authors of the earlier two chapters, including Woodward [73]. To a large degree, Hoover’s view concurs with that of Woodward. However, the two disagree on some of the fundamental issues, such as the explanatory power of a causal relationship. Hoover contends that “econometric models do not explain.” Hoover also contests many of the finer constructs, such as the meanings of “law” and “theory” implied by Woodward, in contrast to an econometrician’s interpretation of these terms.

The reason we brought up the debate between philosophers and economists

is two-fold. First, it gives a snapshot of the ever-contested topic of causality, which only highlights its importance across various disciplines. Second, it exposes a key component of the causality study in general—interventions or changes made to a system. We will illustrate how we incorporate interventions in strategic settings, but first, we will take a detour to some recent happenings in order to further signify the notion of interventions in causality studies.

## 1.2 Causal Probabilistic Inference

In recent times, one of the most celebrated success stories in the study of causality is the development of causal probabilistic inference during the 1990s [95–100]. Applications of causal probabilistic inference can now be seen in very diverse disciplines, such as economics, public policy, sociology, computer science, and various branches of life sciences, to name just a few. Given its emergence in wide-ranging application domains, it may at first be surprising to learn that the issue of causation has been swept under the rug for decades in classical statistics until 1935 when Sir Ronald Fisher’s seminal work on randomized experiments [38, 39] was published [97, p. 339–342]. Correlation, rather than causation, had been the prescriptive concept in statistics all those years. However, correlation alone does not directly answer questions such as: Does smoking *cause* cancer? Or, will increasing taxes *cause* the national debt to go down?

Judea Pearl, the recipient of the ACM Turing award in 2012 and one of the forerunners in the pursuit of studying causality in probabilistic settings, notes that the reason for this apparent neglect of causation in classical statistics is deeply rooted in the inability of probability theory to express causal statements [97, p. 342]. In particular, the language of probability theory is geared toward expressing *observational* inferences, as opposed to causal ones. In an observational inference, we may seek the probability of some events happening *given that* some other events have happened. Probability theory lays out a clear set of rules on how to express and manipulate such an inference question in order to give an answer to this question. In contrast to observational inferences where some events are observed (or given), causal inferences are accompanied by the mechanism of *intervention*. An example of an intervention is to set a random variable  $X$  (over which we have control) to a specific value  $x$ , which Pearl denotes by  $do(X = x)$  [97, p. 23]. An inference question in connection with this intervention would be to ask what would the probability of some events happening be once we perform this *do* operation.

On the surface, the causal inference question having a  $do(X = x)$  operation may seem to be very similar to an observational inference question where

$X = x$  is given. However, there are two notable differences. First, the *do* operation has the power to change the dependency structure among the random variables. Therefore, it can potentially change the joint probability distribution of the random variables. Second and as alluded above, probability theory, in its originality, cannot express this *do* operation mathematically. The study of causality in probabilistic settings has provided us with a mathematical framework that extends probability theory to express and process interventions.

In fact, Judea Pearl takes a broader view of causality than just the *do* operation. According to him, the study of causality can be hierarchically organized in three natural types of queries with increasing levels of difficulty: predictions, interventions, and counterfactuals [97, p. 38]. First, prediction is the type of query where we observe something about the “system,” and taking that observational knowledge into account, we are asked to infer something else that we did not observe. The important aspect in prediction is that we are not allowed to change anything in the system. Changing something in the system, which Judea Pearl often refers to as *surgery*, is permitted in the second type of causal query—interventions. An example is the the *do* operation in probabilistic settings, as outlined above. Counterfactuals, the third type of causal queries, are the most challenging ones in the sense that we are given some observation about the system and asked to infer the outcome of the system if the opposite of that observation, in some sense, were to take place.

The goal of this dissertation is to study causal inferences in game-theoretic settings at the second level of queries, interventions. Since game theory reliably encodes strategic interactions among a set of players, we will call this type of inference *causal strategic inference*.

### 1.3 Causal Strategic Inference

A game in non-cooperative game theory can be described by a set of *players*, a set of *actions* for each player, and a *payoff function* for each player that maps each *joint-action* to a real number. Here, a joint-action is specified by an action for each player. A central solution concept in non-cooperative game theory is Nash equilibrium. A pure-strategy Nash equilibrium (PSNE) can be defined as a joint-action of the players such that every player *plays its best response* to the other players’ actions simultaneously.<sup>1</sup> Here, a best-response of a player to the the other players’ actions is defined by an action that maximizes its payoff

---

<sup>1</sup>There is a more general solution concept known as mixed-strategy Nash equilibrium (MSNE) where each player independently chooses a probability distribution for playing one of its actions so that every player maximizes its *expected* payoff simultaneously.

with respect to the corresponding joint-action. Of course, a best-response of a player may not be unique. Yet, the definition of Nash equilibrium signifies *stability* in the sense that no player has any incentive to unilaterally deviate to a different action. (Consider any pure-strategy Nash equilibrium, specified by a joint-action. If a player could indeed increase its payoff by playing a different action, then the player is not playing its best response in this joint-action. Therefore, this joint-action cannot be a pure-strategy Nash equilibrium.)

As mentioned earlier, our goal is to study interventions, the second type of causal queries, in game-theoretic settings. Interventions are carried out by surgeries. Therefore, one rudimentary question is what types of surgeries we would allow in the context of a game-theoretic setting. Put differently, what is the analog of the *do* operation described above in a game-theoretic setting?

Although game theory explicitly represents the actions of the players, it is different from actions (e.g., the *do* operation) or interventions in the context of causal probabilistic inference. In game theory, actions are adopted or played by the players of the game, who are integral parts of the system. However, in causal probabilistic inference, interventions are performed by someone outside the system, such as an experiment designer. Notably, an intervention involving a set of random variables can be performed if one has control over them. One implication of such an intervention is that it makes *changes* to the original system (and hence the name surgery). Therefore, causal probabilistic inference is a query concerning the *changed system*, although the given input is with respect to the *system before the changes were made*. We will formulate causal strategic inference in an analogous way, but first, we will answer the question regarding the types of surgeries we perform. We will consider the following two possibilities.

### **Surgery 1: Setting the Actions of Some of the Players**

One way of doing a surgery on a game is to restrict the actions of some of the players. For example, we can set the actions of some of the players to particular ones. Now, the question is: how should we interpret this surgery? For example, suppose that player  $i$ 's action has been set to  $a_i$ . Should we modify the game in a way that player  $i$ 's best response is always  $a_i$ , no matter what the other players play? Or, should we keep the game unchanged and rather focus only on those equilibria (if any) where player  $i$  plays  $a_i$  as its best response? Let us consider these two different interpretations.

First, once we set the actions of some of the players, we can modify the original game in the following way. Consider any player  $i$  whose action has been fixed to  $a_i$ . The payoff function of player  $i$  is changed so that  $i$ 's payoff is 1 for any joint-action in which  $i$  plays  $a_i$  and 0 for all other joint-actions.

This change makes sure that the preset actions are indeed best responses of the corresponding players in the modified game (with respect to any action that the other players play). Note that a Nash equilibrium of the modified game may not be a Nash equilibrium of the original game. In particular, the players whose actions have been set may not be playing their best response with respect to the original game. Therefore, the outcome of the modified game is not guaranteed to be stable with respect to the original game. Such an approach has been used in a line of work on finding the most influential nodes in a social network [67], where the goal is to maximize the spread of a “new” behavior (e.g., buying a new product) by selecting a small “seed” set of initial adopters. The underlying mechanism is to set the actions of the seed players to the one denoting the adoption of the new behavior and then let the diffusion process set off. At the end of the diffusion process, however, the seed players may not be playing their best response *with respect to the original game*.

Controlling the actions of some of the players has also been used in other settings. For example, in the setting of network routing games, Sharma and Williamson study the minimum number of users that need to be controlled by a central authority to improve the social welfare of a Nash equilibrium [111]. This is motivated by the case where there are two types of users of a network application: premium users with the privilege of choosing their own route of traffic and ordinary users, who must go by whatever the network administrator has chosen for them [104]. The general problem is to find the minimum fraction of users to be controlled by the network administrator to achieve a desirable objective. Here, being controlled by the network administrator, the ordinary users might be forced to adopt an action that they would not have adopted otherwise.

Second and in contrast to the above point, we can also do interventions by “controlling” the actions of some of the players *without changing the game*. For example, after setting the actions of some of the players, we can ask questions regarding the stable outcomes (e.g., Nash equilibria) where these players play according to the preset actions. An example of an inference question in this approach is to ask how many stable outcomes could possibly result from setting the actions of a subset of the players. We will adopt this notion of intervention in studying causal strategic inference questions in a social network setting in Chapter 2.

## **Surgery 2: Changing the Structure of the Game**

In this type of surgery, we change the game without setting the action of any of the players. Note that the notion of “changing the game” is very much

open ended. It can potentially mean changing the payoff function of a player, removing a player from the game, adding a new player to the game, changing the set of actions of a player, as well as any combination of these. Here, we narrow our focus to the *structure* of a special type of games, namely graphical games in parametric forms (as opposed to normal forms). We use the term structure in this context to refer to the underlying topology of the game. One example of an intervention by changing the structure of a game is to remove a player from the game. A causal strategic inference question under this type of intervention is to infer how the outcome of the game would change due to the intervention.

We study such interventions in Chapter 3, where we model a microfinance market in a game-theoretic way in order to ask causal strategic inference questions, such as what would happen if some of the loss-making government-owned banks are shut down? To answer such a question, we first learn the parameters of the model from real-world data and compute an equilibrium point, which reflects the outcome *before* the removal of any bank. We then do an intervention by changing the structure of the game (i.e., removing the loss-making government-owned banks). After that, we compute an equilibrium point of the resulting game (note that after removing the banks, we *do not* go back and learn the parameters of the model again). The difference in equilibrium outcomes before and after the removal of the banks gives us the desired answer.

## 1.4 Causal Strategic Inference: A Comparative Study

We should first clarify that the idea of interventions in strategic settings is not new. Some of the surgeries we have mentioned above have been studied before in the context of various application scenarios. However, our main objective here is to build a comprehensive framework that wraps around this idea of interventions in strategic settings.

Our particular focus will be on those systems where individuals or entities exhibit *strategic interactions* and affect each other in a *complex* but *network-structured* way. The proposed framework of causal strategic inference is composed of the following components: mathematically *modeling* a complex system, *learning* the parameters of the model from *real-world data*, and designing *algorithms* to predict the effects of interventions. We will now review earlier research relevant to these components, especially from economics literature (related computer science literature will be reviewed in Chapters 2 and 3). As we will see, one of the distinctive features of our approach is our focus



on *computation* (e.g., designing algorithms for equilibrium computation).

### 1.4.1 Causal Strategic Inference in Econometrics

Modeling strategic scenarios and estimating the parameters of the model are active research topics in econometrics. Here, we will conduct a brief review of the literature with the goal of illustrating the difference between our approach and the general approach in econometrics. Since we will not give any detailed specification of our model, the discussion will be at a high level.

In econometrics literature, the flagship application scenario for studying strategic decision-making is the setting where two or more firms simultaneously decide whether to enter a market or not. This decision is strategic, because a firm's decision and hence its expected profit depend on the decisions of the other firms. As a result, game theory has been the prescriptive tool for modeling entry decisions in econometrics.

Within the general game-theoretic framework, there is a variety of entry-decision models in econometrics capturing homogeneous versus heterogeneous firms, complete information versus private information settings, and static versus dynamic games. The literature also shows different ways of addressing some of the inherent issues like the multiplicity of equilibria and equilibrium selection. There is, however, one unifying theme in the literature: almost all of the models are based on the discrete choice model [77, 78]. The discrete choice model in its originality does not allow the utility of an entity to depend on the actions of the other entities. However, the main ingredient of a game-theoretic model is this interdependence of actions. To account for this, the econometrics models that we will review indeed extend the original discrete choice model to what the literature commonly refers to as the *discrete game model*.

#### Bjorn and Vuong's Model of Labor Force Participation

The first discrete game model is attributed to Bjorn and Vuong, who studied the case of simultaneous decision-making by a husband and a wife on whether to enter the labor force or not [17]. This is a two-player game, where each player has two actions. Denoting the action of player  $i \in \{1, 2\}$  by  $x_i \in \{0, 1\}$  (0 denotes not entering and 1 entering the labor force) and that of the other player by  $x_{-i}$ , the payoff of  $i$  is defined using McFadden's random utility model [77] as follows.

$$\tilde{u}_i(x_i, x_{-i}) = u_i(x_i, x_{-i}) + \eta_i(x_i, x_{-i}). \quad (1.1)$$

The above payoff function consists of an observed, deterministic part  $u_i$  and an unobserved (to the researcher), random part  $\eta_i$ . The first accounts for observed attributes of the players, such as age, education level, assets, number of kids, etc. The latter part accounts for factors that the researcher could not observe or did not model, and as a result, it appears as a random shock. Every player  $i$  observes its own random part  $\eta_i$ , but depending on whether it also observes the other player’s random part, the game becomes either a complete information or an incomplete information game, respectively. The model of Bjorn and Vuong is a complete information one. Here, the best response  $x_i^*$  of player  $i$  can be written as follows.

$$x_i^* = 1 \iff \tilde{u}_i(1, x_{-i}) - \tilde{u}_i(0, x_{-i}) > 0. \quad (1.2)$$

In other words, player  $i$ ’s best response is to choose the action that maximizes its payoff with respect to the other player’s action. A pure-strategy Nash equilibrium (PSNE) is given by  $(x_1^*, x_2^*)$  such that both of players are best responding to each other simultaneously.<sup>2</sup> Obviously, there could be three possible types of outcomes in this game: a unique PSNE, multiple PSNE, and no PSNE at all. Bjorn and Vuong view the data as a unique equilibrium. However, if the latter two possibilities of multiplicity and non-existence are ruled out, then they show that the model no longer remains strategic (that is, the best response of one of the two players does not depend on the other player’s action). It rather becomes equivalent to a previously studied *simultaneous equations model with structural shift* where a certain “logical consistency condition” must hold [54, 110]. Now, if the option of the multiplicity of PSNE is kept on table and if the data is viewed as a unique PSNE, an important question is: which one of the multiple possible PSNE is “played” in the data? This is typically known as the *equilibrium selection problem*.

Bjorn and Vuong take a randomized approach to this problem of equilibrium selection. They assign probabilities to all possible pairs of the reaction functions<sup>3</sup> of the two players and express the probability of each PSNE in terms of these probabilities. They then use the maximum likelihood technique to estimate the probability of observing any particular PSNE, along with the

---

<sup>2</sup>Equation (1.2) shows that the best response of a player depends on the *difference* between payoff functions and hence on the difference between the corresponding random parts of Equation (1.1). Along with other simplifying assumptions, Bjorn and Vuong’s main assumption is that this difference between the random parts is a standard normal distribution with a correlation between the two players.

<sup>3</sup>For example, the husband’s reaction function could be one of the followings: choosing action 1 all the time (no matter what the wife has chosen), choosing 0 all the time, choosing whatever the wife has chosen, and choosing the opposite of what the wife has chosen.

other parameters of the model that are not detailed here. They also give a necessary and sufficient condition for these parameters to be *identifiable*, which means that if that condition holds, then for any outcome of the model, the estimated parameters are unique. In other words, different instantiations of the parameters cannot generate the same outcome. In general, identifiability of the parameters is a major focus in econometrics literature.

Being the first of its kind, Bjorn and Vuong’s model is simplistic and does not scale well if the number of players is increased. For instance, if there is a large number of players, then assigning a probability to each possible combination of the reaction functions of all the players would be computationally expensive. As we will see, much of the later literature actually avoids combining the reaction functions of the players.

### Entry Models of Bresnahan and Reiss

Bresnahan and Reiss investigate entry in monopoly markets using two types of models: a simultaneous, game-theoretic model and a sequential decision making model [21]. Their empirical study is based on the markets of automobile dealers with a focus on how market sizes influence entry decisions and whether the second entrant faces entry barrier (i.e., whether the fixed cost and market opportunities for the second entrant are less favorable compared to the first entrant).

This is again a two-player, two-action setting. We will first give a brief overview of the simultaneous-move model of Bresnahan and Reiss. Suppose that  $\tilde{u}_i^M$  and  $\tilde{u}_i^D$  are the payoffs of firm  $i$  in a monopoly and a duopoly market, respectively. We will not go into the details of these payoff functions, but as in the Bjorn-Vuong model, they also comprise of two parts: an observable part and an unobserved, random part. The random part is observed by both the players, although not by the researcher. The entry decision  $(x_1^*, x_2^*)$  would be a pure-strategy Nash equilibrium (PSNE) if and only if the following best response condition holds for all firms  $i = 1, 2$ .

$$x_i^* = 1 \iff (1 - x_{-i}^*)\tilde{u}_i^M + x_{-i}^*\tilde{u}_i^D > 0.$$

Given a particular model, a PSNE outcome could be one of the following five types: monopoly by firm 1, monopoly by firm 2, duopoly, no entrant, and finally, monopoly by either firm 1 or firm 2 (but not duopoly). Once again, the multiplicity of PSNE is deemed challenging, as the authors say, “The presence of non-unique equilibria in game-theoretic models makes it impossible to use standard qualitative choice models to model entrants’ profits.” This is so,

because the data is viewed as a unique PSNE. As Bjorn and Vuong showed earlier, restricting the model to rule out the multiplicity of PSNE results in a model that is not strategic any more. Therefore, the equilibrium selection problem arises inevitably.

Regarding the problem of equilibrium selection, Bresnahan and Reiss observe that if the model is reinterpreted to predict the *number of entrants* instead of the identity of the entrants, then the PSNE outcome is always unique. Another approach to avoid the multiplicity of PSNE is to consider a sequential-move version of the model. It is easy to show that if the firms do not make their decisions simultaneously, then the outcome is unique. The estimation of the model parameters using spatially isolated rural automobile dealership markets shows that the second entrant is not subjected to entry barrier and that its entry does not cause the first entrant's profits by much. In many cases, this is due to the market size being already very big when the second firm enters the market.

In a related paper, Bresnahan and Reiss discuss the issues of the existence and the uniqueness of PSNE in discrete game models [22]. They also discuss how one could deal with mixed-strategy Nash equilibria (MSNE) in discrete game models and how these models can be extended to the cooperative games setting. Although they motivated the issues of MSNE and cooperative games using real world examples, they did not actually apply their ideas to any empirical setting. For instance, they say that "the researcher must exercise care when selecting [certain probability] distributions," which needs to be done on a case by case basis if we would like to consider MSNE.

### **Berry's Model of Entry in Airline Markets**

Airline markets, each consisting of a source-destination pair of cities, have been studied by economists from different points of view. A common example is various explanations of an airline's profit due to its hub and spoke network [72]. Berry took a different approach to studying an airline market, by investigating the effects of strategic entry decisions of an airline on the profitability of the flights in a market [16]. To model the entry decision of an airline in a market, Berry presents a discrete game model that allows for a *large number of heterogeneous* airlines. Apart from the specifics of the model, this is one of the key differences with the previous entry models, such as the ones by Bresnahan and Reiss. Heterogeneity among the airlines can be observed in terms of their flight networks, fleets of aircrafts, etc. Heterogeneity can also be due to unobserved factors. In Berry's model, the payoff of an airline  $i$  in a market  $k$  is defined as follows.

$$u_{ik}(\mathbf{x}_k) = v_k(f(\mathbf{x}_k)) + \phi_{ik}.$$

Here,  $\mathbf{x}_k$  is the vector of the entry decisions of each airline  $i$  in a market  $k$ ,  $x_{ik} \in 0, 1$  denotes an airlines entry decision. The first term in the payoff function is market specific and captures the competitive effect due to the entry decisions of the airlines, and the second term is specific to the airline-market pair and is treated as a single index of profitability. In order to guarantee the existence of a PSNE and the uniqueness of the *number* of PSNE (to deal with the equilibrium selection problem as mentioned above), Berry imposes several assumptions. First, the airlines in a market  $k$  can be sorted according to the profitability index  $\phi_{ik}$ , and this ordering is independent of the entry decisions. Second, the function  $f(\mathbf{x}_k)$  in the first term is defined as the count of entrants, i.e.,  $f(\mathbf{x}_k) = \sum_i x_{ik}$ . Third, the market-specific function  $v_k$  is decreasing in the number of entrants.

As before, each of the two terms in the above payoff function is further decomposed into two parts: one observable part and one unobserved, random part. Again, the setting here is a complete information game. The main challenge in analytically characterizing the probability of a certain number of entrants is due to the large number of airlines, which contributes to an exponential number of integrations over the random parts. Berry proposes two directions to address this. The first one is to impose additional restrictions on the model, such as removing the part of strategic interaction from the payoff function (i.e., an airline’s profit is not affected by the number of entrants). The other direction is to apply simulation estimators [79]. The estimated model shows a strong negative influence of competition on an airline’s profit, which can limit the effectiveness of a policy encouraging the potential entrants.

As Berry points out, his entry model is guided by a “partial equilibrium approach,” where instead of considering an airline’s network of flight routes, the analysis focuses on a pair of source-destination cities. However, we know that the network structure of an airline’s flight routes is one of the most important ingredients of its operation and profitability. In our view, the major challenge in accounting for this network structure is due to the analytical approach to the problem. An alternative to deal with this would be an *algorithmic approach*, which we pursue in this dissertation (although not on the same problem).

Berry’s model has been subsequently extended by others. Ciliberto and Tamer allow a general form of heterogeneity among the airlines that no longer guarantees a unique number of entrants in all the PSNE of the (complete information) game [27]. Without assuming a particular equilibrium selection rule, they *bound* the choice probabilities between an upper and a lower limit. They

estimate the parameters of the model by minimizing the distance between the *set* of choice probabilities between these two bounds and the probabilities estimated from the data. Apart from modeling the airline industry, a different work by Berry et al. presents techniques to estimate the parameters of an oligopolistic market with a wide range of product differentiation with applications to the U.S. automobile market [15].

## Other Models

The three early models that we reviewed above exhibit some of the key aspects of the general econometrics approach to modeling strategic scenarios, such as the adoption of a random utility model [77, 78], analytical characterizations of some of the quantities of interest, a way of dealing with the multiplicity of PSNE (for example, a randomized equilibrium selection mechanism) or a way of avoiding the multiplicity issue altogether (for example, by imposing additional assumptions that would lead to a unique equilibrium or by reinterpreting the model to predict a common property of all PSNE, such as the number of players playing a particular action), and the identifiability of the parameters of the model. The literature has since been enriched with a number of interesting pieces of work extending the previous research as well as injecting new ideas to deal with these challenging tasks. Here, we will briefly review a sample of some of the widely cited research along this line.

## Seim’s Model of Product Differentiation

The early game-theoretic entry models, such as the one by Berry [16], focus mostly on firm-specific profits and the competitive effects of multiple entrants, but do not model product differentiation by the firms. Seim proposes a discrete game model of entry decisions that allows the firms to spatially differentiate their products by choosing, for example, a location of operation. Her empirical study is based on the location choice of video retailers. In contrast to much of the earlier work, she models entry decisions as a game of *incomplete* information, which accounts for a firm’s lack of information about many of the characteristics of another firm, such as that firm’s managerial talent. Interestingly, the reason why many of the earlier models were games of *complete* information is that the incomplete information version was thought to be more challenging. For example, Bresnahan and Reiss say, “Games of private information pose much more complicated estimation issues” [22, p. 60]. However, it later turned out that an equilibrium in an incomplete information game (also known as the Bayes-Nash equilibrium) can be characterized more easily [106] and that the estimation can also be done in a straightforward two-step method

[7]. As a result, modeling a scenario as an incomplete information game often serves the dual purpose of modeling various unobserved idiosyncrasies as well as dealing with that model in a tractable fashion.

Going back to the model of Seim, a market consists of multiple locations, and each firm chooses an entry location if it decides to enter the market. The choices of the firms are made simultaneously. The payoff function of a firm consists of the following terms: an observable location-specific characteristics (such as the population and the income level of the potential customers), an unobserved market-specific random term, a competitive effect term that accounts for the decisions of all the firms, and an unobserved firm and location-specific term that captures the firm's private information about its profitability in that location. The last term is assumed to be independently and identically distributed (iid) draws from a type-1 extreme value distribution, which leads to closed form expressions for a firm's choice probabilities. The goal of the model is to predict the unique number of entrants in a PSNE, although Seim also shows that the PSNE itself is unique under certain additional assumptions. The estimated model shows that video retailers use location choice to their competitive advantage. Also, as the market size increases, the local demand decreases due to the spreading out of the population density. As a result, the number of entrants does not increase by much.

### **Augereau et al.'s Model of Technology Adoption**

Beyond entry decisions, discrete game models have been designed for many other interesting phenomena, such as adoption of a particular technology. Augereau et al. study the adoption of the 56K modem technology by the ISPs in a market during the late 1990s [4]. At that time, there were two competing and incompatible implementations of the 56K standard, one by the U.S. Robotics and the other by Rockwell. If an ISP adopts the U.S. Robotics technology, for example, then its customers must also buy the U.S. Robotics modems to enjoy a high-speed connection. Augereau et al. model the choice of the ISPs as a discrete game of incomplete information, which accounts for the characteristics of the market as well as the ISPs and the simultaneous decision of the ISPs. They show that the ISPs in a market want their choice to be different from their competitors (so that they do not lose their customers to their competitors).

### **Sweeting's Model of the Timing of Radio Commercials**

Whereas Augereau et al.'s result on technology adoption among ISPs can be interpreted as a coordination failure, Sweeting's model of radio stations' choice

of timings for playing commercials tells us the opposite story of coordination and synchronization. In Sweeting’s model, the radio stations in a market choose timings for advertisement strategically with their payoff function capturing market and station specific factors, a competitive factor that takes the form of the proportion of the other stations that choose the same time, and a private information term modeled as a random shock. Sweeting shows how the multiplicity of equilibria can help the identification of the parameters of the model. Estimation is done by the two-step method of Bajari et al. mentioned above [7]. The finding of coordination among the radio stations signifies that the interests of the radio stations are somewhat aligned with the interests of the advertisers. During drivetime hours, the coordination incentive is very strong. Multiplicity of equilibria is also more common during that time.

### **Discrete Game Models of the Banking Sector: ATM Networks**

Game-theoretic models have also been developed to capture decision making in the banking sector. Consider the case of ATM networks for example. From our daily experience, the ATM networks of different banks are incompatible unless we pay a surcharge. This surcharge never covers the cost of the ATM service of a bank. It is rather intended to attract customers to open deposit accounts at the bank. As a result, larger banks often charge more surcharge than smaller banks and credit unions.

Ishii models banks and customers as strategic decision makers in a two-sided market to understand the effects of this surcharge [57]. In brief, the payoff function of a customer for choosing a bank (i.e., having a deposit account in that bank) in a market accounts for the customer’s observable characteristics, the bank’s observable characteristics (e.g., its number of ATMs), the bank’s interest rate (which is determined by the bank strategically), and the unobservables corresponding to the customer and the bank. The banks, on the other hand, maximize their profits in two stages. In the first stage, each bank strategically chooses the number of ATMs to be deployed. In the second stage, it chooses an interest rate to maximize its profit given a PSNE from the previous stage. We will not go into the details of these two stages. The estimation is done by the generalized method of moments (GMM) method. The estimated model captures the phenomenon that when choosing a bank, customers are influenced by the bank’s ATM network size and its surcharge. It also shows that the revenue from the ATM service does not cover its cost. Rather, the incentive for a bank to invest in an ATM network lies in securing a share of the deposit market.

An interesting feature of Ishii’s work is her study of various counterfactuals. For example, what would happen if the surcharges are eliminated by law? In



that case, the model predicts that the market becomes less concentrated. That is, the market share of customer deposits is reallocated from the larger banks to the smaller ones. In response, the larger banks raise the interest rates on deposit, which decreases their profit. In fact, the overall profit of the industry decreases due to the elimination of surcharges. To the contrary, the presence of surcharges encourages banks to expand their ATM networks, although it makes the market share of customer deposits concentrated at the larger banks.

### **Discrete Game Models of the Banking Sector: Adoption of ACH**

Game-theoretic models of decision making have also been developed to study other phenomena in the banking industry, such as the adoption of automated clearinghouse (ACH) technology, which provides an electronic equivalent of paper checks and is commonly used in direct deposits and automated bill payments. Akerberg and Gowrisankaran estimate the magnitude of network effects in ACH adoption [1]. The model has two sides: banks and customers (e.g., small businesses), where the customers are treated as homogeneous. Two banks can do an ACH transaction if both have already adopted this technology. On the other hand, there are two alternatives for a customer. In a one-way transaction, a customer may receive ACH payment without adopting the technology, provided that her bank has adopted it. In a two-way transaction, a customer must also adopt the ACH technology in addition to her bank.

Akerberg and Gowrisankaran define a two-stage game. In the first stage, the banks simultaneously decide whether to adopt ACH or not, and in the second stage, the customers decide on the adoption of ACH, given the decision of the banks. This model does not rule out the multiplicity of equilibria. The way the authors approach equilibrium selection is by estimating the probability of seeing one of the two extreme equilibria (Pareto-best and Pareto-worst), which is obviously a simplification compared to considering all possible equilibria. The estimation of the parameters is done by a simulation method. The estimated model is used for counterfactual policy experiments. The model suggests that government subsidy directed toward the customers is more effective for ACH adoption than that directed toward the banks.

### **Recent Developments**

The subject of discrete game models is an active area of research in econometrics. Recently, Bajari, Hong, and Ryan presented methods for identification and estimation of discrete game models of complete information, which are also applicable to general normal-form games [9]. A key feature of their work is that they estimate both the parameters of the model and the equilibrium

selection mechanism. For the latter, they compute all the equilibria of a game using the algorithm of McKelvey and McLennan [81]. This is certainly a computationally expensive step for large games. In a separate work, Bajari et al. address estimation of discrete game models of incomplete information [8]. Finally, an excellent survey of some of the recent results on the variants of discrete game models (e.g., complete vs. incomplete information, static vs. dynamic games) and their identification, estimation, and equilibrium selection has been presented by Bajar, Hong, and Nekipelov [7].

### 1.4.2 Causal Strategic Inference: Our Approach

Even though we study a completely different set of problems than the ones reviewed above, we share the most important ingredient of strategic decision making with all these problem. However, there are some fundamental differences between our approach to studying strategic settings and that of econometricians in general. Once again, since we have not formally defined any of our models, this discussion will not focus on any detail. It will rather highlight these key differences in the context of causal strategic inference.

#### Analytic vs. Algorithmic Approaches

First, an econometrics approach to dealing with strategic settings is in large part analytic. True, econometricians do provide algorithms (e.g., algorithms for estimating parameters), but for the most part, those algorithms are primarily driven by analytic techniques. See, for example, the two-step estimator for discrete games of incomplete information [7]. In contrast, the main focus of this dissertation is an algorithmic approach to problems. For example, one of the main objectives of Chapter 2 of this dissertation is to design efficient algorithms for equilibrium computation that exploits the special structure of the problem.

Note that we do not claim that an algorithmic approach is better than an analytic approach. However, in certain situations, an algorithmic approach might provide a good alternative to an analytic approach. This is particularly the case when we have large, complex systems with an underlying structure. For example, the network structures of an airline’s flight routes are not often exploited in the econometrics models of airline markets [16]. This could be due to the challenge posed by dealing with a large, heterogeneous system in an analytic manner. As mentioned above, such complex systems are exactly the focus of this dissertation. In Chapters 2 and 3, we will show that looking through an algorithmic lens helps us solve problems that would otherwise be impossible to manage analytically.

## Modeling

A common modeling approach in the econometrics literature that we reviewed is to adopt of the random utility model [77, 78]. Besides giving a reasonable way of modeling unobservables, this approach sometimes also leads to simple closed form expressions that are easy to deal with. For example, the choice probabilities in a discrete game model can be expressed that way when the random parts of the payoff functions are iid type-1 extreme value distributed.

In contrast, the modeling approach in this dissertation is completely different. In Chapter 2, we model influence among the heterogeneous individuals in a social network using a graphical polymatrix game [60, 66], where the interpretation of the payoff functions, which do not contain any random term, is rooted in a well-studied and widely adopted sociology model [49]. Furthermore, in Chapter 3, we model two-sided microfinance markets using the well-studied concept of abstract economies in classical economics [3, 30]. Our model of microfinance markets is network-structured, and a very special case of it can be shown to be one type of Fisher market [37], which has been a subject of intense algorithmic study by computer scientists in recent years.

Again, we do not claim that our models are “superior” to the random utility model in any sense. Rather, with the specific applications that we would like to address here, our modeling approach serves the purpose best while having its root in the relevant social science literature. The key aspects of our models are heterogeneity, network-structure, compact representation, and the ability to capture the strategic interactions among a large number of entities.

## Estimation

Econometrics and computer science (machine learning, in particular) have diverging views on the issue of the estimation of parameters. As we saw in the literature review above, the identification of the parameters of a model (or some function of the parameters) is a major concern in econometrics literature. The reason is that if the parameters are not identified (i.e., different instantiations of the parameters lead to the same outcome), then the estimated model may be very different from the actual system that generates the data, even though the estimated model produces almost the same outcome as the actual system. To ensure identification, additional restrictions are sometimes imposed on the model. Identification with infinite sample is also very common [8, 9].

In contrast, one of the most primitive principles of machine learning is Occam’s razor, which says that if multiple models explain the same observation reasonably well, then we should choose the “simplest” model. Compared to a

complex model, a simpler model usually also shows more generalization power in terms of predicting something that has not been observed before. However, too much simplicity might not capture the observed data well. Therefore, researchers often strike a balance between simplicity and complexity, which is usually guided by the well-known bias-variance tradeoff (a high bias corresponds to too much simplicity and a high variance corresponds to complexity). One technique often employed for this is to estimate the parameters of a model using a portion of the data, not the whole, and then to test its predictive power using the rest of the data. At the end, the best predictive model is chosen.

Estimation in econometrics, on the other hand, is very different from that in machine learning. In econometrics, all the available data is used for the purpose of estimation, and a high variance is a desirable objective. The anxiety about whether the estimated model would perform well in an unforeseen environment is eased with the assurance that the parameters have been identified. However, as we mentioned above, identification often necessitates making strong assumptions.

There are many other contested issues between these two disciplines, which are out of scope for this dissertation. It is not that one of the approaches is good and the other is bad. It is just that they are different. In this dissertation, we have taken a machine learning approach to estimation. In Chapter 2, the graphical structure of influence among the U.S. senators and various related parameters are estimated using the method of Honorio and Ortiz [55]. In Chapter 3, the estimation is done using a bi-level optimization program. In each case, the predictive power of a model with respect to unforeseen events has been the prime focus.

The objective of our estimation is also different from that of the fast growing literature on causal estimation in computer science and statistics. For example, a common technique for understanding the effects of new product features on consumers is known as bucket testing, which basically exposes the feature to a random sample of the population and measures its effect on them. With the advent of online social networks like Facebook, bucket testing can no longer focus on a disconnected random sample of users. It also needs to consider the network structure, because in the context of online social networks, the effect of a new feature is more meaningful when a user as well as some of her friends are exposed to it, compared to only the user being exposed in isolation.

To extend bucket testing to networked settings, Backstrom and Kleinberg propose a graph-theoretic sampling technique that addresses these two competing requirements: samples need to be uniformly random and they also need to be well-connected [6]. Along the same line, Ugander et al. propose graph

cluster randomization techniques to give an efficient algorithm to compute the probability of the exposure of a user [117]. They also show that their techniques can lower the estimator variance. In another notable work, Toulis and Kao propose two techniques for estimating causal peer influence effects—a frequentist approach that can deal with more complex response functions and a Bayesian approach that provides more accurate estimates under network uncertainties [114]. In contrast to this line of work, our goal is not causal estimation. We rather want to estimate models that capture strategic interactions.

### Equilibrium Selection

Almost all of the models we reviewed above exhibit multiplicity of equilibria. Therefore, the question of equilibrium selection naturally arises as the data is often viewed as a single equilibrium. Econometrics literature suggests three main ways of dealing with the equilibrium selection problem [7, 9]. First, the probability that an equilibrium, which is generated by the model, is observed in the data is estimated [9, 17]. Sometimes, instead of considering all possible equilibria, only a few equilibria are considered in this probabilistic approach [1]. Second, the model can sometimes be reinterpreted to give a unique outcome, even if there are multiple equilibria. A typical example is considering the number of entrants instead of the identity of entrants in an equilibrium of an entry market [16, 21]. Third, the choice probabilities can sometimes be bounded between two limits, which guides the selection of an equilibrium [27].

When we study influence among the U.S. senators in Chapter 2, it would naturally appear that the multiplicity of equilibria is a desirable objective. This is because if we model the voting outcomes from Senate as equilibria of a game, then even though these are generated by the same 100 senators, these are not the same. Therefore, the estimation algorithm of Honorio and Ortiz tries to capture a *set* of equilibria rather than one particular equilibrium [55]. They choose a model that maximizes the number of observed voting outcomes that are captured as the equilibria of the model and minimizes the number of unobserved outcome captured as equilibria.

In our study of microfinance markets in Chapter 3, our model can potentially generate multiple equilibria. We select one of these equilibria that is geometrically closest to the observed data. This equilibrium selection mechanism is embedded in the parameter estimation procedure. We also test for the robustness of this mechanism. We find that even if we introduce considerably large magnitudes of noise in the data, this mechanism selects the same equilibrium.

## Interventions

A key component of causal strategic inference as well as causality in general is interventions. All of the econometrics studies reviewed above concern two of the components of causal strategic inference that we mentioned earlier: modeling a strategic scenario and estimating the parameters of the model. However, many of these studies do not perform interventions. There are, of course, exceptions. For example, Isii studies the effect of removing the ATM surcharges [57]. Akerberg and Gowrisankaran show the comparative effects of subsidies to customers and banks on ACH adoption [1].

The main focus of this work is a wide range of interventions. In Chapter 2, we identify a most influential set of individuals in a social network by performing interventions. We also study the “powers” of groups of individuals using interventions. In Chapter 3, we perform various interventions in a microfinance market, such as setting an interest rate cap, removing a bank from the system, providing subsidies to certain banks to make loans more affordable, etc. It should be mentioned here that interventions by removing players is not a new concept. Ballester et al., for example, performed interventions in a criminal network by removing players from it [12].

## 1.5 Organization

This dissertation is divided into two main parts. In the first part (Chapter 2), we model influence in a networked population with the ultimate goal of performing various types of causal strategic inference. Our model is rooted in the widely used threshold models from the mathematical sociology literature [49]. However, instead of the traditional contagion approach, we take a non-cooperative game-theoretic approach that captures strategic aspects of the network. In the second part (Chapter 3), we model microfinance markets in order to study causal questions without the privilege of performing trial-and-error experiments. Again, our model is game-theoretic as we view stable outcomes of the market as equilibrium points. Following is a brief outline of these two main chapters.

### Outline of Chapter 2: Causal Strategic Inference in Social Networks

In Chapter 2, our approach to influence in social networks is essentially a causal one (in a game-theoretic setting). This will become particularly clear if we consider our definition of the most influential individuals in a social network in Sections 2.2.3 and 2.2.2. One instance of the general definition of the most influential nodes is the following. A set of nodes is the most influential with

respect to a desirable outcome (e.g., passing a bill in Senate by 100–0 vote) if their adoption of the desirable behavior (e.g., voting “yes”) *causes* everyone to also adopt that desirable behavior as the mutual best response. That is, a set of the most influential nodes *leads* a complex system to a unique stable outcome that is the same as the desirable outcome.

Another example of causal strategic inference within our study of influence among senators is the following. Does there exist a “small” set of senators who can prevent a filibuster? Again, we interpret this question in a causal way: we are interested in finding a small set of senators whose coalition would *cause* at least 60 senators to vote “yes” in any stable outcome, thereby avoiding a filibuster situation.

Chapter 2 is organized as follows. It begins by relating causal strategic inference to our study of influence in social networks. Section 2.1 gives a very high-level outline of our model and puts it in the context of the existing body of literature in sociology as well as computer science. Section 2.2 defines our model of *influence games* and formulates the most influential nodes problem. The next two sections explore various computational problems in influence games, including the related computational complexity questions. Section 2.5 presents empirical results using both synthetic experimental data as well as the real-world data obtained from congressional voting records and Supreme Court rulings. Appendix 2.A gives a brief exposition of the collective actions and collective behavior literature in sociology, which would provide a broader view of the topics we study in Chapter 2.

### **Outline of Chapter 3: Causal Strategic Inference in Economic Networks**

In this chapter, we develop a model of microfinance markets as a two-sided networked economy. In this model, the branch-banking microfinance institutions (MFIs) want to clear their loan and villagers want to maximize the amount of “diversified” loans that they can borrow, under some constraints. A stable outcome corresponds to both the MFI-side and the village-side achieving their goals.

Once again, our objective in this networked economy model is to perform causal strategic inference that can assist policy makers in making critical decisions. One example of such inference is: What would happen to the market if some of the government-owned MFIs stop operating? We ascribe the following causal interpretation to this question. We are interested in finding a stable outcome that would result from an *intervention* of removing the government-owned MFIs. A causal strategic inference approach to answering this question would be to first learn the parameters of the model having the government-

owned MFIs in it and then to remove those banks from the model and compute an equilibrium point in the resulting market.

Another example is: What would be the effect of giving additional subsidies to some MFI in the market? A similar causal interpretation can be given to this question as well. We basically perform an intervention in the system by injecting subsidies to an MFI and would like to know what equilibrium point could be obtained as a result of this intervention.

Chapter 3 is organized in the following way. Section 3.1 gives a brief overview of a microfinance system and the central mechanisms in it. Section 3.2 presents our model of a microfinance market and various equilibrium properties of this model. Section 3.3 exploits these properties to give an algorithm for equilibrium computation as well as a method for learning the parameters of the model. The next section presents the results of applying these algorithms to the real-world data obtained from microfinance markets in Bangladesh and Bolivia. Section 3.5 answers various causal questions using our model.



## Chapter 2

# Causal Strategic Inference in Social Networks

The influence of an entity on its peers is a commonly noted phenomenon in both online and real-life social networks. In fact, there is growing scientific evidence that suggests that influence can induce behavioral changes among the agents in a network. For example, recent work in medical social sciences posits the intriguing hypothesis that smoking [25], obesity [24], and even happiness [41] is contagious within a social network. These studies have been based on a real-world social network constructed from the data collected during the Framingham Heart Study, a decades old effort to look into the risk factors of cardiovascular diseases [28]. Regardless of the specific problem addressed, the underlying system under study in that research exhibits several core features. First, it is often very *large and complex*, with many individual entities exhibiting different behaviors and interactions. Second, the *network structure of complex interactions* is central. Third, the *directions and strengths of local influences* are highlighted as very relevant to the global behavior of the system as a whole.

The prevalence of systems and problems like the ones just described in the context of social medical science, combined with the obvious issue of often limited control over individuals, raises immediate, broad, difficult, and longstanding policy questions: e.g., *Can we achieve a desired objective, such as reducing the level of smoking, or controlling obesity via targeted, minimal interventions in a system? How do we optimally allocate our often limited resources to achieve the largest impact in such systems?* Clearly, these issues are not exclusive to obesity, smoking, or happiness; similar issues arise in a large variety of settings: drug use, vaccination, crime networks, security, marketing, markets, the economy, and public policy-making and regulations, to name a few.

The work reported in this dissertation is in large part motivated by such questions and their broader implication. Our approach to modeling influence in networks and formulating the problem of identifying the most influential individuals in a network is causal. This causal connotation will be further explained in the next section once we give a high-level overview of our modeling aspects.

Our major contributions here are: (1) a new approach to influence in networks grounded in non-cooperative game theory; (2) *influence games* as a new class of graphical games to model the behavior of individuals in networks; and (3) a theoretical and empirical study of computational aspects of influence games, including an algorithm for the identification of the most influential individuals.

## 2.1 Influence in Networks

A very important problem in social network analysis is the identification of the most “influential” individuals (see, e.g., [67, 119] and the references therein). We now provide a brief and informal description of our approach to influence in networks with the goal of putting it in the context of the existing literature.

### Overview of Our Model of Influence

Consider a social network where each individual has a binary choice of actions, denoted by  $-1$  and  $1$ . Let us represent this network as a directed graph. Each node of this graph has a *threshold level*, which can be positive, negative, or zero, and the threshold levels of all the nodes are not required to be the same. Each arc of this graph is weighted by an *influence factor*, which signifies the level of influence the tail node of that arc has on the head node. Again, the influence factors can be positive, negative, or zero and are not required to be the same (i.e., symmetric) between two nodes.

Given the above network, our model specifies the *best response* of a node (i.e., what action it should choose) with respect to the actions chosen by the other nodes. The best response of a node is to adopt the action  $1$  if the *total influence* on it exceeds its threshold and  $-1$  if the opposite happens. (In the case of a tie, the node is indifferent between choosing  $1$  and  $-1$ , i.e., either would be its best response.) Here, the total influence on a node is calculated as follows. First, sum up the incoming influence factors on the node from the ones who have adopted the action  $1$ . Second, sum up those influence factors that are coming in from the ones who have adopted  $-1$ . Finally, subtract the second sum from the first to get the total influence on that node.

Clearly, in a network with  $n$  nodes, there are  $2^n$  possible *joint actions*, where each joint action is specified by one action for each node. Among all these joint actions, the ones where every node has chosen its best response to everyone else will be called pure-strategy Nash equilibria (PSNE). We mathematically model the *stable outcomes* that can be generated from such a networked system as PSNE.

## Overview of the Most Influential Nodes Problem

Roughly speaking, in our approach, we consider a set of individuals  $S$  in a network to be *most influential, with respect to some objective of interest*, if  $S$  is the *most preferred* subset among all those that satisfy the following condition: were the individuals in  $S$  to choose the behavior  $\mathbf{x}_S$  prescribed to them by some *stable outcome* of the system  $\mathbf{x} \equiv (\mathbf{x}_S, \mathbf{x}_{-S})$  (which achieves the desirable objective of interest), then the *only* stable outcome of the system that remains consistent with their choices  $\mathbf{x}_S$  is  $\mathbf{x}$  itself. Now, there could be many different sets  $S$  that satisfy this condition. For example,  $S$  could comprise of all the individuals, which would not be a desirable. To account for this, we also specify a preference over all subsets of individuals. A typical example of this preference is selecting a set  $S$  with the minimum cardinality.

Said differently, once the nodes in the most influential set  $S$  follow the behavior  $\mathbf{x}_S$  prescribed to them by a stable outcome  $\mathbf{x}$  achieving the objective of interest, they become collectively so influential that their behavior forces every other individual to a unique choice of behavior! Our proposed concept of the most influential individuals is illustrated in Figure 2.1 with a very simple example.

## Causal Strategic Inference

A causal interpretation is inherent in the way we have formulated the most influential nodes problem. Informally speaking, a set of nodes is most influential with respect to some prescribed behavior if these nodes can *cause* everyone else to follow that prescribed behavior. A central aspect of this causal interpretation is stable outcomes, which we model here as PSNE. In a PSNE, everyone adopts his best response to others' actions. In other words, each node's choice of action depends on what others have chosen. Therefore, a set of nodes can lead the whole system to a particular PSNE by their choice of actions.

We view the question of identifying the most influential set of nodes, which constitutes a major part of this chapter, as a causal strategic inference question. Other related questions, such as identifying a set of senators who can prevent filibusters, as explored in Section 2.5.6, can also be viewed similarly.

Another causal strategic inference question that we will address in Section 2.5.7 is how powerful a “gang” of senators is. Here, the causal interpretation is that a powerful gang of senators would be able to *lead* a majority of the senators into reaching a consensus with them.




Figure 2.1: **Illustration of our approach to influence in networks.** Each node has a binary choice of behavior,  $\{-1, +1\}$ , and wants to behave like the majority of its neighbors (and is indifferent if there is a tie). We adopt pure-strategy Nash equilibrium (to be defined later), abbreviated as PSNE, as the notion of stable outcome. The network is shown in (a) and the enumeration of PSNE (a row for each PSNE, where black denotes node’s behavior 1, gray  $-1$ ) in (b). We want to achieve the objective of every node choosing 1 (desirable outcome). Selecting the set of nodes  $\{1, 2, 3\}$  and assigning these nodes behavior prescribed by the desirable outcome (i.e., 1 for each) lead to two consistent stable outcomes of the system, shown in (c) and (d). Thus,  $\{1, 2, 3\}$  cannot be a most influential set of nodes. On the other hand, selecting  $\{1, 6\}$  and assigning these nodes behavior 1 lead to the desirable outcome as the unique stable outcome remaining. Therefore,  $\{1, 6\}$  is a most influential set, even though these two nodes are at the fringes of the network. Furthermore, note that  $\{1, 6\}$  is not most influential in the diffusion setting, since it does not maximize the spread of behavior 1. (It should be mentioned that we study a much richer class of games here than the one shown in this example.)

### 2.1.1 Connection to Rational Calculus Models of Collective Action

The formal study of individual behavior in a collective setting originally began under the umbrella of “collective behavior” in sociology and social psychology. The classical treatment of collective behavior views individuals in a “crowd” as irrational beings with a lowered intellectual and reasoning ability. The

proposition is that an increased level of suggestibility among the individuals facilitates the rapid spread of the homogeneous “mind of the crowd” [19, 71, 94]. Herbert Blumer’s work, in particular, popularized the classical theory of collective behavior well beyond academia and into such domains as police and the armed forces [82, p. 9]. However, this theory was subjected to much criticism primarily because it did not study empirical accounts systematically. Later on, Blumer himself referred to this as a “miserable job” by sociologists [20].

In response to that, Clark McPhail undertook a massive effort, spanning three decades, to record the behavior of individuals in collective settings that he calls “gatherings” in order to distinguish it from (homogeneous) “crowds” in collective behavior (see McPhail [82, ch. 5, 6] for a summary of his two-decade study). His empirical accounts, stored in a range of media formats as technology improved, reveal one common thing—that a gathering consists of individuals with diverse objectives, who nevertheless behave rationally and purposefully. To distinguish this purposive nature of individuals from irrationality in the classical treatment, he calls his study “collective action” and broadly defines it as “any activity that two or more individuals take with or in relation to one another” [83, p. 881]. In short, collective action can be seen as the modern approach, as opposed to the “old” (but not unimportant) approach of collective behavior [85, p. 14–15]. A brief review of collective behavior and collective action literature is included in the Appendix for interested readers.

Many of the rational calculus or economic choice models that were originally proposed for collective behavior, are now discussed under collective action due to the purposive nature of the individuals. Here, we will conduct a very narrow and focused review of the relevant literature in order to place our model in its proper context. Our review will be concentrated around Mark Granovetter’s threshold models [49], which is one of the most influential models of collective action to date.

## Schelling’s Models of Segregation

A notable precursor to Granovetter’s threshold models is Nobel-laureate economist Thomas Schelling’s models of segregation [108, 109]. Schelling’s models account for segregations that take place as a result of *discriminatory individual behavior* as opposed to organized processes (e.g., separation of on-campus residence between graduate and undergraduate students due to a university’s housing policy) or economic reasons (e.g., segregation between the poor and the rich in many contexts). An example of a segregation due to individual choice, or “individually motivated segregation” as Schelling puts it [109, p. 145], is the residential segregation by color in the U.S. Although Schelling’s

models and their analyses expressly focus on this case, these can be applied to many other scenarios as well.

In Schelling’s *spatial proximity model*, if an individual’s *level of tolerance* for population of the opposing type is exceeded in his neighborhood, he moves to another spatial location where he can be “happy.” The dynamics of segregation is studied in this model using a rule of movement for the “unhappy” individuals. The *bounded-neighborhood model* is concerned with one global neighborhood. An individual enters it if it satisfies its level of tolerance constraint and leaves it otherwise. Schelling studies the stability of equilibria and the *tipping* phenomenon in this model when the distribution of tolerances and the population ratio of the two types are varied. An important finding is that in the cases studied, the modal level of tolerance does not correspond to a tipping point. More on Schelling’s models can be found in the Appendix.

### **Berk’s “Gaming” Approach**

Another notable precursor to Granovetter’s models is Berk’s rational calculus approach [14]. Berk strongly criticizes the assumption of individual irrationality which became prevalent in collective behavior literature. He formulates his approach by first giving a detailed empirical account of an anti-war protest at Northwestern University that originated in a town-hall meeting addressing dormitory rent hike. Berk’s description gives accounts of both mundane and exciting happenings during the course of the protest and is recognized as “among the best in the literature” [82, p. 126]. He explains individual decision making through Raiffa’s decision theory principles. To motivate his approach, he first notes that participating individuals were diverse in their disposition and that they exercised their reasoning power. He then broadly classifies the participants into two types—militants (with the desirable action of trashing properties) and moderates (with the desirable action of an anti-war activity, but not trashing). Each participant, militant or moderate, estimates the support in favor of his disposition, and with enough support, he will “act” (e.g., trash properties if he is militant). Clearly, an individual’s estimate of support directly affects his “payoff.” If an individual estimates that there is not enough support to act in favor of his disposition, he can try to persuade others to support his disposition so that he can receive a higher payoff by being able to act. This can be translated as an attempt to change others’ payoff matrices, which is facilitated by the *milling* phase when they communicate and negotiate with each other. The milling phase ends when a consensus or a compromise is reached and becomes common knowledge. A concerted action takes place at that time.

## Granovetter’s Threshold Models

Granovetter’s threshold models [49] are presented in the setting of a crowd, where each individual is deciding whether to riot or not. In the simplest setting, each individual has a threshold and his decision is influenced by the decisions of others—if the number (or the proportion) of individuals already rioting is below his threshold, then he remains inactive, otherwise he engages in rioting. The emphasis is on investigating equilibrium outcomes due to the process of forward recursion [49, p. 1426], given a distribution of the thresholds of the population. It may be mentioned here that forward recursion starts only if there is an individual with a threshold of 0.

Granovetter’s models are inspired by Schelling’s models of segregation. In fact, one can draw a parallel between Schelling’s level of tolerance and Granovetter’s threshold in the following way. In Schelling’s models, an individual leaves a neighborhood if his level of tolerance is exceeded, whereas in Granovetter’s models, an individual becomes active in rioting if his threshold is exceeded. Furthermore, in both models, dynamics is of utmost importance and serves the purpose of explaining *how* an equilibrium collective outcome emerges from individual behavior. However, apart from these similarities, these two models are semantically different and also focus on completely different outlooks. First, Granovetter ascribes a deeper meaning to the concept of threshold. Threshold of an individual is not just “a number that he carries with him” from one situation to another [49, p. 1436]. It rather depends on the situation in question and can even vary within the same situation due to changes occurring in it. Second, in Granovetter’s models, a very small perturbation in the distribution of population threshold may lead to sharply different equilibrium outcomes. Granovetter highlights this property of his models as an explanation of seemingly paradoxical outcomes that goes against the predispositions of the individuals.

Two features of Granovetter’s models make it stand out among the rational calculus models. First, the models are capable of capturing scenarios beyond the classical realm of collective behavior. Granovetter begins by setting up his model to complement the emergent norm theory (see Appendix) by providing an explicit model of how “individual preferences interact and aggregate” to form a new norm [49, p. 1421]. Not only that such an explicit model eliminates the need for implicit assumptions (such as a new norm emerges when the majority of the population align themselves with that norm), it can also capture paradoxical outcomes alluded above that cannot be captured by the implicit assumption on the majority. Beyond the emergent norm theory, Granovetter’s models can capture a wide range of phenomena that do not fall within the classical realm of collective behavior, such as diffusion of innova-

tion, voting, public opinion, and residential segregation, to name a few. The second prominent feature of Granovetter's models is its ease of adaptation when dealing with a networked population. The same mechanism of forward recursion is applicable when the underlying influence structure is specified by a "sociomatrix," which accounts for how much an individual influences another [49, p.1429]. This is particularly useful for studying collective action in the setting of a social network.

### **Criticism of Rational Calculus Models**

An implicit assumption regarding Granovetter's sociomatrix is that the elements of the matrix are non-negative. Otherwise, the process of forward recursion may never terminate, even on the simplest of examples. However, many real world scenarios do exhibit co-existence of both positive and negative influences. For the most part, democrat senators in the U.S. Congress influence their republican colleagues negatively, while they influence colleagues of their own party positively. In residential segregation involving more than two types of individuals, an individual is negatively influenced, in different magnitudes, by individuals belonging to other types. Clearly, such a situation cannot be modeled using a non-negative sociomatrix. Furthermore, if we take a second look at Berk's account, militant individuals positively reinforce each other in their decision to engage in trashing properties, whereas their decision is negatively affected by the moderates (that is, the presence of too many moderates makes it risky for militants to engage in violent action).

Critiques of rational calculus models point out the lack of behavioral adjustment in a "negative feedback" fashion [83, p. 883]. Here, negative feedback is defined in the context of the perceptual control theory that lays the foundation of McPhail's *sociocybernetics theory* of collective action (see the Appendix). In a negative feedback system, an individual can adjust his behavior depending on the discrepancy between the input signal and the desired signal (the sign of this discrepancy has no correlation to negative feedback). In contrast, in a positive feedback system, such control of behavior is not possible. A typical example of a positive feedback system is a chemical chain reaction. An analogue to this is the "domino effect" cited often in rational calculus models [49, p. 1424]. It is true that rational calculus models neither accounts for "errors" as desired by the proponents of the sociocybernetics theory, nor is it well-defined in the context of rational calculus. But it is not the case that "reversal" of behavior, which can be thought of as a crude form of behavioral adjustment, is precluded in rational calculus models. Such a form of behavioral adjustment can certainly be incorporated by allowing negative elements in the sociomatrix, but the challenge lies in the forward recursion process which may



oscillate indefinitely because of those negative elements.

## Our Approach

Although our approach may seem close to the rational calculus models of collective action, particularly to Granovetter’s threshold models, our objective is very much different from that of collective action theory. The focus of collective action theory in sociology is to explain *how* individual behavior in a group leads to collective outcomes. For example, Schelling’s models explain how different distributions of the level of tolerances of individuals lead to residential segregations of different properties. Berk explains how a compromise (such as placing a barricade) evolves within a mixture of rational individuals of different predispositions (militants vs. moderates). Granovetter shows how a little perturbation in the distribution of thresholds can possibly lead to a completely different collective outcome. In short, explaining collective social phenomena is at the heart of all these studies. While this is a scientific pursuit of utmost importance, our focus is rather on an engineering approach to *predicting* stable behavior in a networked population setting. Our approach is not to go through fine-grained details of a process, such as forward recursion, which is often plagued with problems when the sociomatrix contains negative elements. Instead, we adopt the concept of Nash equilibrium to define stable outcomes in a non-cooperative game setting.<sup>1</sup> Said differently, the *path* to an equilibrium is not what we focus on; rather, it is the prediction of the equilibrium itself that we focus on. Next, we justify this approach.

Sociologists have recorded minute details of various collective action scenarios in order to substantiate their theories with empirical accounts. One example is Clark McPhail’s three-decade-old effort in recording a great many gatherings in various media formats. However, in the application scenarios that we are interested in, such as strategic interaction in the U.S. Congress and the U.S. Supreme Court, very little details can be obtained about *how* a collective outcome emerges. For example, the Budget Control Act of 2011 was passed by 74–26 votes in the Senate on August 2, 2011, ending a much debated debt-ceiling crisis. Despite intense media coverage, it would be difficult, if not impossible, to give an accurate account of how this agreement on debt-ceiling was reached. Even if there were an exact account of every conversation and every negotiation that had taken place, it would be extremely challenging to translate such a subjective account into a mathematically defined process, let alone learning the parameters and computing stable outcomes of such a complex model.

---

<sup>1</sup>Other researchers in the social sciences have also taken this approach of using Nash equilibrium in one-shot non-cooperative games as models of behavior [11, 12, 51–53, 69, 70].

In addition, simplistic models of dynamics used in the literature require some restrictions on the underlying model. For example, as mentioned above, the forward recursion process implicitly assumes that the sociomatrix does not have negative elements. However, if we abstract the process, by the concept of Nash equilibrium in our case, we can deal with rich models without such a restriction and at the same time, capture equilibria beyond the ones captured by a simple model of dynamics. In particular, any equilibrium that the process of forward recursion converges to (with any initial configuration) is also captured by our model; but in addition, our model can capture equilibria that the forward recursion process cannot.

The basic intuition behind our approach is deeply rooted in the philosophy of AI and machine learning. For example, without fully understanding or modeling how human beings perform speech recognition, we have been able to device successful speech recognition systems. Not that the scientific question of how we perform speech recognition is not important, but the focus of AI, in general, is to engineer solutions that serve our purpose, not to explain physical phenomena. Interestingly enough, AI can sometimes help us understand physical phenomena, although not purposefully. Just as building aircrafts helped us understand the aerodynamics of a bird’s flight, recent research by psychologist Alison Gopnik suggests that young children, even 2-year-olds, perform Bayesian inference while learning from the environment [46].

In short, we propose an AI-based approach to predicting the behavior of a large, networked population. Our approach does not model the complex behavioral dynamics that takes place in the network, but abstracts it with the solution concept of Nash equilibrium. By doing this, we are able to deal with a rich set of models and focus more on the prediction of stable outcomes.

### **2.1.2 Connection to Literature on the Most Influential Nodes**

To date, the study of influence in a network, by both economists [26, 90] and computer scientists [34, 67], has been rooted in rational calculus models of behavior. Their approach to connecting individual behavior to collective outcome is mostly by adopting the process of forward recursion [49, p. 1426], which is often employed in studying diffusion of innovations [50, p. 168]. As a result, the term “contagion” in these settings has a rational connotation contrary to the early sociology literature on collective behavior, where “contagion” or “social contagion” alludes to irrational and often hysteric nature of the individuals in a crowd [19, 94]. The computational question of identifying the most influential nodes in a network [67], which was originally posed by

Domingos and Richardson [31], has also been studied using forward recursion within the context of rational calculus models. In the setting of Kleinberg and coauthors’ “cascade” or “diffusion models,” each node behaves in one of these two ways—it either adopts a new behavior or does not, and initially, none of the nodes adopts the new behavior. Given a number  $k$ , their formulation of the most influential nodes problem asks us to select a set of  $k$  nodes such that the *spread* of the new behavior is maximized by the selected nodes being the initial adopters of the new behavior. (Note that in their setting, the set of initial adopters, some of whom may have thresholds greater than 0, are *externally* selected in order to set off the forward recursion process, whereas in Granovetter’s setting, the initial adopters must have a threshold of 0.) It should be noted here that the most influential nodes question in the cascade or diffusion settings sometimes concerns *infinite* graphs [67, p. 615], such as Morris’ local interaction games [90, p. 59], whereas we concern ourselves with large but finite graphs here.

The notion of “most influential nodes” considered in this dissertation is different, and is aimed at *supplementing* the traditional line of work with a new game-theoretic perspective. In addition to the overview of our approach mentioned in the last subsection, let us briefly mention a few contrasting points between Kleinberg and coauthors’ approach to identifying the most influential nodes and that of ours.

A subtle aspect of diffusion models is that each node in the network behaves as an independent agent. Any observed influence that a node’s neighbors impose on the node is the result of the same node’s “rational” or “natural” response to the neighbors’ behavior. Thus, in many cases, it would be desirable that the solution to the most influential nodes problem lead us to a stable outcome of the system, in which each node’s behavior is a best response to the neighbors’ behavior. However, if we select a set of nodes with the goal of maximizing the spread of the new behavior then it might very well happen that some of the selected nodes are “unhappy” being the initial adopters of the new behavior relative to their neighbors’ final behavior at the end of forward recursion. For example, a selected node’s best behavioral response could be *not* adopting the new behavior after all. Thus, is not it more natural to require that the desired final state of the system, such as the maximum spread of the new behavior, be stable, in which everyone is “happy” with their behavioral response?

In order to address the question of finding the most influential nodes, the forward recursion process has been modeled as a “monotonic” process in general. Here, a monotonic process refers to the setting where once an agent adopts the new behavior, it cannot go back. It should be noted here that al-

though two versions of the forward recursion processes—progressive and non-progressive—have been discussed in the literature, previous work on finding the most influential nodes uses the progressive or monotonic version [67]. If we think of an application such as reducing the incidence of smoking or obesity, then a model that allows a “change of mind” based on the response of the immediate neighborhood may make more sense. Thus, a notable contrast between the traditional treatment of the most influential nodes problem and that of ours is that we do not restrict the influence among the nodes of the network to non-negative numbers.

In fact, in many applications, both positive and negative influence factors may exist in the same problem instance. Take the U.S. Congress as an example: senators belonging to the same party may have non-negative influence factors on each other (as usually perceived from voting instances on legislation issues), but one senator may (and often does) have a negative influence on another belonging to a different party. While generalized versions of threshold models that allow “reversals” have been derived in the social science literature, to the best of our knowledge, there is no substantive work on the most influential nodes problem in that context.

Finally, the traditional approach to the most influential nodes problem emphasizes modeling the complex dynamics of interactions among the nodes in order to give the final answer, that is, a set of the most influential nodes. In fact, our model is inspired by the same threshold models that are used by them. However, as we have mentioned earlier, our emphasis is not on the dynamics of interactions, but on the stable outcomes in a game-theoretic setting. By doing this, we are able to capture significant, basic, and core *strategic* aspects of complex interaction in networks that naturally appear in many real-world problems (e.g., identifying the most influential senators in the U.S. Congress). Of course, we recognize the importance of the dynamics of interactions on capturing and studying problems of influence at a finer level of detail. Yet, we believe that our approach can still capture significant aspects of the problem even at the coarser level of “steady-state” or stable outcome.

### 2.1.3 A Brief Note on Mechanism Design

On the surface, our approach may seem related to mechanism design in objective, because mechanism design is often motivated by achieving a desirable objective. For example Balcan et al. study how the Nash equilibrium outcome of a game could be “nudged” toward the one with the best welfare (which gives the *price of stability*) by making some additional information globally available [10]. Although we share the underlying goal of achieving a desirable outcome, our approach is conceptually very different. Here, in contrast to a mechanism-

design approach to achieving desirable stable outcomes, we are not interested in changing, defining, or engineering a new system—*the system is what it is*.

We are rather interested in altering the behavior within the same system so as to lead or “tip” it to a desirable stable outcome [51–53, 69, 70]. To this end, we cannot or need not change the system. Instead, we just need to “convince” the right individuals to adopt the behaviors prescribed by the desirable stable outcome: the others would follow “voluntarily” the behavior prescribed by that outcome because no other stable outcome is possible. Note that making the selected individuals follow the prescribed behavior is facilitated by the fact that they will end up “happy:” none will have an incentive to behave any other way after all!

## 2.2 Influence Games

Inspired by the well-studied threshold models [49], we introduce *influence games* as a model of influence in large networked populations. The payoffs of the players in an influence games are defined in a parametric fashion, leading to a compact representation of the game. Various commonly known games on networks, such as the *threshold games of complement* [59, Ch. 9] as well as the game illustrated in Figure 2.1, can be shown to be special instances of influence games. On the other hand, one broad subclass of influence games which we will call *linear influence games* falls within the general class of graphical games [66]. Moreover, we will show in Section 2.2.4 that linear influence games are nothing but polymatrix games [60].

The motivation for defining influence games as a model of behavior is to be able to compute the stable outcomes of a system as well as to identify the most influential entities in a network *relative to a particular objective*. We will first define a general game-theoretic framework to accomplish these tasks. We will then define and use linear influence games as a special type of influence games.

### 2.2.1 General Game-Theoretic Model

Let us first formalize *influence games* as a general model of behavior. Let  $n$  be the number of individuals in the population. For simplicity, we restrict our attention to binary behavior, a common assumption in most of the work in this area. Thus,  $x_i \in \{-1, 1\}$  denotes the *behavior* of individual  $i$ , where  $x_i = 1$  indicates that  $i$  “adopts” a particular behavior and  $x_i = -1$  indicates  $i$  “does not adopt” the behavior. Some examples of behavior of this kind are supporting a particular political measure, candidate or party; holding a par-

ticular view or belief; vaccinating against a particular disease; installing virus protection software (and keeping it up-to-date); acquiring fire/home insurance; becoming overweight; taking up smoking; becoming a criminal or participating in criminal activity; among many others.

**Definition 2.2.1.** Denote by  $f_i : \{-1, 1\}^{n-1} \rightarrow \mathbb{R}$  the function that quantifies the “influence” of other individuals on  $i$ . In influence games, we define the payoff function  $u_i : \{-1, 1\}^n \rightarrow \mathbb{R}$  quantifying the preferences of each player  $i$  as  $u_i(x_i, \mathbf{x}_{-i}) \equiv x_i f_i(\mathbf{x}_{-i})$ , where  $\mathbf{x}_{-i}$  denotes the vector of all joint-actions excluding that of  $i$ .

Given  $\mathbf{x}_{-i} \in \{-1, 1\}^{n-1}$ , the best-response correspondence  $\mathcal{BR}_i^{\mathcal{G}} : \{-1, 1\}^{n-1} \rightarrow 2^{\{-1, 1\}}$  of a player  $i$  of an influence game  $\mathcal{G}$  is defined as follows.

$$\mathcal{BR}_i^{\mathcal{G}}(\mathbf{x}_{-i}) \equiv \arg \max_{x_i \in \{-1, 1\}} u_i(x_i, \mathbf{x}_{-i}).$$

Therefore, for all individuals  $i$  and any possible behavior  $\mathbf{x}_{-i} \in \{-1, 1\}^{n-1}$  of the other individuals in the population, the best-response behavior  $x_i^*$  of individual  $i$  to the behavior  $\mathbf{x}_{-i}^*$  of others satisfies

$$\begin{aligned} f_i(x_{-i}^*) > 0 &\implies x_i^* = 1, \\ f_i(x_{-i}^*) < 0 &\implies x_i^* = -1, \text{ and} \\ f_i(x_{-i}^*) = 0 &\implies x_i^* \in \{-1, 1\}. \end{aligned}$$

Informally, “positive influences” lead an individual to adopt the behavior, while “negative influences” lead the individual to “reject” the behavior; the individual is indifferent if there is “no influence.” A *stable outcome* of the system, by which we formally mean a *pure-strategy Nash equilibrium (PSNE)* of the corresponding influence game  $\mathcal{G}$ , is a behavior assignment  $\mathbf{x}^* \in \{-1, 1\}^n$  that satisfies all those conditions: Each player  $i$ ’s behavior  $x_i^*$  is a (simultaneous) best-response to the behavior  $\mathbf{x}_{-i}^*$  of the rest. Denote the set of PSNE of game  $\mathcal{G}$  by

$$\mathcal{NE}(\mathcal{G}) \equiv \{\mathbf{x}^* \in \{-1, 1\}^n \mid x_i^* \in \mathcal{BR}_i^{\mathcal{G}}(\mathbf{x}_{-i}^*) \text{ for all } i\}.$$

## 2.2.2 Most Influential Nodes: Problem Formulation

In formulating the most influential nodes problem in a network, we depart from the traditional model of diffusion and adopt influence games as the model of strategic behavior among the nodes in the network.

**Definition 2.2.2.** Let  $\mathcal{G}$  be an influence game,  $g : \{-1, 1\}^n \times 2^{[n]} \rightarrow \mathbb{R}$  be the goal or objective function mapping a joint-action and a subset of the players

in  $\mathcal{G}$  to a real number quantifying the general preferences over the space of joint-actions and players' subsets, and  $h : 2^{[n]} \rightarrow \mathbb{R}$  be the set-preference function mapping a subset of the players to a real number quantifying the a priori preference over the space of players' subsets. Denote by  $\mathcal{X}_g^*(S) \equiv \arg \max_{\mathbf{x} \in \mathcal{NE}(\mathcal{G})} g(\mathbf{x}, S)$  the optimal set of PSNE of  $\mathcal{G}$ , with respect to  $g$  and a fixed subset of players  $S \subset [n]$ . We say that a set of nodes/players  $S^* \subset [n]$  in  $\mathcal{G}$  is most influential with respect to  $g$  and  $h$ , if

$$S^* \in \arg \max_{S \subset [n]} h(S), \text{ s.t., } |\{\mathbf{x} \in \mathcal{NE}(\mathcal{G}) \mid \mathbf{x}_S = \mathbf{x}_S^*, \mathbf{x}^* \in \mathcal{X}_g^*(S)\}| = 1.$$

As mentioned earlier, we can interpret the players in  $S^*$  to be collectively so influential that they are able to restrict every other player's choice of action to a unique one: the action prescribed by some desirable stable outcome  $\mathbf{x}^*$ .

An example of a goal function  $g$  that captures the objective of achieving a specific stable outcome  $\mathbf{x}^* \in \mathcal{NE}(\mathcal{G})$  is  $g(\mathbf{x}, S) \equiv \mathbb{1}[\mathbf{x} = \mathbf{x}^*]$ . Another example that captures the objective of achieving a stable outcome with the largest number of individuals adopting the behavior is  $g(\mathbf{x}, S) \equiv \sum_{i=1}^n \frac{x_i + 1}{2}$ .

A common example of the set-preference function  $h$  that captures the preference for sets of small cardinality is to simply define  $h$  such that  $h(S) > h(S')$  iff  $|S| < |S'|$ .

### 2.2.3 Linear Influence Games

A simple instantiation of the general influence game model just described is the case of linear influences.

**Definition 2.2.3.** *In a linear influence game (LIG), the influence function of each individual  $i$  is defined as  $f_i(\mathbf{x}_{-i}) \equiv \sum_{j \neq i} w_{ji} x_j - b_i$  where for any other individual  $j$ ,  $w_{ji} \in \mathbb{R}$  is a weight parameter quantifying the “influence factor” that  $j$  has on  $i$ , and  $b_i \in \mathbb{R}$  is a threshold parameter for  $i$ 's level of “tolerance” for negative effects.*

It follows from Definition 2.2.1 that although the influence function of an LIG is linear, its payoff function is quadratic. Furthermore, the following argument shows that an LIG is a special type of graphical game in *parametric* form. In general, the influence factors  $w_{ji}$  induce a directed graph, where nodes represent individuals, and therefore, we obtain a graphical game having a *linear* (in the number of edges) representation size, as opposed to the *exponential* (in the maximum degree of a node) representation size of general graphical games in normal form [66]. In particular, there is a directed edge (or arc) from individual  $j$  to  $i$  iff  $w_{ji} \neq 0$ .

## 2.2.4 Connection to Polymatrix Games

Polymatrix games [60] are defined to be  $n$ -player noncooperative games where a player's total payoff is the sum of the *partial payoffs* received from the other players. For any joint action  $\mathbf{x}$ , player  $i$ 's payoff is given by  $M_i(x_i, \mathbf{x}_{-i}) \equiv \sum_{j \neq i} \alpha_{ji}(x_j, x_i)$ , where  $\alpha_{ji}(x_j, x_i)$  is the partial payoff that  $i$  receives from  $j$  when  $i$  plays  $x_i$  and  $j$  plays  $x_j$ . Note that this partial payoff is local in nature and is not affected by the choice of actions of the other nodes. We will consider polymatrix games with only binary actions  $\{1, -1\}$  here.

The following property shows an equivalence between LIGs and 2-action polymatrix games. Thus, our computational study of LIGs directly carries over to 2-action polymatrix games.

**Proposition 2.2.4.** *LIGs are equivalent to 2-action polymatrix games with respect to the set of PSNE.*

*Proof.* Assume that the number of players  $n > 1$ ; otherwise, the statement holds trivially. We first show that given any instance of an LIG, we can design a polymatrix game that has the same set of PSNE. In an LIG instance, player  $i$ 's payoff is given by

$$\begin{aligned} u_i(x_i, \mathbf{x}_{-i}) &= x_i \left( \sum_{j \neq i} w_{ji} x_j - b_i \right) \\ &= x_i \sum_{j \neq i} \left( w_{ji} x_j - \frac{b_i}{n-1} \right) \\ &= \sum_{j \neq i} \left( x_i w_{ji} x_j - \frac{x_i b_i}{n-1} \right). \end{aligned}$$

Thus, constructing a polymatrix game instance by defining  $\alpha_{ji}(x_j, x_i) \equiv x_i w_{ji} x_j - \frac{x_i b_i}{n-1}$ , we have the same set of PSNE in both instances.

Next, we show the reverse direction. Player  $i$ 's payoff in a 2-action polymatrix game is given by



$$\begin{aligned}
M_i(x_i, \mathbf{x}_{-i}) &= \sum_{j \neq i} \alpha_{ji}(x_j, x_i) \\
&= \sum_{j \neq i} (\mathbb{1}[x_i = 1] \alpha_{ji}(x_j, 1) + \mathbb{1}[x_i = -1] \alpha_{ji}(x_j, -1)) \\
&= \sum_{j \neq i} \left( \frac{1+x_i}{2} \alpha_{ji}(x_j, 1) + \frac{1-x_i}{2} \alpha_{ji}(x_j, -1) \right) \\
&= \frac{x_i}{2} \sum_{j \neq i} (\alpha_{ji}(x_j, 1) - \alpha_{ji}(x_j, -1)) + \frac{1}{2} \sum_{j \neq i} (\alpha_{ji}(x_j, 1) + \alpha_{ji}(x_j, -1)).
\end{aligned}$$

Note that the second term above does not have any effect on  $i$ 's choice of action. Thus, we can re-define the payoff of player  $i$ , without making any change to the set of PSNE of the original polymatrix game, as follows.

$$\begin{aligned}
M'_i(x_i, \mathbf{x}_{-i}) &= \frac{x_i}{2} \sum_{j \neq i} (\alpha_{ji}(x_j, 1) - \alpha_{ji}(x_j, -1)) \\
&= \frac{x_i}{2} \left( \sum_{j \neq i} (\mathbb{1}[x_j = 1] \alpha_{ji}(1, 1) + \mathbb{1}[x_j = -1] \alpha_{ji}(-1, 1)) - \right. \\
&\quad \left. \sum_{j \neq i} (\mathbb{1}[x_j = 1] \alpha_{ji}(1, -1) + \mathbb{1}[x_j = -1] \alpha_{ji}(-1, -1)) \right) \\
&= \frac{x_i}{2} \left( \sum_{j \neq i} \left( \frac{1+x_j}{2} \alpha_{ji}(1, 1) + \frac{1-x_j}{2} \alpha_{ji}(-1, 1) \right) - \right. \\
&\quad \left. \sum_{j \neq i} \left( \frac{1+x_j}{2} \alpha_{ji}(1, -1) + \frac{1-x_j}{2} \alpha_{ji}(-1, -1) \right) \right) \\
&= \frac{x_i}{4} \left( \sum_{j \neq i} x_j (\alpha_{ji}(1, 1) - \alpha_{ji}(-1, 1) - \alpha_{ji}(1, -1) + \alpha_{ji}(-1, -1)) \right. \\
&\quad \left. + \sum_{j \neq i} (\alpha_{ji}(1, 1) + \alpha_{ji}(-1, 1) - \alpha_{ji}(1, -1) - \alpha_{ji}(-1, -1)) \right).
\end{aligned}$$

Therefore, we can construct an LIG that has exactly the same set of PSNE as the polymatrix game, in the following way. For any player  $i$ , define  $b_i \equiv$

$-\sum_{j \neq i} \frac{1}{4}(\alpha_{ji}(1, 1) + \alpha_{ji}(-1, 1) - \alpha_{ji}(1, -1) - \alpha_{ji}(-1, -1))$ , and for any player  $i$  and any other player  $j$ , define  $w_{ji} \equiv \sum_{j \neq i} \frac{1}{4}(\alpha_{ji}(1, 1) - \alpha_{ji}(-1, 1) - \alpha_{ji}(1, -1) + \alpha_{ji}(-1, -1))$ .  $\square$

### 2.2.5 Learning Influence Games

One of the core components of causal strategic inferences is learning the parameters of the model from real-world data. As we showed earlier in Section 1.4, the econometrics literature is rich with techniques for estimating the parameters of game-theoretic models, but it mostly relates to the games based on random utility models [77, 78]. Our model, in contrast, is rooted in sociology literature on threshold models. Also, considering the scale and the complexity of our model, an analytic approach typically taken in econometrics would be unmanageable in our case (without imposing additional restrictions). For these practical reasons, we have adopted a computational approach to learning the parameters of our model. In particular, we have adopted the machine learning scheme given by Honorio and Ortiz [55].

In contrast to many other game-theoretic model, LIG has the special property that it focuses on a *set* of PSNE rather than a single PSNE. For example, when we consider congressional voting and Supreme Court rulings in Sections 2.5.4 and 2.5.3, we will see that we want our model to capture each of a set of voting outcomes as a PSNE of the game. The machine learning scheme of Honorio and Ortiz addresses this issue by estimating the parameters of the model in a way that maximizes the number of observed outcomes captured as PSNE of the model and minimizes the number of unobserved outcomes also captured as PSNE [55]. They formulate a simultaneous logistic regression technique to achieve this. Among multiple models that fulfill this objective reasonably well, they select the one that is the simplest (that is, in graph-theoretic terms, they prefer a sparse graph to a dense graph). For more details, the reader is referred to [55].

## 2.3 Equilibria Computation in Influence Games

We first study the problem of computing and counting PSNE in LIGs. We show that several special cases of LIGs present us with attractive computational advantages, while the general problem is intractable unless  $P = NP$ . We present heuristics to compute PSNE in general LIGs.

### 2.3.1 Nonnegative Influence Factors

When all the influence factors are non-negative, an LIG is supermodular [84, 113]. In particular, the game exhibits what is called *strategic complementarity* [23]. Hence, the best-response dynamics converges in at most  $n$  rounds. From this, we obtain the following result.

**Proposition 2.3.1.** *The problem of computing a PSNE is in  $P$  for LIGs on general graphs with only non-negative influence factors.*

This property implies certain monotonicity of the best-response correspondences. More specifically, for each player  $i$ , if any subset of the other players “increases his/her strategy” by adopting the new behavior, then player  $i$ ’s best-response cannot be to abandon adoption (i.e., move from 1 to  $-1$ ). In other words, once a player adopts the new behavior, it has no incentive to go back. This monotonicity property also follows directly from the linear threshold model. Strategic complementarity implies other interesting characterizations of the structure of PSNE in LIGs and the behavior of best-response dynamics. For example, it is not hard to see that such games always have a PSNE: If we start with the complete assignment in which either everyone is playing 1, or everyone is playing  $-1$ , parallel/synchronous best-response dynamics converges after at most  $n$  rounds [84]. If both best-response processes starting with all  $-1$ ’s and all 1’s converge to the same PSNE, then the PSNE is unique. Otherwise, any other PSNE of the game must be “contained” between the two different PSNE. We can also view this from the perspective of constraint propagation with monotonic constraints [105].

### 2.3.2 Special Influence Structures and Potential Games

Several special subclasses of LIGs are potential games [87]. This connection guarantees the existence of PSNE in such games.

**Proposition 2.3.2.** *If the influence factors of an LIG  $\mathcal{G}$  are symmetric (i.e.,  $w_{ji} = w_{ij}$ , for all  $i, j$ ), then  $\mathcal{G}$  is a potential game.*

*Proof.* We show that the game has a cardinal potential function,

$$\Phi(\mathbf{x}) = \sum_{t=1}^n x_t \left( \sum_{i \neq t} \frac{x_i w_{it}}{2} - b_t \right). \quad (2.1)$$

Consider any player  $j$ . The difference in  $j$ ’s payoff for  $x_j = 1$  and  $x_j = -1$

(assuming all other players play  $\mathbf{x}_{-j}$  in both cases) is

$$\begin{aligned}
& u_j(1, \mathbf{x}_{-j}) - u_j(-1, \mathbf{x}_{-j}) \\
&= 1 \times \left( \sum_{i \neq j} x_i w_{ij} - b_j \right) - (-1) \times \left( \sum_{i \neq j} x_i w_{ij} - b_j \right) \\
&= 2 \times \left( \sum_{i \neq j} x_i w_{ij} - b_j \right).
\end{aligned} \tag{2.2}$$

Next, the difference in the potential function when  $j$  plays 1 and  $-1$  is

$$\begin{aligned}
& \Phi(1, \mathbf{x}_{-j}) - \Phi(-1, \mathbf{x}_{-j}) \\
&= 1 \times \left( \sum_{i \neq j} \frac{x_i w_{ij}}{2} - b_j \right) + \sum_{t \neq j} x_t \left( \sum_{i \neq t} \mathbb{1}[i \neq j] \frac{x_i w_{it}}{2} - b_t \right) \\
&+ \sum_{t \neq j} x_t \left( \sum_{i \neq t} \mathbb{1}[i = j] \frac{1 \times w_{it}}{2} - b_t \right) - \\
&(-1) \times \left( \sum_{i \neq j} \frac{x_i w_{ij}}{2} - b_j \right) - \sum_{t \neq j} x_t \left( \sum_{i \neq t} \mathbb{1}[i \neq j] \frac{x_i w_{it}}{2} - b_t \right) \\
&- \sum_{t \neq j} x_t \left( \sum_{i \neq t} \mathbb{1}[i = j] \frac{(-1) \times w_{it}}{2} - b_t \right) \\
&= 2 \times \left( \sum_{i \neq j} \frac{x_i w_{ij}}{2} - b_j \right) + 2 \times \left( \sum_{t \neq j} \frac{x_t w_{jt}}{2} \right) \\
&= 2 \times \left( \sum_{i \neq j} x_i w_{ij} - b_j \right).
\end{aligned} \tag{2.3}$$

The last line follows due to the symmetric weights (i.e.,  $w_{ij} = w_{ji}$ ).  $\square$

If, in addition, the threshold  $b_i = 0$  for all  $i$ , the game is a *party-affiliation game*, and computing a PSNE in such games is PLS-complete [36].

The following result is on a large class of games that we call *indiscriminate LIGs*, where for every player  $i$ , the influence weight,  $w_{ij} \equiv \delta_i \neq 0$ , that  $i$  imposes on every other player  $j$  is the same. The interesting aspect of this result is that these LIGs are potential games despite being possibly *asymmetric* and exhibiting strategic substitutability (due to negative influence factors).

**Proposition 2.3.3.** *Let  $\mathcal{G}$  be an indiscriminate LIG in which all  $\delta_i$  for all  $i$ ,*

have the same sign, denoted by  $\rho \in \{-1, +1\}$ . Then  $\mathcal{G}$  is a potential game with the following potential function  $\Phi(\mathbf{x}) = \rho \left[ (\sum_{i=1}^n \delta_i x_i)^2 - 2 \sum_{i=1}^n b_i \delta_i x_i \right]$ .

*Proof.* It is sufficient to show that the sign of the difference in the individual utilities of any player due to changing her action unilaterally, is the same as the sign of the difference in the corresponding potential functions. For any player  $j$ , the first difference is

$$\begin{aligned} & 1 \times \left( \sum_{i \neq j} \delta_i x_i - b_j \right) - (-1) \times \left( \sum_{i \neq j} \delta_i x_i - b_j \right) \\ &= 2 \left( \sum_{i \neq j} \delta_i x_i - b_j \right). \end{aligned} \tag{2.4}$$

The potential function when  $j$  plays 1,

$$\begin{aligned} & \Phi(x_j = 1, \mathbf{x}_{-j}) \\ &= \rho \left[ \left( \sum_{i \neq j} \delta_i x_i + \delta_j \times 1 \right)^2 - 2 \sum_{i \neq j} b_i \delta_i x_i - 2b_j \delta_j \times 1 \right] \\ &= \rho \left[ \left( \sum_{i \neq j} \delta_i x_i \right)^2 + \delta_j^2 + 2 \left( \sum_{i \neq j} \delta_i x_i \right) \delta_j - 2 \sum_{i \neq j} b_i \delta_i x_i - 2b_j \delta_j \right]. \end{aligned}$$

The potential function when  $j$  plays  $-1$ ,

$$\begin{aligned} & \Phi(x_j = -1, \mathbf{x}_{-j}) \\ &= \rho \left[ \left( \sum_{i \neq j} \delta_i x_i + \delta_j \times (-1) \right)^2 - 2 \sum_{i \neq j} b_i \delta_i x_i - 2b_j \delta_j \times (-1) \right] \\ &= \rho \left[ \left( \sum_{i \neq j} \delta_i x_i \right)^2 + \delta_j^2 - 2 \left( \sum_{i \neq j} \delta_i x_i \right) \delta_j - 2 \sum_{i \neq j} b_i \delta_i x_i + 2b_j \delta_j \right]. \end{aligned}$$

Thus, the difference in the potential functions,

$$\Phi(x_j = 1, \mathbf{x}_{-j}) - \Phi(x_j = -1, \mathbf{x}_{-j}) = 4\rho\delta_j \left( \sum_{i \neq j} \delta_i x_i - b_j \right). \tag{2.5}$$

Since  $\rho\delta_j > 0$ , the quantities given in (2.4) and (2.5) have the same sign.  $\square$

### 2.3.3 Tree-Structured Influence Games

The following result follows from a careful, non-trivial modification of the **TreeNash** algorithm [66]. Note that the running time of the **TreeNash** algorithm is exponential in the degree of a node and thus also exponential in the representation size of an LIG! In contrast, our algorithm is linear in the maximum degree and thereby linear in the representation size of an LIG. The complete proof follows a proof sketch.

**Theorem 2.3.4.** *There exists an  $O(nd)$  time algorithm to find a PSNE, or to decide that there exists none, in LIGs with tree structures, where  $d$  is the maximum degree of a node.*

*Proof Sketch.* We use similar notations as in [66]. The modification of the **TreeNash** involves efficiently (in  $O(d)$  time, not  $O(2^d)$ ) determining the existence of a witness vector and constructing one, if it exists, at each node during the downstream pass, in the following way.

Suppose that an internal node  $i$  receives tables  $T_{ki}(x_k, x_i)$  from its parents  $k$ , and that  $i$  wants to send a table  $T_{ij}(x_i, x_j)$  to its unique child  $j$ . If for some parent  $k$  of  $i$ ,  $T_{ki}(-1, x_i) = 0$  and  $T_{ki}(1, x_i) = 0$ , then  $i$  sends the following table entries to  $j$ :  $T_{ij}(x_i, -1) = 0$  and  $T_{ij}(x_i, 1) = 0$ . Otherwise, we first partition  $i$ 's set of parents into two sets in  $O(d)$  time:  $Pa_1(i, x_i)$  consisting of the parents  $k$  of  $i$  that have a unique best response  $\hat{x}_k$  to  $i$ 's playing  $x_i$  and  $Pa_2(i, x_i)$  consisting of the remaining parents of  $i$ . We show that  $T_{ij}(x_i, x_j) = 1$  iff

$$x_i(x_j w_{ji} + \sum_{k \in Pa_1(i, x_i)} \hat{x}_k w_{ki} + \sum_{t \in Pa_2(i, x_i)} \underbrace{(2 \times \mathbb{1}[x_i w_{ti} > 0] - 1)}_{t\text{'s action in witness vector}} w_{ti}) \geq 0,$$

from which we get a witness vector, if it exists. □

Following is the complete proof of Theorem 2.3.4.

*Proof.* We denote any node  $i$ 's action by  $x_i \in \{-1, 1\}$ , its threshold by  $b_i$ , and the influence of any node  $i$  on another node  $j$  by  $w_{ij}$ . Furthermore, the set of parents of a node  $i$  is denoted by  $Pa(i)$ . The two phases of the modified TreeNash algorithm are described below.

1. **Downstream phase.** In this phase each node sends a table to its unique child. We denote the table that node  $i$  sends to its child  $j$  as  $T_{ij}(x_i, x_j)$ , which is indexed by the actions of  $i$  and  $j$ , and define the set

of conditional best-responses of a node  $i$  to a neighboring node  $j$ 's action  $x_j$  as  $BR_i(j, x_j) \equiv \{x_i \mid T_{ij}(x_i, x_j) = 1\}$ . If  $|BR_i(j, x_j)| = 1$  then we will abuse this notation by letting  $BR_i(j, x_j)$  be the unique best-response of  $i$  to  $j$ 's action  $x_j$ .

The downstream phase starts at the leaf nodes. Each leaf node  $l$  sends a table  $T_{lk}(x_l, x_k)$  to its child  $k$ , where  $T_{lk}(x_l, x_k) = 1$  if and only if  $x_l$  is a conditional best-response of  $l$  to  $k$ 's choice of action  $x_k$ . Suppose that an internal node  $i$  obtains tables  $T_{ki}(x_k, x_i)$  from its parents  $k \in Pa(i)$ , and that  $i$  needs to send a table to its child  $j$ . Once  $i$  receives the tables from its parents, it first computes (in  $O(d)$  time) the following three sets that partition the parents of  $i$  based on the size of their conditional best-response sets when  $i$  plays  $x_i$ .

$$Pa_r(i, x_i) \equiv \{k \text{ s.t. } k \in Pa(i) \text{ and } |BR_k(i, x_i)| = r\}, \text{ for } r = 0, 1, 2.$$

This is how  $i$  computes the table  $T_{ij}(x_i, x_j)$  to be sent to  $j$ :  $T_{ij}(x_i, x_j) = 1$  if and only if there exists a *witness vector*  $(x_k)_{k \in Pa(i)}$  that satisfies the following two conditions:

*Condition 1.*  $T_{ki}(x_k, x_i) = 1$  for all  $k \in Pa(i)$ .

*Condition 2.* The action  $x_i$  is a best-response of node  $i$  when every node  $k \in Pa(i)$  plays  $x_k$  and  $j$  plays  $x_j$ .

There are two cases.

**Case I:**  $Pa_0(i, x_i) \neq \emptyset$ . In this case, there exists some parent  $k$  of  $i$  for which both  $T_{ki}(-1, x_i) = 0$  and  $T_{ki}(1, x_i) = 0$ . Therefore, there exists no witness vector that satisfies Condition 1, and  $i$  sends the following table entries to  $j$ :  $T_{ij}(x_i, x_j) = 0$ , for  $x_j = -1, 1$ .

**Case II:**  $Pa_0(i, x_i) = \emptyset$ . In this case, we will show that there exists a witness vector for  $T_{ij}(x_i, x_j) = 1$  satisfying Conditions 1 and 2 if and only if the following inequality holds (which can be verified in  $O(d)$  time). Below, we will use the *sign* function  $\sigma$ :  $\sigma(x) = 1$  if  $x > 0$ , and  $\sigma(x) = -1$  otherwise.

$$x_i \left( w_{ji}x_j + \sum_{k \in Pa_1(i, x_i)} w_{ki}BR_k(i, x_i) + \sum_{k \in Pa_2(i, x_i)} w_{ki}\sigma(x_i w_{ki}) - b_i \right) \geq 0. \quad (2.6)$$

In fact, if Inequality (2.6) holds then we can construct a witness vector in the following way: If  $k \in Pa_1(i, x_i)$ , then let  $x_k = BR_k(i, x_i)$ , otherwise, let  $x_k = \sigma(x_i w_{ki})$ . Since each parent  $k$  of  $i$  is playing its conditional best-response  $x_k$  to  $i$ 's choice of action  $x_i$ , we obtain,  $T_{ki}(x_k, x_i) = 1$  for all  $k \in Pa(i)$ . Furthermore, Inequality (2.6) says that  $i$  is playing its best-response  $x_i$  to each of its parent  $k$  playing  $x_k$  and its child  $j$  playing  $x_j$ .

To prove the reverse direction, we start with a witness vector  $(x_k)_{k \in Pa(i)}$  such that Conditions 1 and 2 specified above hold. In particular, Condition 2 can be written as:

$$x_i \left( w_{ji} x_j + \sum_{k \in Pa(i)} w_{ki} x_k - b_i \right) \geq 0. \quad (2.7)$$

The following line of arguments shows that Inequality (2.6) holds.

$$\begin{aligned} & x_i w_{ki} \sigma(x_i w_{ki}) \geq x_i w_{ki} x_k, \text{ for any } k \in Pa_2(i, x_i) \\ \Rightarrow & x_i \sum_{k \in Pa_2(i, x_i)} w_{ki} \sigma(x_i w_{ki}) \geq x_i \sum_{k \in Pa_2(i, x_i)} w_{ki} x_k \\ \Rightarrow & x_i \left( w_{ji} x_j + \sum_{k \in Pa_1(i, x_i)} w_{ki} BR_k(i, x_i) + \sum_{k \in Pa_2(i, x_i)} w_{ki} \sigma(x_i w_{ki}) - b_i \right) \\ & \geq x_i \left( w_{ji} x_j + \sum_{k \in Pa(i)} w_{ki} x_k - b_i \right) \\ \Rightarrow & x_i \left( w_{ji} x_j + \sum_{k \in Pa_1(i, x_i)} w_{ki} BR_k(i, x_i) + \sum_{k \in Pa_2(i, x_i)} w_{ki} \sigma(x_i w_{ki}) - b_i \right) \\ & \geq 0, \text{ using Inequality (2.7)}. \end{aligned}$$

In addition to computing the table  $T_{ij}$ , node  $i$  stores the following witness vector  $(x_k)_{k \in Pa(i)}$  for each table entry  $T_{ij}(x_i, x_j)$  that is 1: if  $k \in Pa_1(i, x_i)$ , then  $x_k = BR_k(i, x_i)$ , otherwise,  $x_k = \sigma(x_i w_{ki})$ . The downstream phase ends at the root node  $z$ , and  $z$  computes a unary table  $T_z(x_z)$  such that  $T_z(x_z) = 1$  if and only if there exists a witness vector  $(x_k)_{k \in Pa(z)}$  such that  $T_{kz}(x_k, x_z) = 1$  for all  $k \in Pa(z)$  and  $x_z$  is a best-response of  $z$  to  $(x_k)_{k \in Pa(z)}$ .

The time complexity of the downstream phase is dominated by the com-



putation of the table at each node, which is  $O(d)$ . We visit every node exactly once. So, the downstream phase is completed in  $O(nd)$ . Note that if there does not exist any PSNE in the game then all the table entries computed by some node will be 0.

2. **Upstream phase.** In the upstream phase, each node sends instructions to its parents about which actions to play, along with the action that the node itself is playing. The upstream phase begins at the root node  $z$ . For any table entry  $T_z(x_z) = 1$ ,  $z$  decides to play  $x_z$  itself and instructs each of its parents to play the action in the witness vector associated with  $T_z(x_z) = 1$ . At an intermediate node  $i$ , suppose that it has been instructed to play  $x_i$  by its child  $j$  which itself is playing  $x_j$ . The node  $i$  looks up the witness vector  $(x_k)_{k \in Pa(i)}$  associated with  $T_{ij}(x_i, x_j) = 1$  and instructs its parents to play according to that witness vector. This process propagates upward, and when we reach all the leaf nodes, we obtain a PSNE for the game. Note that we can find a PSNE in this phase if and only if there exists one.

In the upstream phase, each node sends  $O(d)$  instructions to its parents. Thus, the upstream phase takes  $O(nd)$  time, and the whole algorithm takes  $O(nd)$  time.  $\square$

### 2.3.4 Hardness Results

Computing PSNE in a general graphical game is known to be computationally hard [47]. However, that result does not imply intractability in our problem, nor do the proofs seem easily adaptable to our case. LIGs are a special type of graphical game with quadratic payoffs, or in other words a graphical, *parametric* poly-matrix game [60], and thus have a more succinct representation than general graphical games ( $O(nd)$  in contrast to  $O(n2^d)$ , where  $d$  is the maximum degree of a node). Next, we show that various interesting computational questions on LIGs are intractable, unless  $P = NP$ .

The central hardness question on LIGs (and also on 2-action polymatrix games) is settled by 1(a) below. Related to the most influential nodes problem formulation, 1(b) states that given a subset of players, it is NP-complete to decide whether there exists a PSNE in which this subset of players adopts the new behavior. A similar statement is given in 1(c).

A prime feature of our formulation of the most influential nodes is the uniqueness of the desirable stable outcome when the set of the most influential nodes adopt their behavior according to the desirable stable outcome. Deciding whether a given set of players fulfills this criterion (in the special case of a

desirable outcome where this set of players adopts the new behavior) is shown to be co-NP-complete in 2.

As we will see later, in order to compute a set of the most influential nodes, it suffices to be able to count the number of PSNE of an LIG (to be more specific, it suffices to count the number of PSNE extensions for a given partial assignment to the players' actions). This problem is shown to be #P-complete in 3. Note that the #P-completeness result for LIGs even with star structure is in contrast to the polynomial-time counterpart for general graphical games with tree graphs, for which not only deciding the existence of a PSNE is in P, but also counting PSNE on general graphical games with tree graphs is in P. This result can be better appreciated by considering the representation sizes of LIGs and tree-structured graphical games, which are linear and exponential in the maximum degree, respectively.

Below, we first summarize the hardness results with an outline of proof, followed by complete proofs of individual statements.

**Theorem 2.3.5.** 1. *It is NP-complete to decide the following questions in LIGs.*

- (a) *Does there exist a PSNE?*
- (b) *Given a designated non-empty set of players, does there exist a PSNE consistent with those players playing 1?*
- (c) *Given a number  $k \geq 1$ , does there exist a PSNE with at least  $k$  players playing 1?*

2. *Given an LIG and a designated non-empty set of players, it is co-NP-complete to decide if there exists a unique PSNE with those players playing 1.*

3. *It is #P-complete to count the number of PSNE, even for special classes of the underlying graph, such as a bipartite or a star graph.*

*Proof Sketch.* The complete proofs can be found in the Appendix. The proof of 1(a) reduces the 3-SAT problem to an LIG that consists of a player for each clause and each variable of the 3-SAT instance. The influence factors among these players are designed such that the LIG instance possesses a PSNE if and only if the 3-SAT instance has a satisfying assignment. Since the underlying graph of the LIG instance is always bipartite, we obtain as a corollary that the NP-completeness of that existence problem holds even for LIGs on bipartite graphs.

The proofs of 1(b), 1(c), and 2 use reductions from the *monotone one-in-three SAT* problem. For 1(b), given a monotone one-in-three SAT instance

$I$ , we construct an LIG instance  $J$  having a player for each clause and each variable of  $I$ . Again, we design the influence factors in such a way that  $I$  is satisfiable if and only if  $J$  has a PSNE. The reduction for 1(c) builds upon that of 1(b) with specifically designed *extra players* and additional connectivity in the LIG instance. Again, the gadgets used in the proof of 1(c) are extended for the proof of 2.

The proof of 3 uses reductions from the 3-SAT and the #KNAPSACK problem. The reduction from the 3-SAT problem is the same as that used in 1(a), and proof of the #P-hardness of the bipartite case is by showing that the number of solutions to the 3-SAT instance is the same as the number of PSNE of the LIG instance. On the other hand, to prove the claim of #P-completeness of counting PSNEs of LIGs having star graphs, we give a reduction from the #KNAPSACK problem. Given a #KNAPSACK instance, we create an LIG instance with a star structure among the players and with specifically designed influence factors such that the number of PSNE of the LIG instance is the same as the number of solutions to the #KNAPSACK instance.  $\square$

### 2.3.5 Complete Proofs of Hardness Results

To enhance the clarity of the proofs we have reduced existing NP-complete problems to LIGs with binary actions  $\{0, 1\}$ , instead of  $\{-1, 1\}$ . We next show, via a linear transformation, that any LIG with actions  $\{0, 1\}$  can be reduced to an LIG with the same underlying graph, but with actions  $\{-1, 1\}$ .

**Reduction from  $\{0, 1\}$ -action LIG to  $\{-1, 1\}$ -action LIG.** Consider any  $\{0, 1\}$ -action LIG instance  $I$ , where the influence factors and the thresholds are denoted by the symbols  $w$  and  $b$ , respectively (see Definition 2.2.3). We next construct a  $\{-1, 1\}$ -action LIG instance  $J$  with the same players that are in  $I$  and with influence factors  $w'_{ji} \equiv \frac{w_{ji}}{2}$  (for any  $i$  and any  $j \neq i$ ), thresholds  $b'_i \equiv b_i - \sum_{j \neq i} \frac{w_{ji}}{2}$  (for any  $i$ ). We show that  $\mathbf{x}$  is a PSNE of  $I$  if and only if  $\mathbf{x}'$  is a PSNE of  $J$ , where  $x'_i = 2x_i - 1$  for any  $i$ .

By definition,  $\mathbf{x}$  is a PSNE of  $I$  if and only if for any player  $i$ ,

$$\begin{aligned}
& x_i \left( \sum_{j \neq i} x_j w_{ji} - b_i \right) \geq (1 - x_i) \left( \sum_{j \neq i} x_j w_{ji} - b_i \right) \\
& \Leftrightarrow (2x_i - 1) \left( \sum_{j \neq i} x_j w_{ji} - b_i \right) \geq 0 \\
& \Leftrightarrow x'_i \left( \sum_{j \neq i} \frac{x'_j + 1}{2} w_{ji} - b_i \right) \geq 0 \\
& \Leftrightarrow x'_i \left( \sum_{j \neq i} x'_j \frac{w_{ji}}{2} - \left( b_i - \sum_{j \neq i} \frac{w_{ji}}{2} \right) \right) \geq 0 \\
& \Leftrightarrow x'_i \left( \sum_{j \neq i} x'_j w'_{ji} - b'_i \right) \geq 0,
\end{aligned}$$

which is the equivalent statement of  $\mathbf{x}'$  being a PSNE of  $J$ .  $\square$

**Theorem 2.3.6.** *It is NP-complete to decide if there exists a PSNE in an LIG.*

*Proof.* Since we can verify whether a joint action is a PSNE or not in polynomial time, the problem is in NP. We use a reduction from the 3-SAT problem to show that the problem is NP-hard.

Let  $I$  be an instance of the 3-SAT problem. Suppose that  $I$  has  $m$  clauses and  $n$  variables. For any variable  $i$  we define  $C_i$  to be the set of clauses in which  $i$  appears, and for any clause  $k$  we define  $V_k$  to be the set of variables appearing in clause  $k$ . For any clause  $k$  and any variable  $i \in V_k$ , let  $l_{k,i}$  be 1 if  $i$  appears in  $k$  in non-negated form and 0 otherwise. We now build an LIG instance  $J$  from  $I$ . In this game every clause as well as every variable is a player. Each clause  $k$  has arcs to variables in  $V_k$ , and each variable  $i$  has arcs to clauses in  $C_i$ . The structure of the graph is illustrated in Figure 2.2. We next define the thresholds of the players and the influence factors on the arcs. For any clause  $k$ , let its threshold be  $1 - \epsilon - \sum_{i \in V_k} (1 - l_{k,i})$ . Here,  $\epsilon$  is a constant, and  $0 < \epsilon < 1$ . For any variable  $i$  let its threshold be  $\sum_{k \in C_i} (1 - 2l_{k,i})$ . The weight on the arc from any clause  $k$  to any variable  $i \in V_k$  is defined to be  $1 - 2l_{k,i}$ , and that from any variable  $i$  to any clause  $k \in C_i$  is  $2l_{k,i} - 1$ . We denote the action of any clause  $k$  by  $z_k \in \{0, 1\}$  and that of any variable  $i$  by  $x_i \in \{0, 1\}$ .

First, we prove that if there exists a satisfying truth assignment in  $I$  then there exists a PSNE in  $J$ . Consider any satisfying truth assignment  $S$  in  $I$ .




Figure 2.2: *Illustration of the structure of an LIG instance from a 3-SAT instance (each undirected edge represents two arcs of opposite directions between the same two nodes). In this example, the 3-SAT instance is  $(i_1 \vee i_2 \vee i_3) \wedge (\neg i_2 \vee i_3 \vee i_4) \wedge (\neg i_3 \vee i_4 \vee \neg i_5)$ .*

Let the players in  $J$  choose their actions according to their truth values in  $S$ , that is, 1 for *true* and 0 for *false*. Clearly, every clause player is playing 1. Next, we show that every player in  $J$  is playing its best response under this choice of actions.

We now show that no clause has incentive to play 0, given that the other players do not change their actions. In the solution  $S$  to  $I$ , every clause has a literal that is *true*. Therefore, in  $J$  every clause  $k$  has some variable  $i \in V_k$  such that  $x_i = l_{k,i}$ . We have to show that the total influence on  $k$  is at least the threshold of  $k$ :

$$\begin{aligned}
& \sum_{i \in V_k} x_i (2l_{k,i} - 1) \geq 1 - \epsilon - \sum_{i \in V_k} (1 - l_{k,i}) \\
\Leftrightarrow & \sum_{i \in V_k} (x_i (2l_{k,i} - 1) + (1 - l_{k,i})) \geq 1 - \epsilon \\
\Leftrightarrow & \sum_{i \in V_k} (x_i l_{k,i} + (1 - x_i) (1 - l_{k,i})) \geq 1 - \epsilon.
\end{aligned}$$

Since for some  $i \in V_k$ ,  $x_i = l_{k,i}$ , the above inequality holds strictly, that is,

$$\sum_{i \in V_k} (x_i l_{k,i} + (1 - x_i) (1 - l_{k,i})) > 1 - \epsilon.$$

Therefore, every clause  $k$  must play 1.

We need to show that no variable player has incentive to deviate, given that the other players do not change their actions. The total influence on any variable player  $i$  is  $\sum_{k \in C_i} z_k(1 - 2l_{k,i}) = \sum_{k \in C_i} (1 - 2l_{k,i})$  (since  $z_r = 1$  for every clause  $r$ ). The threshold of  $i$  is  $\sum_{k \in C_i} (1 - 2l_{k,i})$ . Thus, every variable player  $i$  is indifferent between choosing actions 1 and 0 and has no incentive to deviate.

We now consider the reverse direction, that is, given a PSNE in  $J$  we show that there exists a satisfying assignment in  $I$ . We first show that at any PSNE, every clause must play 1. If this is not the case, suppose, for a contradiction, that for some clause  $r$ ,  $z_r = 0$ . Since  $r$ 's best response is 0 (this is a PSNE), we obtain

$$\begin{aligned} \sum_{i \in V_r} x_i(2l_{r,i} - 1) &\leq 1 - \epsilon - \sum_{i \in V_r} (1 - l_{r,i}) \\ \Leftrightarrow \sum_{i \in V_r} (x_i l_{r,i} + (1 - x_i)(1 - l_{r,i})) &\leq 1 - \epsilon. \end{aligned}$$

Therefore, for every variable player  $j \in V_r$ ,  $x_j \neq l_{r,j}$ . Furthermore, for any  $j \in V_r$ ,  $j$  does not have any incentive to deviate. Using these properties of a PSNE we will arrive at a contradiction, and thereby prove that  $z_r$  must be 1.

Consider any variable player  $j \in V_r$ , and let the difference between  $j$ 's total incoming influence and its threshold be  $U_j$ . We get

$$\begin{aligned} U_j &= \sum_{k \in C_j} z_k(1 - 2l_{k,j}) - \sum_{k \in C_j} (1 - 2l_{k,j}) = \sum_{k \in C_j} ((1 - z_k)(2l_{k,j} - 1)) \\ \Leftrightarrow U_j &= \sum_{k \in C_j} ((1 - z_k)(2l_{k,j} - 1)\mathbf{1}[l_{k,j} = 1]) + \sum_{k \in C_j} ((1 - z_k)(2l_{k,j} - 1)\mathbf{1}[l_{k,j} = 0]) \\ \Leftrightarrow U_j &= \sum_{k \in C_j} ((1 - z_k)\mathbf{1}[l_{k,j} = 1]) - \sum_{k \in C_j} ((1 - z_k)\mathbf{1}[l_{k,j} = 0]). \end{aligned}$$

At any PSNE, if  $x_j = 1$  then  $U_j \geq 0$ ; otherwise,  $U_j \leq 0$ . Thus, the best

response condition for variable  $j$  gives us

$$\begin{aligned}
& \sum_{k \in C_j} ((1 - z_k) \mathbf{1}[l_{k,j} = x_j]) \geq \sum_{k \in C_j} ((1 - z_k) \mathbf{1}[l_{k,j} \neq x_j]) \\
\Leftrightarrow & \sum_{k \in C_j - \{r\}} ((1 - z_k) \mathbf{1}[l_{k,j} = x_j]) + (1 - z_r) \mathbf{1}[l_{r,j} = x_j] \geq \\
& \sum_{k \in C_j - \{r\}} ((1 - z_k) \mathbf{1}[l_{k,j} \neq x_j]) + (1 - z_r) \mathbf{1}[l_{r,j} \neq x_j] \\
\Leftrightarrow & \sum_{k \in C_j - \{r\}} ((1 - z_k) \mathbf{1}[l_{k,j} = x_j]) \geq \\
& \sum_{k \in C_j - \{r\}} ((1 - z_k) \mathbf{1}[l_{k,j} \neq x_j]) + 1, \text{ since } l_{r,j} \neq x_j.
\end{aligned}$$

The above inequality cannot be true, because the left hand side is always 0 (if  $l_{k,j} = x_j$  then  $z_k$  must be 1 at any PSNE), and the right hand side is  $\geq 1$ . Thus, we have obtained a contradiction, and  $z_r$  cannot be 0.

So far, we have shown that at any PSNE  $z_k = 1$  for any clause player  $k$ . To complete the proof, we now show that for every clause player  $k$ , there exists a variable player  $i \in V_k$  such that  $x_i = l_{k,i}$ . If we can show this then we can translate the semantics of the actions in  $J$  to the truth values in  $I$  and thereby obtain a satisfying truth assignment for  $I$ .

Suppose, for the sake of a contradiction, that for some clause  $k$  and for all variable  $i \in V_k$ ,  $x_i \neq l_{k,i}$ . Since  $z_k = 1$ , we find that

$$\begin{aligned}
& \sum_{i \in V_k} x_i(2l_{k,i} - 1) \geq 1 - \epsilon - \sum_{i \in V_k} (1 - l_{k,i}) \\
\Leftrightarrow & \sum_{i \in V_k} (x_i l_{k,i} + (1 - x_i)(1 - l_{k,i})) \geq 1 - \epsilon \\
\Leftrightarrow & 0 \geq 1 - \epsilon, \text{ which gives us the desired contradiction.}
\end{aligned}$$

□

The proof of Theorem 2.3.6 reduces the 3-SAT problem to an LIG where the underlying graph is bipartite. Thus, we obtain the following corollary.

**Corollary 2.3.7.** *It is NP-complete to decide if there exists a PSNE in an LIG on a bipartite graph.*

The proof of Theorem 2.3.6 directly leads us to the following result that the counting version of the problem is #P-complete.

**Corollary 2.3.8.** *It is #P-complete to count the number of PSNE of an LIG.*

*Proof.* The proof follows from the proof of Theorem 2.3.6. Membership of this counting problem in #P is easy to see. Using the same reduction as in the proof of Theorem 2.3.6, we find that each satisfying truth assignment (among the  $2^n$  possibilities) to the variables of the 3-SAT instance  $I$  can be mapped to a distinct PSNE of the LIG instance  $J$ . Furthermore, we have seen that at each PSNE in  $J$ , every clause player must play 1. Thus, for each of the  $2^n$  joint strategies of the variable players (while having the clause players play 1), if the joint strategy is a PSNE then we can map it to a distinct satisfying assignment in  $I$ . Moreover, each of these two mappings are the inverse of the other. Therefore, the number of satisfying assignments of  $I$  is the same as the number of PSNE in  $J$ . Since counting the number of satisfying assignments of a 3-SAT instance is #P-complete, counting the number of PSNE of an LIG, even on a bipartite graph, is also #P-complete.  $\square$

While Corollary 2.3.8 shows the hardness of counting the number of PSNE of an LIG on a general graph, we can show the same hardness result even on special classes of graphs, such as star graphs:

**Theorem 2.3.9.** *Counting the number of PSNE of an LIG on a star graph is #P-complete.*

*Proof.* Since we can verify whether a joint strategy is a PSNE in polynomial time, the problem is in #P. We will show #P-hardness using a reduction from #KNAPSACK, which is the problem of counting the number of feasible solutions in a 0-1 Knapsack problem: Given  $n$  items, the weight  $a_i \in \mathcal{Z}^+$  of each item  $i$ , and the maximum capacity of the sack  $W \in \mathcal{Z}^+$ , #KNAPSACK asks how many ways we can pick the items to satisfy  $\sum_{i=1}^n a_i x_i \leq W$ , where  $x_i = 1$  if the  $i$ -th item has been picked, and  $x_i = 0$  otherwise. Given an instance  $I$  of the #KNAPSACK problem with  $n$  items, we construct an LIG instance  $J$  on a star graph with  $n + 1$  nodes. Let us label the nodes  $v_0, \dots, v_n$ , where  $v_0$  is connected to all other nodes. We define the influence factors among the nodes as follows: the influence of  $v_0$  to any other node  $v_i$ ,  $w_{v_0 v_i} = 1$ , and the influence in the reverse direction,  $w_{v_i v_0} = -a_i$ . The threshold of  $v_0$  is defined as  $b_{v_0} = -W$ , and the threshold of every other node  $v_i$ ,  $b_{v_i} = 1$ . We denote the action of any node  $v_i$  by  $x_i \in \{0, 1\}$ . Note that at any PSNE of  $J$ ,  $v_0$  must play 1. Otherwise, if  $v_0$  plays 0 then all other nodes must also play 0, and this implies that  $v_0$  must play 1, giving us a contradiction.

We prove that the number of feasible solutions in  $I$  is the same as the number of PSNE in  $J$ . For any  $(x_1, \dots, x_n) \in \{0, 1\}^n$  in  $I$ , we map each  $x_i$  to the action selected by  $v_i$  in  $J$ , for  $1 \leq i \leq n$ . As proved earlier, the action



of  $v_0$  must be 1 at any PSNE. Furthermore, when  $v_0$  plays 1, all other nodes become indifferent between playing 0 and 1. Thus, the number of PSNE in  $J$  is the number of ways of satisfying the inequality  $\sum_{i=1}^n w_{v_i v_0} x_i \geq b_{v_0}$ , which is equivalent to  $\sum_{i=1}^n a_i x_i \leq W$ . Thus the number of PSNE in  $J$  is equal to the number of feasible solutions in  $I$ .  $\square$

The following three theorems show the hardness of several other variants of the problem of computing a PSNE of an LIG.

**Theorem 2.3.10.** *Given an LIG, along with a designated subset of  $k$  players in it, it is NP-complete to decide if there exists a PSNE consistent with those  $k$  players playing the action 1.*

*Proof.* It is easy to see that the problem is in NP, since a succinct yes certificate can be specified by a joint action of the players, where the designated players play 1, and it can be verified in polynomial time whether this is a PSNE or not.

We show a reduction from the monotone one-in-three SAT problem, a known NP-complete problem, to prove that the problem is NP-hard. An instance of the monotone one-in-three SAT problem consists of a set of  $m$  clauses and a set of  $n$  variables, where each clause has exactly three variables. The problem asks whether there exists a truth assignment to the variables such that each clause has exactly one variable with the truth value of *true*. Given an instance of the monotone one-in-three SAT problem, we construct an instance of LIG as follows (please refer to Figure 2.3 for an illustration). For each variable we have a *variable player* in the game, and for each clause we have a *clause player*. Each variable player has a threshold of 0, and each clause player has a threshold of  $\epsilon$ , where  $0 < \epsilon < 1$ . We now define the connectivity among the players of the game. There is an arc with weight (or influence)  $-1$  from a variable player  $u$  to another variable player  $v$  if and only if, in the monotone one-in-three SAT instance, both of the corresponding variables appear together in at least one clause. Also, for each clause  $t$  and each variable  $w$  appearing in  $t$ , there is an arc from the variable player (corresponding to  $w$ ) to the clause player (corresponding to  $t$ ) with weight 1. Furthermore, we assign  $k = m$ , and assume that the designated set of players is the set of clause players. We also assume that the action 1 in the LIG corresponds to the truth value of *true* in the monotone one-in-three SAT problem and 0 to *false*.

Note that the way we have constructed the LIG, at most one variable player per clause can play the action 1 at any PSNE. To see this, assume, for contradiction, that at some PSNE two variable players  $u$  and  $v$ , both connected to the same clause  $t$ , are playing the action 1. Then the influence on either of these two variable players is  $\leq -1$ , which is less than its threshold 0, and this




Figure 2.3: *Illustration of the NP-hardness reduction of Theorem 2.3.10. The monotone one-in-three SAT instance is  $(i_1 \vee i_2 \vee i_3) \wedge (i_2 \vee i_3 \vee i_4) \wedge (i_3 \vee i_4 \vee i_5)$ . The threshold of each variable player is 0, and that of each clause player is  $\epsilon$ .*

contradicts the PSNE assumption. Also, note that at any PSNE, each clause player will play the action 1 if and only if at least one of the variable players connected to it plays 1.

First, we show that if there exists a solution to the monotone one-in-three SAT instance then there exists a PSNE in the LIG where the set of clause players play 1. A solution to the monotone one-in-three SAT problem implies that each clause has the truth value of *true* with exactly one of its variables having the truth value of *true*. We claim that in the LIG, every player playing according to its truth assignment, is a PSNE. First, observe that the variable players do not have any incentive to change their actions, since the ones playing 1 are indifferent between playing 0 and 1 (because the total influence = 0 = threshold) and the remaining must play 0 (because the total influence is  $\leq -1 < \text{threshold}$ ). Since each clause has one of its variables playing 1, each clause player must play 1 (because  $1 > \epsilon$ ). This concludes the first part of the proof.

We next show that if there exists a PSNE with the clause players playing 1 then there exists a solution to the monotone one-in-three SAT instance. Consider any PSNE where the clause players are playing 1. Since each clause player is playing 1, *at least* one of the three variable players connected to the clause player is playing 1. Furthermore, as we have shown earlier, no two variables belonging to the same clause can play 1 at any PSNE. Thus, for each clause player, *at most* one variable player connected to it is playing 1. Therefore, for every clause player, exactly one variable player connected to it

is playing 1. Translating the semantics of the actions to the truth values of the variables and the clauses, we obtain a solution to the monotone one-in-three SAT instance.  $\square$

**Theorem 2.3.11.** *Given an LIG and a number  $k \geq 1$ , it is NP-complete to decide if there exists a PSNE with at least  $k$  players playing the action 1.*

*Proof.* Clearly, the problem is in NP, since we can verify a whether a joint action is a PSNE or not in polynomial time.

For the proof of NP-hardness, once again we show a reduction from the monotone one-in-three SAT problem. Please see Figure 2.4 for an illustration. Given an instance  $I$  of the monotone one-in-three SAT problem, we first build an LIG as shown in the proof of Theorem 2.3.10. We then add  $m(m - 1)$  additional players, named *extra players*, to the game, where  $m$  is the number of clauses in  $I$ . Each of these extra players is assigned a threshold of  $\epsilon$ , where  $0 < \epsilon < 1$ . The way we connect the extra players to the other players is as follows: From each clause player we introduce  $m - 1$  arcs, each weighted by 1, to  $m - 1$  distinct extra players. That is, no two clause players have arcs to the same extra player. Finally, we set  $k = m^2$ . We denote this instance of LIG by  $J$ .




Figure 2.4: *Illustration of the NP-hardness reduction of Theorem 2.3.11.*

We prove that for any solution to  $I$  there exists a PSNE with  $k$  players playing 1 in  $J$ . Suppose that each of the variable and clause players is playing according to their corresponding truth value in the solution to  $I$ . None of the

variable players has any incentive to change its action, because exactly one variable player connected to each clause player is playing 1. For the same reason, the clause players, each playing 1, also do not have any incentive to deviate. Considering the extra players, each of these players must play 1, because each of the clause players is playing 1. The total number of clause and extra players is  $k$ . Therefore, we have a Nash equilibrium where at least  $k$  players are playing 1.

On the other direction, consider any PSNE in  $J$  with at least  $k$  players playing 1. We claim that all the clause and extra players are playing 1 at this PSNE. If this is not true then at least one of these players is playing 0. This implies that at least one clause player is playing 0, because conditioned on a PSNE, whenever a clause player plays 1, all the extra players connected to it also plays 1. Furthermore, by our construction at most one of the variable players connected to each clause player can play 1. So, the total number of players playing 1 is  $\leq (m - 1)(m + 1) < m^2$  (at most  $m - 1$  clause players are playing 1, and for each of these clause players,  $m - 1$  extra players, 1 variable player, and the clause player itself are playing 1), which contradicts our assumption that  $m^2$  players are playing 1. Thus, at any PSNE with  $k$  players playing 1, it must be the case that every clause player is playing 1. This leads us to a solution for  $I$ .  $\square$

**Theorem 2.3.12.** *Given an LIG and a designated set of  $k \geq 1$  players, it is co-NP-complete to decide if there exists a unique PSNE with those players playing the action 1.*

*Proof.* Two distinct joint actions (PSNE), each having the same  $k$  players playing 1, can serve as a succinct no certificate, and we can check in polynomial time if these two joint actions are indeed PSNE or not.

Suppose that  $I$  is an instance of the monotone one-in-three SAT problem. We reduce  $I$  to an instance  $J$  of our problem in polynomial time and show that  $J$  has a “no” answer if and only if  $I$  has a “yes” answer.

Given  $I$ , we start constructing an LIG in the same way as in Theorem 2.3.10 (see Figures 2.5 and 2.3). Assign  $k = m^2$ . Now, add two new players, named the *all-satisfied-verification player* and the *none-satisfied-verification player*, which have threshold values of  $m - \epsilon$  and  $-\epsilon$ , respectively. We add arcs from every clause player to these two new players, and the arcs to the all-satisfied-verification player are weighted by 1, and the ones to the none-satisfied-verification player are weighted by  $-1$ .

In addition, add  $k = m^2$  new players, named *extra players*, and let these players constitute the set of designated players. Assign a threshold value of  $\epsilon$  to each of these extra players, and introduce new arcs, each with weight 1, from the all-satisfied-verification player and the none-satisfied-verification player to

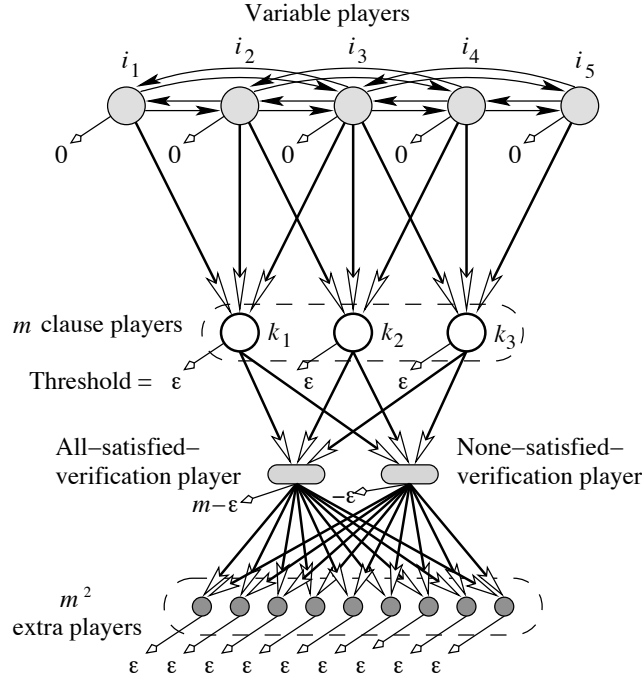


Figure 2.5: *Illustration of the NP-hardness reduction (Theorem 2.3.12). For the monotone one-in-three SAT instance of Figure 2.3, we first obtain the same construction as in Theorem 2.3.10. We add two extra players, the all-satisfied-verification player and the none-satisfied-verification player, whose tasks are to verify if all clauses are satisfied and if no clause is satisfied, respectively. These two players are connected to  $m^2$  extra players.*

every extra player. The resulting LIG is the instance  $J$  of the problem in question.

Note that at any PSNE the all-satisfied-verification player plays 1 if and only if every clause player plays 1, and the none-satisfied-verification player plays 1 if and only if no clause player plays 1. Furthermore, at any PSNE, each extra player plays 1 if and only if either every clause player plays 1 or no clause player plays 1. Therefore, we find that every extra player playing 1, the none-satisfied-verification player playing 1, and every other player playing 0 is a PSNE, and we denote this equilibrium by  $E_0$ . We claim that there exists a different PSNE where every extra player plays 1 if and only if  $I$  has a solution.

Suppose that there exists a solution  $S_I$  to  $I$ . It can be verified that making the all-satisfied-verification player play 1, none-satisfied-verification player play 0, every extra player play 1, and choosing the actions of the clause and the variable players according to the corresponding truth values in  $S_I$  gives us a

PSNE that we call  $E_1$ . Thus  $J$  has two PSNE  $E_0$  and  $E_1$ , where the  $k$  extra players play 1 in both cases.

Considering the reverse direction, suppose that there exists no solution to  $I$ . This implies that at any PSNE in  $J$  all clause players can never play 1, otherwise we could have translated the PSNE to a satisfying truth assignment for  $I$ . This further implies that the all-satisfied-verification player always plays 0. The none-satisfied-verification player plays 1 if and only if none of the clause players plays 1. Thus, every extra player plays 1 if and only if no clause player plays 1, if and only if no variable player plays 1. Therefore,  $E_0$  is the only PSNE in  $J$  with the  $k$  extra players playing 1.  $\square$

### 2.3.6 Heuristics for Computing and Counting Equilibria

The fundamental computational problem at hand is that of computing PSNE in LIGs. We have just seen that various computational questions pertaining to LIGs on general graphs, sometimes even on bipartite graphs, are NP-hard. We now present a heuristic to compute PSNE of an LIG on a general graph.

A natural approach to finding all the PSNE in an LIG would be to perform a backtracking search. However, a naive backtracking method that does not consider the structure of the graph would be destined to failure in practice. Thus, we need to order the node selections in a way that would facilitate pruning the search space.

The following is an outline of a backtracking search procedure that we have used in practice. The first node selected by the procedure is a node with the maximum outdegree. Intuitively, this node is the “most constraining” (see, e.g., Chapter 5 of [105]) in terms of the number of nodes that a node directly influences. Subsequently, we select a node  $i$  that will most likely show that the current partial joint action cannot lead to a PSNE and explore the two actions of  $i$ ,  $x_i \in \{-1, 1\}$  in a suitable order. A good node selection heuristic that has worked well in our experiments is to select the one that has the maximum influence on any of the already selected nodes.

Suppose that the nodes are selected in the order  $1, 2, \dots, n$  (wlog). After selecting node  $i + 1$  and assigning it an action  $x_{i+1}$ , we determine if the partial joint action  $\mathbf{x}_{1:(i+1)} \equiv (x_1, \dots, x_{i+1})$  can possibly lead to a PSNE and prune the corresponding search space if not. Note that a “no” answer to this requires a proof that one of the players  $j$ ,  $1 \leq j \leq i + 1$ , can never play  $x_j$  according to the partial joint action  $\mathbf{x}_{1:(i+1)}$ . A straightforward way of doing this is to consider each player  $j$ ,  $1 \leq j \leq i + 1$ , and compute the quantities  $\gamma_j^+ \equiv \sum_{k=1, k \neq j}^{i+1} x_k w_{kj} + \sum_{k=i+2}^n |x_k w_{kj}|$  and  $\gamma_j^- \equiv \sum_{k=1, k \neq j}^{i+1} x_k w_{kj} - \sum_{k=i+2}^n |x_k w_{kj}|$ ,

and then test if the logical expression  $((\gamma_j^- > b_j) \wedge x_j = -1) \vee ((\gamma_j^+ < b_j) \wedge x_j = 1)$  holds, in which case we can discard the partial joint action  $\mathbf{x}_{1:(i+1)}$  and prune the corresponding search space. Furthermore, it may happen that due to  $\mathbf{x}_{1:(i+1)}$ , the choices of actions of some not-yet-selected players have become restricted. To this end, we apply **NashProp** [93] with  $\mathbf{x}_{1:(i+1)}$  as the starting configuration, and see if the choices of the other players have become restricted because of  $\mathbf{x}_{1:(i+1)}$ . Although each round of updating the table messages in **NashProp** takes exponential time in the maximum degree in general graphical games, we can show in a way similar to Theorem 2.3.4 that we can adapt the table updates to the case of LIGs so that it takes polynomial time.

## A Divide-and-Conquer Approach

To further exploit the structure of the graph in computing the PSNE, we propose a divide-and-conquer approach that relies on the following separation property of LIGs.

**Property 2.3.13.** *Let  $G = (V, E)$  be the underlying graph of an LIG and  $S$  be a vertex separator of  $G$  such that removing  $S$  from  $G$  results in  $k \geq 2$  disconnected components:  $G_1 = (V_1, E_1), \dots, G_k = (V_k, E_k)$ . Let  $G'_i$  be the subgraph of  $G$  induced by  $V_i \cup S$ , for  $1 \leq i \leq k$ . Consider the LIGs on these (smaller) graphs  $G'_i$ 's, where we retain all the weights of the original graph, except that we treat the nodes in  $S$  to be indifferent (that is, we remove all the incoming arcs to these nodes and set their thresholds to 0). Computing the set of PSNE on  $G'_i$ 's and then merging the PSNE (by performing outer-joins of joint actions and testing for PSNE in the original LIG), we obtain the set of all PSNE of the original game.*

*Proof Sketch.* First, since the joint actions are tested for PSNE in the original LIG, the output will never contain a joint action that is not a PSNE. Second, since the nodes in  $S$  are made indifferent in the LIGs on  $G'_i$ ,  $1 \leq i \leq k$ , no PSNE of the original LIG can get omitted from the result of the outer-join operation.  $\square$

To obtain a vertex separator, we first find an edge separator (using well-known tools such as METIS [65]), and then convert the edge separator to a vertex separator (by computing a maximum matching on the bipartite graph spanned by the edge separator). We then use this vertex separator to compute PSNE of the game in the way outlined in Property 2.3.13. The benefits of this approach are two-fold: (1) for graphs that have good separation properties (such as preferential-attachment graphs), we have found this approach to be computationally effective in practice; and (2) this approach leads to an *anytime*

*algorithm* for enumerating or counting PSNE: Observe that ignoring some edges from the edge separator may result in a smaller vertex separator, which greatly reduces the computation time of the divide-and-conquer algorithm at the expense of producing only a subset of all PSNE. (The reason we obtain a subset of all PSNE is that the edges that are ignored from the edge separator are not permanently removed from the original graph, and that after merging, every resulting joint action is tested for PSNE in the *original* game, not in the game where some of the edges were temporarily removed. As a result, some the original PSNE may not be included in the final output. At the same time, we can never have a joint action in the final output that is not a PSNE.) We can obtain progressively better result as we ignore less number of edges from the edge separator.

## 2.4 Computing the Most Influential Nodes

We now focus on the problem of computing the most influential set of nodes with respect to a specified desirable PSNE and a preference for sets of minimal size. In the discussion below, we also assume, *only* for the purpose of establishing and describing the equivalence to the *minimum hitting set problem* [64], that we are given the set of all PSNE. (As we will see, a counting routine is all that our algorithm requires, not a complete list of PSNE.) We give a hypergraph representation of this problem that would lead us to a logarithmic-factor approximation by a natural greedy algorithm.

Let us start by building a hypergraph that can represent the PSNE of a binary-action game. The nodes of this hypergraph are the player-action tuples of the game. Thus, for an  $n$ -player, binary-action game, we have  $2n$  nodes in the hypergraph. That is, for each player  $i$  of the game, there are two nodes in the hypergraph: one in which  $i$  plays  $-1$  (tuple  $(i, -1)$ , colored red in Figure 2.6) and the other in which  $i$  plays  $1$  (tuple  $(i, 1)$ , colored black). For every PSNE  $\mathbf{x}$  we construct a hyperedge  $\{(i, x_i) \mid 1 \leq i \leq n\}$ . Let us call this hypergraph the *game hypergraph*. By construction, a set of players  $S$  play the same joint-action  $\mathbf{a}_S \in \{-1, 1\}^{|S|}$  in two distinct PSNE  $\mathbf{x}$  and  $\mathbf{y}$  of the LIG if and only if both of the corresponding hyperedges  $e_{\mathbf{x}}$  and  $e_{\mathbf{y}}$  (resp.) of the game hypergraph contains  $T = \{(i, a_i) \mid i \in S\}$ .

We can use the above property to translate the most influential nodes selection problem, given all PSNE, to an equivalent combinatorial problem on the corresponding game hypergraph  $H$ . Let  $e_{\mathbf{x}^*}$  be the hyperedge in  $H$  corresponding to the desirable PSNE  $\mathbf{x}^*$ . Let us call  $e_{\mathbf{x}^*}$  the *goal hyperedge*. Then the most influential nodes selection problem is the problem of selecting a minimum-cardinality set of nodes  $T \subseteq e_{\mathbf{x}^*}$  such that  $T$  is contained in no






Figure 2.6: A hypergraph representation of three PSNE in a 9-player game with binary actions. The PSNE shown here are the followings:  $(1, -1, -1, 1, -1, -1, 1, -1, -1)$  (triangle),  $(-1, -1, -1, -1, -1, -1, 1, 1, 1)$  (rectangle), and  $(-1, -1, -1, -1, 1, -1, 1, 1, 1)$  (6-gon).

other hyperedge of  $H$  (recall that we are dealing with a set-preference function that captures the preference for sets of minimal cardinality). Let us call the latter problem the *unique hyperedge problem*. Using the notation above, the equivalence relationship between the influential nodes selection problem (given the set of all PSNE) and the unique hyperedge problem can be stated as follows. The set  $S \subseteq \{1, \dots, n\}$  is a (feasible) solution to the most influential nodes selection problem if and only if  $T = \{(i, x_i^*) \mid i \in S\}$  is a (feasible) solution to the unique hyperedge problem.

We now show that the unique hyperedge problem is equivalent to the minimum hitting set problem. Immediate consequences of this result are that the unique hyperedge problem is not approximable within a factor of  $c \log h$  for some constant  $c > 0$ , and that it admits a  $(1 + \log h)$ -factor approximation [61, 103], where  $h$  is the total number of PSNE.

**Theorem 2.4.1.** *The unique hyperedge problem having  $2n$  players and  $h$  hyperedges is equivalent to the minimum hitting set problem having  $n$  nodes and  $h$  hyperedges.*

*Proof.* Let us consider an instance  $I$  of the unique hyperedge problem, given by a game hypergraph  $G = (V, E)$ , where  $V$  is the set of  $2n$  nodes and  $E$  is the set of  $h$  hyperedges, along with a specification of the goal hyperedge  $e_{\mathbf{x}^*}$ . Given  $I$ , we now construct an instance  $J$  of the minimum hitting set problem,

specified by the hypergraph  $G' = (e_{x^*}, \{e_{x^*}\} \cup \{\bar{e} \cap e_{x^*} \mid e \in E \text{ and } e \neq e_{x^*}\})$ , where  $\bar{e}$  indicates the complement set of the hyperedge  $e$ . Thus, the nodes of  $G'$  are exactly the  $n$  nodes of  $e_{x^*}$  and the hyperedges of it are constructed from the complement hyperedges of  $G$  except  $e_{x^*}$ , which is present in both  $G$  and  $G'$ . We show that a set  $S$  of nodes is a feasible solution to  $I$  if and only if it is a feasible solution to  $J$ .

If  $S$  is a feasible solution to  $I$  then  $S \subseteq e_{x^*}$  (because in the unique hyperedge problem, we are only allowed to select nodes from the goal hyperedge) and  $S \not\subseteq e$  for any hyperedge  $e \neq e_{x^*}$  of  $G$  (otherwise, the uniqueness property is violated). This implies that for any hyperedge  $e \neq e_{x^*}$  of  $G$ , there exists a node  $v \in S$  such that  $v \notin e$ , which further implies that  $v \in \bar{e} \cap e_{x^*}$ . Thus, every hyperedge of  $G'$ , including  $e_{x^*}$ , of course, has at least one of its nodes selected in  $S$ , and therefore,  $S$  is a feasible solution to  $J$ . On the other hand, if  $S$  is a feasible solution to  $J$  then for any hyperedge of  $G'$ , at least one of its nodes has been selected in  $S$ . That is, for any hyperedge  $e \neq e_{x^*}$  of  $G$ , we have  $e' \equiv \bar{e} \cap e_{x^*}$  as the corresponding complementary hyperedge in  $G'$ , and there exists a node  $v \in S$  such that  $v \in e'$ , which implies that  $v \notin e$ . Thus,  $S \not\subseteq e$  for any hyperedge  $e \neq e_{x^*}$  of  $G$ . Furthermore, all the nodes of  $S$  have been selected from  $e_{x^*}$  of  $G$ . Thus,  $e_{x^*}$  is the unique hyperedge of  $G$  containing the nodes of  $S$ .

To prove the reverse direction, we start with an instance  $J$  of the minimum hitting set problem, specified by a hypergraph  $G' = (V, E)$ , where  $V$  is a set of  $n$  nodes and  $E$  is a set of  $h$  hyperedges. Without the loss of generality, we assume that  $E$  contains the hyperedge  $e^*$  consisting of all the nodes of  $V$ . We now construct an instance  $I$  of the unique hyperedge problem that has a hypergraph  $G$  with  $2n$  nodes and  $h$  hyperedges. The node set of  $G$  literally consists of two copies of the nodes of  $V$ , denoted by  $V \times \{1, -1\}$ . We now construct the hyperedges of  $G$ . For each hyperedge  $e \neq e^*$  of the minimum hitting set instance, we include a hyperedge  $e' \equiv \bar{e} \times \{1\} \cup e \times \{-1\}$  in  $G$ , and for the hyperedge  $e^*$  of  $J$ , we include the hyperedge  $e^* \times \{1\}$  in  $G$ . Thus, the game hypergraph can be defined as  $G = (V \times \{1, -1\}, \{e^* \times \{1\}\} \cup \{\bar{e} \times \{1\} \cup e \times \{-1\} \mid e \in E \text{ and } e \neq e^*\})$ . Finally, we designate  $e^* \times \{1\}$  as the goal hyperedge of  $I$ . We will show that  $S \subseteq V$  is a feasible solution to  $J$  if and only if  $S \times \{1\}$  is a feasible solution to the unique hyperedge problem instance  $I$ . The set  $S$  is a feasible solution to  $J$  if and only if for every hyperedge  $e \neq e^*$  of  $G'$ , there exists a node  $v \in S$  such that  $v \in e$  (note that  $S \subseteq e^*$ ). This is equivalent to saying that for every hyperedge  $e \times \{1\} \neq e^* \times \{1\}$  of  $G$ , there exists a node  $v \in S \times \{1\}$  such that  $v \notin e \times \{1\}$ . Using the fact that  $S \times \{1\} \subseteq e^* \times \{1\}$ ,  $S \times \{1\}$  is a feasible solution to  $I$ .  $\square$

The adaptation of the well-known hitting set approximation algorithm for

our problem can be outlined as follows: At each step, select the least-degree node  $v$  of the goal hyperedge, remove the hyperedges that do not contain  $v$ , remove  $v$  from the game hypergraph, and include  $v$  in the solution set, until the goal hyperedge becomes the last remaining hyperedge in the hypergraph. In the context of the original LIG, at every round, this algorithm is essentially picking the node whose assignment would reduce the set of PSNE consistent with the current partial assignment the most. Hence, the algorithm only requires a subroutine to *count* the PSNE extensions for some given partial assignment to the players’ actions, not an *a priori* full list or enumeration of all the PSNE. Of course, it may require a complete list of PSNE in the worst case.

## 2.5 Experimental Results

We have performed empirical studies on several types of LIGs, namely, random LIGs, preferential-attachment LIGs, LIGs created to model potential interactions in two different real-world scenarios: those among the U.S. Supreme Court Justices, and those among the U.S. senators. While the first two types of LIGs have been constructed artificially, the latter two have been learned from real-world data using machine learning techniques [55].

### 2.5.1 Random Influence Games

As a first attempt, we have created instances of random graphs using the Erdős-Rényi model. The number of nodes have been varied from 10 to 30, and the probability of including an edge has also been varied. Assuming binary actions: 1 and  $-1$ , the threshold  $b_i$  and the influence factors  $w_{ji}$  of the incoming arcs of each node  $i$  have been chosen uniformly at random from a unit hyperball. That is, for each node  $i$ ,  $b_i^2 + \sum_{j \in N(i)} w_{ji}^2 = 1$ , where  $N(i)$  is the set of nodes having arcs toward  $i$ . Then, the sign of each threshold, as well as each weight, has been chosen to be either  $+$  or  $-$  with a probability of 0.5. We have applied the heuristic given earlier to find the set of all PSNE in these random graphs. Our experiments show that in all of these random LIGs, the number of PSNE, almost always, is very small—usually one or two, and sometimes none.

We have also studied LIGs on uniform random directed graphs. While constructing the random graphs, we have independently chosen each arc with a probability of 0.50, and assigned it a weight of  $-1$  with a probability  $p$  (named *flip probability*) and 1 with probability  $1 - p$ . Several interesting findings have emerged from our study of this parameterized family of LIGs on

uniform random graphs. The results are summarized in tabular forms in the Appendix. For various flip probabilities, we have independently generated 100 uniform random graphs of 25 nodes each, and for each of these random graphs, we have first computed all PSNE using our heuristic. We have then applied the greedy approximation algorithm to obtain a set of the most influential nodes in each graph and compared the approximation results to the optimal ones.

Unless  $p$  is either 0 or 1, the existence of a PSNE cannot be guaranteed. In our experiments, we have found that, in fact, for  $p = 0.50$ , the probability of not having a PSNE is highest (around 5%), and as we go toward the two extremes of  $p$ , the probability of not having a PSNE decreases. We have reported the games with at least one equilibrium in this experimental study, since these are the games that we are interested in for computing the most influential nodes. Another interesting finding with respect to the number of PSNE is that this number is very small when  $p = 0$ , that is when all the arcs have weight 1, and it is large when  $p = 1$ , although quite small (on average, a fraction  $5.81 \times 10^{-6} \approx 2^{-17.29}$ ) relative to the total number of  $2^{25}$  possible joint actions. Also, the average number of nodes of the search tree that the backtracking method visits per equilibrium computation is relatively small on the two extremes of  $p$ , compared with  $p$  around 0.5. Note that the backtracking method does a very good job with respect to the number of search-tree nodes visited in searching the  $2^{25}$  space. In fact, our experiments have shown that the addition of the NashProp heuristic on top of the node selection heuristic considerably speeds up the search. Finally, we have found that although the approximation algorithm has a logarithmic factor worst-case bound, most often the results of the approximation algorithm are very close to the optimal solution.

As shown in Figure 2.7, the number of PSNE usually increases if we have more negative-weighted arcs than positive ones, although the number of PSNE is still very small relative to the maximum potential number as remarked earlier. We have further found that although the approximation algorithm for influential nodes selection problem has a logarithmic factor worst-case bound, most often the result of the approximation algorithm is very close to the optimal solution. For example, for the random games having all negative influence factors, in 87% of the trials the approximate solution size  $\leq$  optimal size +1, and in 99% of the trials the approximate solution size  $\leq$  optimal size +2 (see the Appendix for more details in a tabular form).




Figure 2.7: *PSNE computation on random LIGs. The vertical bars denote 95% confidence intervals.*

## 2.5.2 Preferential-Attachment LIGs

We have also experimented with LIGs based on preferential-attachment graphs primarily because of its power to explain the structure of many real-world social networks in a generative fashion [2]. In order to construct these graphs, we have started with three nodes in a triangle and then progressively added each node to the graph, connecting it with three existing nodes with probabilities proportionate to the degrees. We have made each connection bidirectional and imposed the same weighting scheme as above: with the flip probability  $p$ , the weight of an arc is  $-1$  and with probability  $1 - p$  it is  $1$ . The threshold of each node has been set to  $0$ . We have observed that for  $0 < p < 1$ , these games have very few PSNE, while for  $p = 0$  and  $p = 1$  the number of PSNE is considerably larger than that. Furthermore, these games show very good separation properties, making the computation amenable to the divide-and-conquer approach. We show the average number of PSNE and the average computation time for graphs of sizes 20 to 50 nodes in Figure 2.8 for  $p = 1$  (each average is over 20 trials). Note that in contrast to the random LIGs, preferential-attachment graphs show an exponential increase in the number of PSNE as the number of nodes increase, although the number of PSNE is still a very small fraction of the maximum potential number.

## 2.5.3 Illustration: Supreme Court Rulings

We have used our model to analyze the influence among the Justices of the U.S. Supreme Court. This is one of the application scenarios where the strate-

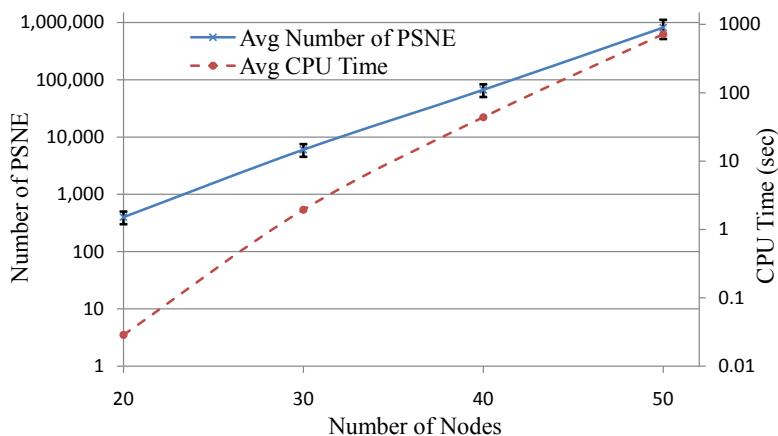


Figure 2.8: *PSNE* computation on preferential-attachment LIGs (*y*-axis is in log scale). The vertical bars denote 95% confidence intervals.

gic aspects of influence is of prime importance. Two distinctive features of such a scenario are: first, the individual outcomes (in this case, the decisions of the Justices on each case) can be modeled as outcomes of a one-shot non-cooperative game (which in our case is LIG), and second, the physical interpretation of the diffusion process is not so much clear in such a scenario as it is in applications like viral marketing.

## Data

We have obtained data from the Supreme Court Database.<sup>2</sup> Although the database captures fine-grained details of the cases, for our purpose we have only focused on the variable `varVote`. Again, the votes of the Justices are not simple yes/no instances. Instead, each vote can have eight distinct values. However, for practical purposes, we can attach a simple yes/no interpretation to the values of the votes, as shown in Table 2.1. For example, some of the votes are interpreted as the majority vote, but in principle, if all the votes are yes/no, then such an interpretation is never required.

In Table 2.1, “majority” in the third column signifies that we have interpreted the corresponding Justice’s vote as yes or no, whichever occurs most among the other Justices. Also, among the natural courts we studied, we did not encounter voting instances where `varVote` has a value of 8.

We next present our study of the natural court (with timeline 1994–2004) comprising of Justices WH Rehnquist, JP Stevens, SD O’Connor, A Scalia, AM Kennedy, DH Souter, C Thomas, RB Ginsburg, and SG Breyer.

<sup>2</sup><http://scdb.wustl.edu/>

varVote	Original Meaning	Our Interpretation
1	Voted with majority	Yes
2	Dissent	No
3	Regular concurrence	Yes
4	Special concurrence	Yes
5	Judgment of the Court	Yes
6	Dissent from a denial or dismissal of certiorari, or dissent from summary affirmation of an appeal (Interpreted as absent from voting in final outcome)	Majority
7	Jurisdictional dissent (Interpreted as absent from voting in final outcome)	Majority
8	Justice participated in an equally divided vote	—

Table 2.1: *Interpretation of Votes*

## Learning LIG

The data for the above natural court consists of 971 voting instances (each voting instance consists of the votes of all nine Justices). Many of these instances are repeated. For example, the most repeated instance is where all the Justices voted yes, which occurred 438 times. The second most repeated instance, which occurred 85 times, is where five of the Justices, namely, Justices Scalia, Thomas, Rehnquist, O’Connor, and Kennedy voted yes, while the others voted no.

We have used  $L_2$ -regularized logistic regression (simultaneous classification) to learn an LIG for this data. In this learning procedure, the data is viewed as a noisy version of PSNE. The objective of learning here is to capture as much of the observations as PSNE of the model while minimizing the number of PSNE that were not observed in data. The regularization parameter prefers a sparser graph compared to a dense one. A more detailed exposition of this learning technique, along with theoretical justifications, can be found in [55]. The influence factors and the biases of the LIG learned in this way are shown in a tabular form in the Appendix. A pictorial representation of the same LIG is shown in Figure 2.9.

The learned LIG represents 589 of the 971 voting instances as PSNE. As expected, it represents the frequently repeated voting instances (such as the ones mentioned above). A graphical representation of the LIG is shown in Figure 2.10. We have clustered the nodes on the traditional perception that Justices Scalia, Thomas, Rehnquist, and O’Connor are “conservative;” Justices




Figure 2.9: *Pictorial representation of LIG learned from data—the non-diagonal elements represent influence factors and the diagonal elements biases. The numbering of the players (from 1 to 9) corresponds to the Justices in this order: Justices A Scalia, C Thomas, WH Rehnquist, SD O’Connor, AM Kennedy, SG Breyer, DH Souter, RB Ginsburg, and JP Stevens. The darker the color of a cell, the more negative is the corresponding number. For example, the most negative number ( $-0.2634$ ) occurs in cell (5,5) (i.e., the bias of Justice Kennedy). The most positive number ( $0.4282$ ) occurs in cell (1,2) (i.e., the influence factor from Justice Scalia to Justice Thomas) and the number closest to zero is 0.001 in cell (2,4).*

Breyer, Souter, Ginsburg, and Stevens are “liberal;” and Justice Kennedy is a “moderate.” As illustrated in Figure 2.10, negative influence factors occur only between players of two different clusters.

### Most Influential Nodes

Analysis of the PSNE of this LIG shows that there is a set of two nodes that is most influential with respect to achieving the objective of every Justice voting yes. This most influential set consists of one node from the set {Scalia, Thomas} and another one from the set {Breyer, Souter, Ginsburg, Stevens}. Furthermore, any one node from the set {Breyer, Souter, Ginsburg, Stevens} is alone most influential with respect to achieving the objective of a 5-4 vote mentioned above (i.e., the second most repeated instance in the data).



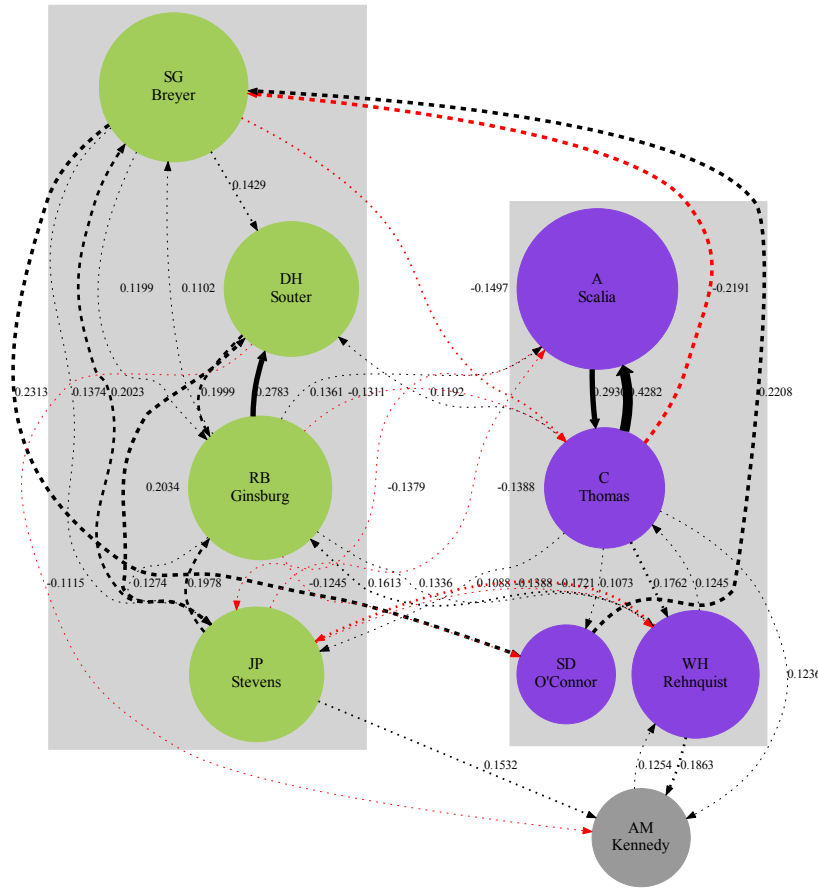


Figure 2.10: *Graphical representation of LIG learned from data. Larger node sizes indicate higher thresholds (more stubborn). Positive influence factors are drawn as black arcs and negative as red. Thicker arcs represent higher value of influence factors. While the learned LIG is a complete graph, we have only drawn approximately half of the arcs (i.e., we are not showing the “weakest” arcs in this graph).*

### 2.5.4 Illustration: Congressional Voting

We further illustrate our computational scheme in another real-world scenario where the strategic aspects of the agents’ behavior are of prime importance. We first learned the LIGs among the senators of the 101st and the 110th U.S. Congress [55]. The 101st Congress LIG consists of 100 nodes, each representing a senator, and 936 weighted arcs among these nodes. On the other hand, the 110th Congress LIG has the same number of nodes, but it is a little sparser

than the 101st one, having 762 arcs. In these LIGs, each node can play one of the two actions: 1 (yes vote) and  $-1$  (no vote). A bird's eye view of the 110th Congress LIG is shown in Figure 2.11 and a part of it is magnified and shown in Figure 2.12.

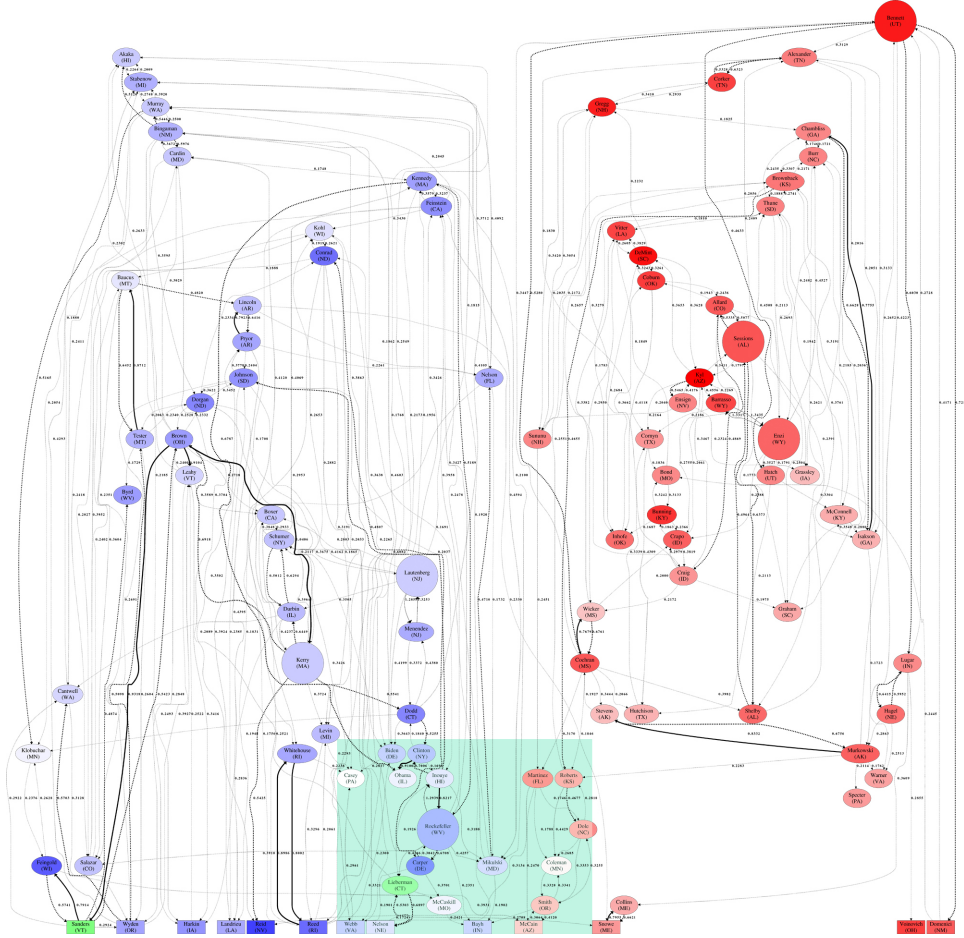


Figure 2.11: *LIG for the 110th U.S. Congress: darker color of nodes represent higher threshold (more stubborn); thicker arcs denote influence factors of higher magnitude (only half of the original arcs with the highest magnitude of influence factors are shown here); circles denote most influential senators; rectangles denote cut nodes used in the divide-and-conquer algorithm. The shaded part of it has been magnified for better visualization in Figure 2.12.*

First, we have applied the divide-and-conquer algorithm that exploits the nice separation properties of these LIGs, to find the set of all PSNE (this has been done for convenience; as discussed earlier, counting alone would have been sufficient). We have obtained a total of 143,601 PSNE for the 101st




Figure 2.12: A part of the LIG for the 110th U.S. Congress: blue nodes represent Democrat senators, red Republican, and white independent; darker color of nodes represent higher threshold (more stubborn); thicker arcs denote influence factors of higher magnitude; circled node (Senator Rockefeller) denotes one of the most influential senators; rectangles at the bottom denote cut nodes used in the divide-and-conquer algorithm.

Congress graph and 310,608 PSNE for the 110th one. Note that the number of PSNE in these games is extremely small (e.g., a fraction  $2.45 \times 10^{-25} \approx 2^{-81.76}$  for the 110th Congress) relative to the maximum possible  $2^{100}$  joint actions. Regarding the computation time, solving the 110th Congress using the divide-and-conquer approach takes about seven hours, whereas solving the same without this approach, simply relying on the backtracking search, takes about 15 hours on a modern quad-core desktop computer.

Next, we have computed the most influential senators using the approximation algorithm outlined earlier. We have obtained a solution of size five for the 101st Congress graph, which we have verified to be an optimal solution. This solution consists of Senators Rockefeller (Democrat, WV), Sarbanes (Democrat, MD), Thurmond (Republican, SC), Symms (Republican, ID), and Dole (Republican, KS). Interestingly, none of the maximum-degree nodes has been selected. Similarly, the six most influential senators of the more recent 110th Congress (January 2007–January 2009) are Kerry (Democrat, MA), Bennett

(Republican, UT), Sessions (Republican, AL), Enzi (Republican, WY), Rockefeller (Democrat, WV), and Lautenberg (Democrat, NJ).

### 2.5.5 Contrasting with Diffusion Model

Our one-shot noncooperative game-theoretic model is fundamentally different from the diffusion model. However, comparing the most influential Senators (110th Congress) obtained using our setting to that obtained using diffusion revealed some striking similarity that we cannot yet explain fully. It should first be noted that both of these analyses have been done using the same influence factors and thresholds that we obtained from learning LIGs. In particular, for the diffusion setting, at each iteration we select a node  $u$  that achieves the maximum spread of action 1, force  $u$  to adopt action 1, let all but the previously selected nodes modify their actions as best responses to  $u$ 's adoption of action 1. We repeat this until every node adopts action 1.<sup>3</sup> *Note that because of negative influence factors, cycling may occur and this procedure may never come to a stop.* However, in our case, even in the presence of negative influence factors, we did not encounter such cycling. Furthermore, it is well known that the above recipe produces a provable approximation algorithm for the cascade model with submodular spread function [67], but *this claim of approximation guarantee vanishes as soon as we have negative influence factors.*

We can visualize all possible choices of the most influential nodes that an algorithm can make as a directed acyclic graph, as shown in Figures 2.13 and 2.14.

Although Figure 2.13 looks more complicated than Figure 2.14 (due to the appearance of the same node in different source-sink paths of the dag at different levels), comparing these we find that not only a set of six nodes are most influential in both cases but also most of the nodes are common between these two distinct frameworks. More remarkably, some of these common nodes are selected at the same iteration in both frameworks. One question that we can ask is does a set of most influential nodes in the LIG setting also remain most influential in the diffusion setting? We have exhaustively tested all possible sets of the most influential nodes (Figure 2.13) and settled the answer in the negative for each set. Interestingly, if we add the “Alexander R TN” node to *any* of the most influential sets in the LIG setting, the resulting set becomes most influential in the diffusion setting. The apparent similarity in results between the two models gives rise to an intriguing question asking us to

---

<sup>3</sup>At the end, we also perform a post-processing step, where we try to remove one of the selected nodes to test if the remaining nodes are still most influential.




Figure 2.13: *Most influential nodes in our setting. This directed acyclic graph (dag) illustrates all possible options for node selection that our approximation algorithm considers. A source node represents a node selected in the first iteration and a sink node represents a node selected in the last step. Any directed path from a source to a sink represents a sequence of nodes selected in successive iterations by our algorithm. All nodes in the same level and having the same parent, are tied in an iteration of the algorithm. Also note that the same node can appear in different paths of the dag at different levels.*

connect these two mathematically and algorithmically different formulations, which is out of scope for this dissertation.

## 2.5.6 Filibuster

Beyond predicting stable behavior and identifying the most influential nodes in a network, our model can be used to study other interesting aspects of a networked population. One example is the filibuster phenomenon in the U.S. Congress, where a senator uses his or her right to hold floor for an indefinite time in an effort to delay the passing of a bill. It can be broken by the procedure of “cloture” which refers to gathering a majority of at least 60 votes among the current 100 senators. However, not every possible cloture scenario of 60 or more “yes” votes may be a *stable* outcome due to influence among the senators. The set of the ones that are indeed stable in the sense of PSNE will be called the *stable cloture set*.

An interesting general question that we can ask is whether there exists a

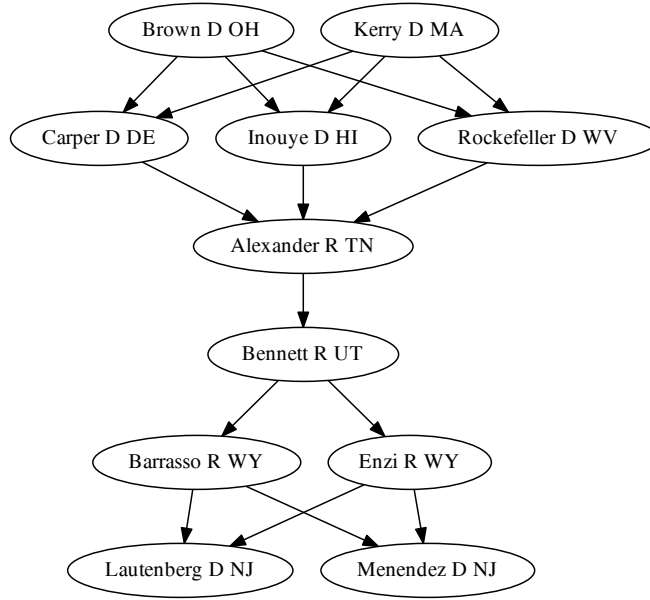


Figure 2.14: *Most influential nodes in the diffusion setting. Each directed path from a source to a sink represents a sequence of nodes that*

small coalition of senators that can break filibusters. We can also think of preventing a filibuster from the democratic or the republican perspective (i.e., favoring the respective party). Let us formally define the problem.

### Problem Formulation

Given the set  $\mathcal{S}$  of all stable outcomes (i.e., PSNE) and a subset  $\mathcal{C}$  of these stable outcomes, find a minimal set  $T$  of players such that

$$T \in \arg \max_{V \subset \{1, \dots, n\}} \{|P_{\mathcal{S}}(V)| \text{ s.t. } P_{\mathcal{S}}(V) \subseteq \mathcal{C}\},$$

where  $P_{\mathcal{S}}(V)$  is the set of PSNE-extensions of the nodes in  $V$  playing action 1, i.e.,

$$P_{\mathcal{S}}(V) = \{\mathbf{x} \text{ s.t. } \mathbf{x} \in \mathcal{S}, x_i = 1 \forall i \in V\}.$$

In this setting, we can use  $\mathcal{C}$  to denote the stable cloture set, consisting of stable outcomes that can prevent a filibuster (i.e., every PSNE in  $\mathcal{C}$  contains at least 60 “yes” votes and thus, can induce a cloture). When we consider

the notion of preventing filibuster that favors a specific party,  $\mathcal{C}$  is defined as consisting of exactly those PSNE that contain 60 or more “yes” votes (thereby representing cloture scenarios) and in addition, are supported (through “yes” votes) by the majority of the senators affiliated with that party. Other definitions are also possible as long as the stable cloture set  $\mathcal{C}$  is well-defined.

Now, we would like to select a *minimal* set of senators such that the set  $P_{\mathcal{S}}(V)$  of the PSNE-extensions of these senators’ voting “yes” is contained in  $\mathcal{C}$  (i.e., their voting “yes” can only lead to a stable cloture scenario, thereby preventing filibuster). In addition, we would also like to achieve a maximum *stable-cloture cover*, that is, we wish to achieve the maximum possible set  $P_{\mathcal{S}}(V)$  so that we are able to capture as many of the stable cloture scenarios as possible. In this formulation, we set up the objective to select a minimal, not minimum, set of senators in order to keep the formulation simple by avoiding bicriteria optimization (minimum set of senators vs. maximum stable-cloture cover). Further note that adding an extra senator to the set of selected senators can only reduce the stable-cloture cover due to additional constraints.

The above problem formulation guarantees a nonempty solution  $T$  if there exists some PSNE in  $\mathcal{C}$  that is not “dominated” by any PSNE in  $\mathcal{S} \setminus \mathcal{C}$ . Here, a PSNE  $\mathbf{x}$  dominates another PSNE  $\mathbf{y}$  if for every  $i$ ,  $y_i = 1 \implies x_i = 1$ .

## A Heuristic

We can modify the approximation algorithm for identifying the most influential nodes to design a heuristic for this problem in the following way. At each iteration, we select a node such that adding it to the set of already selected nodes minimizes the number of PSNE-extensions of the selected nodes playing 1 that are in  $\mathcal{S} \setminus \mathcal{C}$ . If there is a tie among several nodes in this step, then we can store these nodes in order to explore all solutions that this heuristic can produce. We stop when the above number of PSNE-extensions within  $\mathcal{S} \setminus \mathcal{C}$  goes to 0. We then perform a minimality test by excluding nodes from the selected set of nodes and testing whether the resulting set can be a solution. Note that although we can select the “best” solution (in terms of the coverage of  $\mathcal{C}$ ) among the ones found due to ties, this heuristic does not guarantee an approximation of the maximum coverage of  $\mathcal{C}$ .

## Experimental Results on the 110th Congress

For the 110th Congress,  $\mathcal{C}$  consists of 15,288 and 10,029 stable cloture scenarios (i.e., PSNE) with respect to the democratic and the republican parties, respectively. Overall, the total number of stable cloture scenarios is 15,595, and most of these are common in both democratic and republican cases. With

respect to the democratic party, the best solutions found by the above heuristic are Senators {Brown (D, OH), Roberts (R, KS), and Graham (R, SC)} and {Kerry (D, MA), Roberts (R, KS), and Graham (R, SC)}, both of which cover 1,500 of the 15,288 stable cloture scenarios. The optimal solutions found by a brute-force procedure are Senators {Brown (D, OH), Craig (R, ID), and Dole (R, NC)} and {Kerry (D, MA), Craig (R, ID), and Dole (R, NC)}, both covering 1,728 stable cloture scenarios. With respect to republican party, the heuristic gives these two solutions as the best, each covering 40 of 10,029 stable cloture scenarios: Senators {Brown (D, OH), Bennett (R, UT), and Gregg (R, NH)} and Senators {Kerry (D, MA), Bennett (R, UT), and Gregg (R, NH)}. The optimal solution for this case is Senators {Bennett (R, UT), Conrad (D, ND), and Sessions (R, AL)}, which covers 138 stable cloture scenarios.

### Application of Diffusion Models to this Problem

We can once again contrast our approach with that of diffusion to highlight two notable shortcomings of the latter. First, the notion of stable-cloture cover is not well-defined in the diffusion setting. The forward recursion mechanism central to diffusion models begins with a set of initial adopters (those senators selected to vote “yes” in our case) and propagates the effects of behavioral changes throughout the network until it reaches a steady state (i.e., no change occurs). However, this mechanism focuses on how the dynamics of behavioral changes evolves, not on the count of steady states that are consistent with a given set of players being among the adopters (not necessary early adopters), which is required for stable-cloture covers. In contrast, stable-cloture cover is well-defined in our approach.

Second and most important, even if we allow reversals of actions due to negative influence factors, forward recursion may produce an *unstable* outcome (i.e., not a PSNE). Although Granovetter’s original model precludes this by requiring the initial adopters to have a threshold of 0 [49], subsequent development allows forward recursion to start with a set of initial adopters whose thresholds are not necessarily 0 [67]. Next, we illustrate this point using our experimental results.

As justified above, in our experimental setting regarding diffusion models, we omit the notion of maximum stable-cloture cover and thereby forgo the measure of goodness of a solution. We only concentrate on finding a set of initial adopters that can drive the forward recursion process to *some* stable cloture scenario (i.e., a PSNE in  $\mathcal{C}$ ). Our experimental procedure is outlined below.

For  $k = 1, 2, \dots$ , do the following. For all possible sets of  $k$  senators, start forward recursion with these  $k$  senators forced to play 1 all the time and other



senators initially playing  $-1$  (but are permitted to switch between  $1$  and  $-1$  later on). When a steady state is reached, verify if there are at least 60 senators who are playing  $1$  in this state. If this is the case, then further verify if the  $k$  senators who are forced to play  $1$  are indeed playing their best response with respect to others' actions, which is the condition for the cloture scenario being stable. Stop iterating over  $k$  once stable cloture scenarios are found.

In our problem instances, which contain both positive and negative influence factors, it is very much possible that forward recursion oscillates indefinitely. However, it did not happen in our experiments. We tried all possible sets of  $k \leq 3$  initial adopters, but failed to reach any cloture scenario (stable or unstable). We then tried all possible quadruplets of initial adopters. With respect to democratic party, 1,189 different quadruplets led the forward recursion process to a cloture scenario, but nearly half of these quadruplets (536 to be exact) led to unstable outcomes. Essentially, those unstable outcomes were due to some of the initial adopters not playing their best response in voting “yes”—all other nodes were indeed playing their best response (otherwise, the process would not terminate).

Therefore, beyond just emphasizing the stability of an outcome, our approach also captures certain phenomena that cannot be captured using the traditional approach.

### 2.5.7 Influence of the “Gangs” of Senators

We can use our model to study the influence of the gangs of senators that are often assigned the task of formulating critical policies on contentious issues. We can ask questions like what outcomes could potentially be generated if a certain gang of senators vote “yes” on a bill. Although we do not have a way of systematically validate our model, answering questions like these can give anecdotal validation that our model is capturing certain strategic aspects of the system.

One such gang of senators, known as the *gang-of-six senators* received much media spotlight in 2011 while trying to reach a deal on debt-ceiling between the Democrats and the Republicans. It is a bipartisan group consisting of Senators Chambliss (R, GA), Coburn (R, OK), Crapo (R, ID), Conrad (D, ND), Durbin (D, IL), Warner (D, VA). Despite their best efforts, they were unable to succeed in reaching a deal, which led to the so-called debt-ceiling crisis toward the end of July 2011. Using our model to infer the set of stable outcomes consistent with the gang-of-six senators voting “yes,” we find that there are 11,106 stable outcomes in this set. Among these, 3,007 stable outcomes consist of less than 50 “yes” votes in each outcome and 10,000 stable outcomes less than 60 “yes” votes. So, according to our model, the gang-of-six senators are not collectively

“powerful enough” to “force” a majority outcome or to prevent a filibuster scenario. This empirical result is consistent with what happened in reality.

Another interesting happening is the formation of the *gang-of-eight* senators in 2012. This new gang of senators consisted of the old gang-of-six senators and two new members: Senators Michael Bennet (D, CO) and Mike Johanns (R, NE). They were given the task of reaching a deal to avoid the much-dreaded “fiscal cliff” situation by December 2012. Essentially, the gang-of-six was strengthened by the addition of the two new senators, which supports the above empirical observation.

## 2.6 Conclusion

In this chapter, we have studied influence and stable behavior from a new game-theoretic perspective. To that end, we introduced a rich class of games, named influence games, to capture the core strategic component of complex interactions in a network. We characterized the computational complexity of computing and counting PSNE in LIGs. We proposed practical, effective heuristics to compute PSNE in such games and demonstrated their effectiveness empirically. Besides predicting stable behavior, we gave a framework for computing the most influential nodes and its variants. We also gave a provable approximation algorithm for the most influential nodes problem.

Although our models are inspired by earlier works by sociologists, at the heart of our whole approach is abstracting the complex dynamics of interactions by the solution concept of PSNE, which allowed us to deal with richer problem instances (e.g., the ones with negative influence factors) as well as to tread into new problem settings beyond identifying the most influential nodes. We conclude this chapter by outlining several interesting lines of future work.

First, we leave several computational problems open. We have shown that counting the number of PSNE even in a star-type LIG is  $\#P$ -complete, but does there exist an FPRAS for the counting problem? The computational complexity of indiscriminant LIGs, which we conjecture to be PLS-complete, is unresolved. Also, computing mixed-strategy Nash equilibria of LIGs, even for special types such as trees, remains an open question.

Second, we can apply our models to a general setting of interventions where we study the effects of changes in node thresholds, connectivity, or influence factors, usually without the possibility of having the corresponding behavioral data. The following is an illustrative example of it in the context of the 111th U.S. Congress. After the death of Senator Ted Kennedy, who was a democratic senator from the state of Massachusetts, a republican senator named Scott Brown was elected in his place. Not only that it was Senator Brown’s first

appointment in Senate, he was also the first republican from Massachusetts to be elected to Senate for a long time. Without any behavioral data at that time, we could perform interventions in our model under various assumptions of thresholds, connectivity, and influence factors regarding Senator Brown, with the general goal of predicting stable outcomes and investigating the effects of this intervention in various settings, such as the filibuster scenario or the setting of the most influential senators.

Another example of intervention, in the context of the Framingham heart study alluded at the beginning of this chapter, is the following. Suppose that we would like to implement a policy of targeted interventions in order to reduce smoking by some margin. Using our model, we can modify the thresholds of the selected targets and predict how it could affect the overall level of smoking.

Besides interventions, we can also use our model to analyze past happenings, such as the role of the bipartisan gang-of-six senators in reaching an agreement during the U.S. debt ceiling crisis.<sup>4</sup> We know that this bipartisan coalition could not prevent the crisis, but can we shed more light on it using the stable outcomes predicted by our model? Questions like these and many others shape the long-term goal of this research.

## 2.A Brief Review of Collective Behavior and Collective Action in Sociology

In sociology, the umbrella of *collective behavior* is very broad and encompasses an incredibly rich set of models explaining various aspects of a wide range of social phenomena such as revolutions, movements, riots, strikes, disaster, panic, diffusion of innovations (e.g., fashion, adopting contraceptives, electronic gadgets, or even religion) just to name a few. In fact, the richness of just one subfield of collective behavior, termed *micro-level theories of collective behavior*, led Montgomery to comment in his book [88, p. 67], “The variety of theories focusing on the micro level is confusing, but is an indication of the complexity and variations in the process by which movements emerge or perhaps fail to emerge...” Sociologists Marx and McAdam, in their concise introductory book on collective behavior [76], contend that unlike many other fields of sociology, the field of collective behavior is not easy to define, partly because of the varied opinion of scholars, beginning from a very narrow perspective and ranging up to such a wide and all-encompassing perspective (e.g., Robert Park and Herbert Blumer’s) that there is virtually no need to have this as an individual field in sociology. Quoting from their book, “The

---

<sup>4</sup>[http://en.wikipedia.org/wiki/United\\_States\\_debt-ceiling\\_crisis](http://en.wikipedia.org/wiki/United_States_debt-ceiling_crisis)

field of collective behavior is like the elephant in Kipling’s fable of the blind persons and the elephant. Each person correctly identifies a separate part, but all fail to see the whole animal.”

Despite this, Marx and McAdam point out the traditional disposition to categorize collective behavior as a “residual field” in sociology. That is, the study of collective behavior consists of those elements of behavior (e.g., fads, fashion, crazes), organization (e.g., social movement), group (e.g., crowd), individual (e.g., psychological states such as panic), etc. that do not readily fit into well-established and commonly observed social structures. Similarly, collective behavior is defined in Goode’s textbook [45, p. 17] as the “relatively spontaneous, unstructured, extrainstitutional behavior of a fairly large number of individuals.”

### 2.A.1 Classical Treatment of Collective Behavior: Mass Hysteria

The classical treatment of collective behavior views individuals in a crowd as non-rational and transformed into being hysteric by the collective environment. The central tenet of the early work by Gustave Le Bon’s is that individuals in a crowd share a “mind of the crowd” and their psychological state and behavior in the collective setting is guided by the “psychological law of the mental unity of crowds” [71, p. 5]. In Le Bon’s account, an individual in a crowd may retain some of the ordinary characteristics he shows in isolation, but the emphasis is on the extraordinary characteristics that emerge only in a crowd due to “a sentiment of invincible power,” contagion, and most importantly, due to the individuals being prone to taking suggestions as if they were hypnotic subjects. Examples of such extraordinary characteristics of a crowd are “impulsiveness,” “incapacity to reason,” and “the absence of judgment,” to name a few [71, p. 16]. In sum, individuals in a crowd are depleted of their intellectual capacity and become uniform in their psychological state. This leads the crowd to an identical direction of collective behavior that may be heroic or criminal, depending on the type of “hypnotic suggestion” alluded above (although Le Bon gives examples of heroic crowds [71, p. 14], for the most part, he rather portrays crowds with a negative connotation).

Le Bon’s work influenced subsequent developments for around half a century. His proposition of transformation of individuals in a crowd was upheld by Park and Burgess, who put forward the concept of *circular reaction* [94], which was later refined by Herbert Blumer [19]. Circular reaction refers to a reciprocal process of social interaction that explains *how* crowd members become uniform in their behavior, something that Le Bon could not really explain. In

this process, an individual’s behavior stimulates another individual to behave alike, and when the latter individual does so, it reinforces the stimulation that the former individual acted upon. In circular reaction, individuals do not act rationally or intellectually. That is, they do not reason about the action of others; they only align themselves with the behavior of others. This is different from *interpretative interaction*, another mechanism that Blumer defined to explain routine group behavior (e.g., a group of individuals shopping in a mall) as opposed to collective behavior (e.g., social movement). In interpretative interaction, individuals react (perhaps differently) to their *interpretation* of others’ action, not the action itself. Therefore, individuals can be treated as rational beings in interpretative interaction. (For clarity of presentation, we differentiate between routine group behavior and collective behavior. To the contrary, Blumer, as well as Park in his earlier work, viewed collective behavior as encompassing a wide range of social phenomena, including routine group behavior. Within the continuum of collective behavior in Blumer’s view, the presence or the absence of rationality of individuals earmarks two distinct mechanisms named interpretative interaction and circular reaction, respectively.)

In Blumer’s account, a crowd goes through several well-defined stages before a collective behavior finally emerges. The three underlying mechanisms that facilitate transitions among these stages are circular reaction, collective excitement, and social contagion, where one can roughly think collective excitement as a more intense form of circular reaction and social contagion as even more intense [82, p. 11].<sup>5</sup>

## 2.A.2 Emergent Norm Thoery

Although Blumer’s account of collective behavior received wide-spread acceptance even beyond academia [82, p. 9], many of the underlying assumptions in it, as well as in the general mass hysteria theory, were deemed unrealistic by others. Arguably, a collective behavior is participated by individuals with *different* objectives in mind, and *changes* in their individual behavior can be observed throughout the process of a collective behavior. The reader is referred to [85, p. 26–27] for a beautiful example in the context of the 1967 anti-war demonstration in Washington, DC, which was participated by nearly 250,000

---

<sup>5</sup>The mechanism of “social contagion” as defined by Blumer or the “social contagion theory” [74, p. 11] in general is not to be confused with the term “social contagion” that computer scientists use [116]. Although both have their roots in epidemics, in the former case, individuals are transformed into being more suggestible that facilitates “rapid, unwitting, and non-rational dissemination” of behavior [19], whereas in the latter case, individuals act rationally.

people. While many of the participants might have been there to genuinely voice their opinion against the war, some might have been looking for “excitement, drug, or sex.” Yet again, individuals playing different roles, such as protest leaders, street vendors, and the police, behaved differently. Therefore, the assumption of complete uniformity behavior in the classical mass hysteria treatment is very much a stretch. Furthermore, the assumption of hysteric crowd in the classical approach has also been called into question. One notable critique of the mass hysteria theory comes from Ralph Turner and Lewis Killian [115]. They view individuals in a crowd as behaving under normative constraints and showing “differential expression.” However, when the crowd is faced with an extraordinary situation that is not adequately guided by the established norms of the society, a new norm emerges. They call this the *emergent norm* and contend that it is the emergent norm that gives the “illusion of unanimity.”

### 2.A.3 Collective Action

The goal-oriented nature of collective behavior was further highlighted by sociologists studying social movements during the 1970s and 80s. In order to distinguish their approach from the traditional approach to collective behavior dominated by the assumption of irrational and aimless nature of crowds, they used the term *collective action* signifying “people acting together in pursuit of common interests” [112]. Strikingly, based on a series of systematic observations, Clark McPhail’s contends that the goal-oriented nature of crowds is not limited to social movements and revolutions alone, but is a feature of various other types of crowds. In his book *the Myth of the Madding Crowd*, he uses two decades of empirical observations pertaining to a multitude of crowd settings to formulate a theory of collective behavior that is recognized as a significant paradigm shift [82, ch. 5, 6]. To distance himself from the term “crowds,” which has already gained several meanings depending on whose theory is being considered, he gives his formulation in the setting of “gatherings.” But first, he places a justifiably strong emphasis on the definition of collective behavior. His “working definition of collective behavior” can be stated briefly as the study of “two or more persons engaged in one or more behaviors (e.g., locomotion, ...) judged common or concerted on one or more dimensions (e.g., direction, velocity, ...)”<sup>6</sup> [82, p. 159].

The broad nature of McPhail’s definition, although based on extensive em-

---

<sup>6</sup>Note that it is the “behavior,” not a specific action, that needs to be common or concerted. For example, when a group of people are chatting together, their behavior is concerted, even though they are not speaking identical words.

pirical evidence, was not readily accepted as a definition of collective behavior. Even modern textbooks on collective behavior try to conserve the classical appeal of collective behavior. For example, in reference to McPhail's definition, Goode writes in his textbook, "In the view of most observers, myself included, many gatherings are *not* sites of collective behavior (most casual and conventional crowds, for example), and much collective behavior does not take place in gatherings of any size (the behavior of most masses and publics, for example)" [45, p. 17]. Ironically, this is the very viewpoint that McPhail seeks to portray as a myth. Perhaps to further distance himself from the traditional viewpoint, McPhail later began to use the term *collective action* instead of collective behavior (for example, in a recent encyclopedia article, McPhail refers to the above mentioned definition as that of collective action [83]). According to David Miller, the modern view on the distinction between collective action and collective behavior is beyond simply terminological. Collective action is given the status of a "new" theory in sociology, while collective behavior is marked as "old," but not unimportant. [85, p. 14–15]. Miller also points out that sociology textbooks are likely to talk about collective behavior only whereas recent journal articles on collective action.

McPhail's approach to collective action is known as *social behavioral interactionist (SBI)* approach. As much as it agrees with the emergent norm theory in terms of the diversity of individual objectives in a collective setting, it does not agree with the concept of an emergent norm suppressing this diversity. The SBI approach studies gatherings in three phases of its life cycle: the assembling process, collective action within the assembled gathering, and the dispersal process [82, p. 153]. Although each of these three phases is rich and interdependent, the goal is to manage the complexity of collective action as a whole by focusing on the recognizable parts of it. Interestingly, the underlying mechanism to explain collective action is drawn from the *perceptual control theory* [82, Ch. 6]. McPhail adapts this theory to formulate his *sociocybernetics theory* of collective action. In brief, an individual receives sensory inputs, compares the input signal to its desired signal,<sup>7</sup> and adjusts its behavior in response to the discrepancy. Behavior of individuals affects the "environment," which in turn affects the input signal, thereby completing a loop. An important aspect of this theory is that various external factors (or "disturbances") may drive an individual to make different behavioral adjustments at different points in time even if the discrepancy between the input signal and the desired signal remains the same.

---

<sup>7</sup>In contrast to engineering control systems theory, the desired signal is not external, but set by individuals themselves (which is also highlighted by the term *cybernetics*).

## Additional Notes on Schelling’s Models

Schelling’s models assume that individuals behave in a discriminatory way. For example, individuals are aware, consciously or unconsciously, about the types of other individuals in their neighborhood and behave (i.e., stay in the neighborhood or leave) according to their preference. This is different from organized processes (e.g., separation of on-campus residence between graduate and undergraduate students due to a university’s housing policy) or economic reasons (e.g., segregation between the poor and the rich in many contexts) [108, 109]. An example of a segregation due to individual choice, or “individually motivated segregation” as Schelling puts it [109, p. 145], is the residential segregation by color in the U.S. In fact, Schelling’s models and their analyses expressly focus on this case. Yet, Schelling’s theory can be applied to many other scenarios as well, since it explains, at an abstract level, how *collective* outcomes are shaped from *individual* choice. It should be mentioned here that connecting individual actions to collective outcomes is a mainstream theme of research in collective action.

Schelling introduces two basic models to study the dynamics of segregation among individuals of two types [109]. The first model is named the *spatial proximity model* where individuals are initially positioned in a spatial configuration (such as a line or a stylized two-dimensional area) and individuals of the same type share a common “level of tolerance,” which quantifies the upper limit on the percentage of an individual’s opposite type in his local neighborhood that he can put up with. Here, an individual’s local neighborhood is defined with reference to the individual’s position in the specified spatial configuration. The dynamics of segregation is studied in this model using a rule of movement for the “unhappy” individuals. For example, an individual whose level of tolerance has been exceeded, moves to the closest location where the tolerance constraint can be satisfied. For equal number of individuals of each type and a fixed local neighborhood size, Schelling first studies how clusters evolve from the initial configuration of a random placement of individuals on a straight line. He then generalizes the experimentation by varying various model parameters such as neighborhood size, level of tolerance, ratio of individuals of the two types, etc. Notable findings are that decreasing the local neighborhood size leads to a decrease in the average cluster size and that for unequal number of individuals of the two types, decreasing the relative size of the minority leads to an increase in the average minority cluster size.

Schelling extends this experimentation to a different setting of a two-dimensional checkerboard. The individuals are randomly distributed on the squares of the checkerboard, leaving some of the squares unoccupied. An individual’s local neighborhood is defined by the squares around it and an unhappy



individual moves to the “closest” unoccupied square (leaving its original square unoccupied) that can satisfy its tolerance constraint. In addition to studying clustering properties by varying various model parameters, two new classes of individual preferences have been studied—congregationist preferences and integrationist preferences. In a congregationist preference, an individual only wants to have at least a certain percentage of neighbors of its own type and does not care about the presence of individuals of the opposite type in its neighborhood. Experimentation shows that even when each individual is happy being a minority in its neighborhood (e.g., having three neighbors of its own type out of eight), the dynamics of segregation leads to a configuration as if the individuals wished to be majority in their neighborhoods. In an integrationist preference, individuals have both an upper and a lower limit on the level of tolerance. Dynamics is much more complex in this case and leads to clusters of unoccupied squares.

Schelling’s second model, named the *bounded-neighborhood model*, is concerned with one global neighborhood. An individual enters it if it satisfies its level of tolerance constraint and leaves it otherwise. The level of tolerance is no longer fixed for each type and the distribution of tolerances among individuals of each type is given. The emphasis on the stability of equilibria when the distribution of tolerances and the population ratio of the two types are varied. For example, under a certain linear distribution of tolerances and a population ratio of 2 : 1, there exist only two stable equilibria, each consisting of individuals of one type only, whereas a mixture of individuals of both types can arise as a stable equilibrium under a different setting. This model has been adapted for the study of the *tipping* phenomenon with one notable constraint, that is, the capacity of the neighborhood is fixed. An example of a tipping phenomenon is when a neighborhood consisting of only one type of individuals is later inhabited by some individuals of the opposite type and as a result, the entire population of the original type evacuates the neighborhood. An important finding is that in the cases studied, the modal level of tolerance does not correspond to a tipping point.

## 2.B Experimental Results in Tabular Form

Table 2.2 shows experimental data of PSNE computation on uniform random directed graphs. Of particular interest is the result that the number of PSNE usually increases when the flip probability  $p$  is increased, i.e., when the number of arcs with negative influence factors is increased.

Table 2.3 illustrates the experimental result that the logarithmic-factor approximation algorithm for identifying the most influential individuals performs

Table 2.2: *PSNE computation on uniform random directed graphs. Offsets of 95% confidence intervals are shown in parenthesis.*

$p$	# of equilibria Avg (95% CI)	# of node visits/equilibrium Avg (95% CI)	Avg CPU time (sec) for computing all equilibria
0.00	2.18 (0.16)	35379.73 (3349.77)	1.81
0.125	3.72 (0.50)	22756.15 (2673.56)	1.57
0.25	13.00 (1.92)	9796.30 (1748.76)	1.9
0.375	19.42 (2.88)	7380.97 (1870.11)	1.95
0.50	14.40 (2.17)	9826.61 (1696.52)	2.04
0.625	19.78 (3.40)	8167.60 (1450.48)	2.07
0.75	28.76 (4.34)	6335.18 (1963.11)	2.21
0.875	67.14 (9.52)	4064.06 (1539.33)	3.27
1.00	194.96 (28.47)	1879.45 (235.90)	5.23

very well in practice.

Finally, Table 2.4 shows the influence factors and thresholds of the LIG among the U.S. Supreme Court Justices, which are learned using the U.S. Supreme Court dataset.

Table 2.3: *Computation of the most influential nodes: Comparison between the solution given by the Approximate-UniqueHyperedge algorithm and the optimal one in random directed graphs. Offsets of the 95% confidence intervals are shown in parenthesis.*

$p$	Approx soln size Avg (95% CI)	Opt soln size Avg (95% CI)	Approx = Opt % of times	Approx $\leq$ Opt+1 % of times	Approx $\leq$ Opt+2 % of times
0.00	1.06 (0.05)	1.06 (0.05)	100	100	100
0.125	1.50 (0.11)	1.48 (0.10)	98	100	100
0.25	2.57 (0.18)	2.14 (0.11)	62	96	99
0.375	2.75 (0.18)	2.35 (0.13)	62	98	100
0.50	2.47 (0.16)	2.09 (0.10)	66	96	100
0.625	2.82 (0.18)	2.31 (0.13)	52	97	100
0.75	3.02 (0.19)	2.48 (0.14)	51	95	100
0.875	3.83 (0.20)	3.04 (0.15)	36	86	99
1.00	4.72 (0.23)	3.95 (0.18)	37	87	99

	Scalia	Thomas	Rehnquist	O'Connor	Kennedy	Breyer	Souter	Ginsburg	Stevens
Scalia	0.0120	0.4282	0.0317	0.0717	0.0721	0.0772	-0.0321	0.1362	-0.1388
Thomas	0.2930	-0.1020	0.1245	-0.0010	0.0183	-0.1497	0.0839	-0.1311	0.0965
Rehnquist	0.0671	0.1762	-0.0834	0.0973	0.1254	0.0921	-0.0861	0.1336	-0.1388
O'Connor	0.0580	0.1073	0.1045	-0.2522	-0.0537	0.2313	0.0325	-0.1245	-0.0359
Kennedy	0.0666	0.1236	0.1863	-0.0255	-0.2634	0.0548	-0.1115	-0.0149	0.1532
Breyer	0.1009	-0.2191	0.0570	0.2208	-0.0061	-0.0209	0.0627	0.1102	0.2023
Souter	0.0368	0.1192	-0.0476	0.0762	-0.0338	0.1429	-0.0619	0.2783	0.2034
Ginsburg	0.0779	-0.1000	0.1613	-0.0962	-0.0089	0.1199	0.1999	-0.0381	0.1978
Stevens	-0.1379	0.1088	-0.1721	-0.0568	0.1053	0.1374	0.0932	0.1274	-0.0611

Table 2.4: *LIG learned from data. Each non-diagonal element in each row represents the influence factor the row player receives from the column players (for example, Justice Scalia's influence factor on Justice Thomas is 0.2930 and that in the opposite direction is 0.4282). The diagonal elements represent thresholds of the corresponding row players.*

## Chapter 3

# Causal Strategic Inference in Economic Networks

In the previous chapter, we applied causal strategic inference to a social network setting in order to answer various interesting questions, such as identifying the most influential nodes. The key idea was to model influence among the nodes in a game-theoretic way and then to use the solution concept of Nash equilibrium, which we used to model stable outcomes, to answer those questions.

In this chapter, we apply the same framework of causal strategic inference to a different setting, namely networked microfinance economies. We model a networked microfinance economy as a two-sided abstract economy, which is classically known as an extension of games. As in the previous chapter, we characterize stable outcomes from this economy as equilibrium points and answer various interesting questions with respect to the equilibrium points. The motivation for applying causal strategic inference to such an economic network lies in the inability of performing trial-and-error experiments to find the answers to many policy-level questions.

### 3.1 Microfinance Systems

Although the history of microfinance systems takes us back to as early as the 18th century, the foundation of the modern microfinance movement was laid in the 1970s by Muhammad Yunus, a then-young Economics professor in Bangladesh. It was a time when the newborn nation was struggling to recover from a devastating war and an ensuing famine. A blessing in disguise may it be called, it led Yunus to design a small-scale experimentation on micro-lending as a tool for poverty alleviation. The feedback from that experimentation gave

Yunus and his students the insight that micro-lending mechanism, with its social and humanitarian goals, could successfully intervene in the informal credit market that was predominated by opportunistic moneylenders. Although far from experiencing a smooth ride, the microfinance movement has nevertheless been a great success story ever since, especially considering the fact that it began with just a small, out-of-pocket investment on 42 clients and boasts a staggering 100 million poor clients worldwide at present [121]. Yunus and his organization Grameen Bank have recently been honored with the Nobel peace prize “for their efforts to create economic and social development from below.”<sup>1</sup>

A puzzling element in the success of microfinance programs is that while commercial banks dealing with well-off customers struggle to recover loans, microfinance institutions (MFI) operate without taking any collateral and yet experience very low default rates! The central mechanism that MFIs use to mitigate risks is known as the *group lending with joint-liability contract*. Roughly speaking, loans are given to groups of clients, and if a person fails to repay her loan, then either her partners repay it on her behalf or the whole group gets excluded from the program. Besides risk-mitigation, this mechanism also helps lower MFI’s cost of monitoring clients’ projects. This and other interesting aspects of microfinance systems, such as efficiency and distribution of intervening informal credit markets, failure of pro-poor commercial banks, gender issues, subsidies, etc., have been beautifully delineated by de Aghion and Morduch in their book [29].

Individual components of microfinance systems have been subjected to various theoretical and empirical studies. Although these studies do not directly relate to our study of microfinance markets, these, nevertheless, help us grasp the essence of the operating principles of the overall system. For example, the *adverse selection problem* has been analyzed by Ghatak and Guinnane [44]. This problem arises when loans that are targeted toward safe clients, end up in risky clients instead. The failure of many loan programs by state-owned banks in developing countries has been attributed to this problem [29]. It has been shown that in the absence of collateral, if the clients are allowed to choose their own partners, then an *assortative matching* takes place that mitigates the adverse selection problem [44]. An empirical validation of this result can be found in a World Bank study conducted in Peru [121], where the authors designed 11 different games of two basic types—repeated one-shot games and dynamic games—which were played for seven months by actual human subjects. Experimental results show that if the players are allowed to choose their own partners (similar to the real-world scenario), then like-minded players are

---

<sup>1</sup>[http://www.nobelprize.org/nobel\\_prizes/peace/laureates/2006/](http://www.nobelprize.org/nobel_prizes/peace/laureates/2006/)

paired together. This assortative matching induces risk aversion.

Another potential problem for a microfinance program is *moral hazard*, which is the problem of enforcing sincere efforts from the borrowers in utilizing the loan in their projects, as well as eliciting truthfulness about the outcome of the projects (in case of a failed project, borrowers do not have any obligation to repay their loans). The theoretical model of Ghatak and Guinnane shows that if a non-monetary social sanction is effective, then group lending with a joint-liability contract improves repayment rates [44]. The authors also show that group lending mechanism can reduce the bank's cost of auditing by letting peers in a group audit each other's case and by limiting the bank's auditing to only the case where the whole group claims to be unable to pay.

In our model, we will assume that assortative matching and joint-liability contracts would mitigate the risks of adverse selection and moral hazard. We will further assume that due to these mechanisms, there would be no default on loans. This assumption of complete repayment of loans may seem to be very much idealistic. However, besides the theoretical results cited above, practical evidence also suggests very high repayment rates. For example, Grameen Bank's loan recovery rate is 99.46% [92].

Tangentially related to our study of microfinance markets is the empirical study of whether competition among the MFIs has a substantial effect on the interest rates. As reported by Porteous, competition among the MFIs has had little effect on the interest rates in countries like Bangladesh and Uganda, whereas it has driven the interest rates down in Bolivia [101]. Furthermore, it has been reported that many MFIs do not use interest rates as a "weapon for competition," as much as they use other features such as larger loan size, shorter waiting period, and flexibility in loan repayment [120]. We will draw on these research results to formulate the diversification term in the village-side objective function.

The goal of our work is different from several other studies on microfinance economies. For example, Banerjee et al. study the how microfinance participation spreads through social networks [13]. They show that the level of microfinance participation goes higher if the initial adopters have higher eigenvector centrality. They also study the roles of microfinance participants and non-participants in information dissemination. In contrast, our goal is to model a microfinance market and use microfinance participation data to estimate the model in order to facilitate policy experiments. Our work is also different from the World Bank study mentioned above [121]. In that study, the goal is to understand the effects of communication on group formation in various types of games. In contrast, we model customers at a more aggregate level of villages. Our goal is also very much different from theirs.

## Brief Overview of Our Model

We will now give a high-level overview of our model, postponing the formal definition and justification for the next section. Ours is a model of a two-sided networked economy, consisting of microfinance institutions (MFIs) and villages. Each MFI has branches in a subset of villages and has the same interest rate in all of its branches (which is the case in practice). Each village can only interact with the MFIs having branches there. The MFIs are driven by the humanitarian goal of clearing their supply of loans, i.e., they want the interest rates to be such that supply equals demand. The objective of each village is to borrow the *maximum amount of diversified loans* (to be formally defined in Section 3.2) such that they would be able to repay the loans with interest. The villages invest their loans in productive projects and they apply the revenues from those projects toward repayment of their loans. An equilibrium point is specified by interest rates and loan allocations such that the objectives of the MFIs and the villages are achieved.

## Connection to Existing Models of Networked Economies

The model we are going to introduce is essentially one of networked economy. In recent times, there has been an intense, inter-disciplinary research effort in the area of networked economies, primarily undertaken by the economics and the computer science communities and in many instances jointly by researchers from these two communities. Subjects of investigation have ranged from generalizing abstract economies in a graphical setting [62], modeling networked markets, such as labor markets and trades (see, for example, Chapter 10 of [58]), designing mechanisms with desirable properties for such markets [5], analyzing how properties such as competition [18] and price variation [63] are influenced by the underlying network structure, to the most fundamental algorithmic question of computing an equilibrium point in such settings [62, 107]. Although we postpone a formal description of our model, we will now place our model in the context of the existing ones at a very high level.

Let us begin with the Fisher model [37], which consists of a set of buyers and a set of divisible goods sold by one central seller (i.e., a fully connected network). The buyers come to the market with some initial endowments of money, and each has a utility function over bundles of goods. Given the prices of the goods, their objective is to use their endowment to purchase a bundle of goods that maximizes their utility. An equilibrium point consists of the unit prices and the allocations of goods such that each buyer fulfills his objective and in addition, there is no excess demand or excess supply (i.e.,



the market clears). A *graphical* Fisher model with one good [63] consists of a set of buyers and *a set of sellers*. All the sellers sell the same good, but the important aspect of this model is that each buyer has access to a subset of the sellers, not necessarily all the sellers. An equilibrium point in this graphical setting is defined similar to the original one. An important distinction between our model and graphical Fisher model is that our model allows buyers (i.e., villages) to invest the goods (i.e., loans) in productive projects, thereby generating revenue that can be used to pay for the goods (i.e., repay the loans). In other words, the crucial modeling parameter of “endowment” is no longer a constant in our case. Furthermore, in our model, the villages have a very different objective function than the one in a Fisher model [63, 118]. There is, however, an interesting connection between our model and that of Fisher through an Eisenberg-Gale convex program formulation, which we will show in Section 3.2.

Arrow and Debreu gave a very generalized model of economy in their seminal work on competitive economies [3]. In fact, a Fisher economy is a special case of an Arrow-Debreu economy (see, for example, [118]). Arrow and Debreu’s proof of the existence of an equilibrium point uses Debreu’s concept of abstract economy [30], which interestingly generalizes Nash’s non-cooperative games [91] in the following way. In an abstract economy, not only a player’s payoff but also her domain of actions are affected by another player’s choice of action. Putting our model in the context of an abstract economy, a village’s set of possible demands for loan from an MFI depends on the MFI’s interest rate. For example, in the simplest setting of one MFI operating in one village, the village cannot ask for unlimited amount of loan from the MFI if the MFI’s interest rate is above a certain bound. However, for the same reasons cited in the previous paragraph (i.e., variable endowment), the classical Arrow-Debreu model or more specifically, a recently developed graphical extension to the Arrow-Debreu model [62], does not capture our setting.

Our work is different from various other works on networked economies [5, 18, 35, 42, 68, 107] from the perspectives of modeling, problem specification, and application. For example, Kranton and Minehart model buyer-seller exchange economies as networks with an emphasis on the emergence of links in such networks [68]. They show that although buyers and sellers are modeled as self-interested non-cooperative agents, “efficient” network structures are necessarily equilibrium outcomes and that for a restricted case, equilibrium outcomes are necessarily efficient. In a related work of significant implications, Even-Dar *et al.* completely characterize the set of all buyer-seller network structures that are equilibrium outcomes in their model of exchange economies [35]. In contrast to these works, we do not study network forma-

tion here, i.e., we treat the spatial structure of the branch-banking MFIs as exogenous.

### **Our Contribution**

The significance of our work lies in mathematically modeling real-world microfinance markets and then complementing it with an AI approach to learning the parameters of the model and designing algorithms for it. We are unaware of any work in the literature that formally models microfinance markets as a networked economy, let alone studying how it can be used in conjunction with historical data to make interventions and policy decisions. Following are our main results. (1) We show that a special case of our problem is equivalent to an Eisenberg-Gale convex program, which leads us to the proof of the existence of an equilibrium point as well as the uniqueness of the equilibrium interest rates. (2) We give a primal-dual based constructive proof of the existence of an equilibrium point in the general setting. A cornerstone of our proof is showing that the strategic complementarity property [23] is inherent in our model. (3) We use real-world microfinance data from Bangladesh and Bolivia to learn respective models. Our parameter estimation procedure takes into account the important practical problem of equilibrium selection. (4) We demonstrate how our model can be applied to formulating a variety of important policy decisions. For example, what are the effects of an interest rate ceiling? How does an intervention by introducing new MFIs affect the market? What would happen if some of the government-owned MFIs, which are known for their inefficient operation, are removed from the market? Or, what would happen if some MFI decides to close several of its branches?

## **3.2 A Model of Microfinance Markets**

We model microfinance markets as a two-sided market consisting of MFIs and villages. Each MFI has a certain number of branches, each branch being located in a distinct village and dealing with borrowers within that village only. Similarly, each village can only interact with the MFIs present there. Each MFI has a fixed amount of loan and wants to disburse all of it. On the other hand, the villages want to maximize an objective function that involves the total amount of loan it receives and how diversified that loan portfolio is, subject to the constraint that after investing the loan, it would be able to repay it with its accrued interest.

## Notation

There are  $n$  MFIs and  $m$  villages.  $V_i$  is the set of villages where MFI  $i$  operates and  $B_j$  is the set of MFIs that operate in village  $j$ .  $T_i$  is the finite total amount of loan available to MFI  $i$  to be disbursed.  $g_j(l) := d_j + e_j l$  is the revenue generation function of village  $j$  (parameterized by the loan amount  $l$ ), where the initial endowment  $d_j > 0$  (i.e., each village has other sources of income [29, Ch. 1.3]) and the rate of revenue generation  $e_j \geq 1$  are constants.  $r_i$  is the flat interest rate at which MFI  $i$  gives loan and  $x_{j,i}$  is the amount of loan borrowed by village  $j$  from MFI  $i$ . Finally, the villages have a *diversification parameter*  $\lambda \geq 0$  that quantifies how much they want their loan portfolios to be diversified.<sup>2</sup> The problem statement is given below.

### 3.2.1 Problem Formulation

Following are the inputs to the problem. First, for each MFI  $i$ ,  $1 \leq i \leq n$ , we are given the total amount of money  $T_i$  that the MFI has and the set  $V_i$  of villages that the MFI has branches. Second, for each village  $j$ ,  $1 \leq j \leq m$ , we are given the parameters  $d_j > 0$  and  $e_j > 1$  of the village's revenue generation function<sup>3</sup> and the set  $B_j$  of the MFIs that operate in that village.

Each MFI  $i$  wishes to set its interest rate  $r_i$  in a way that it is able to lend all of its available money  $T_i$ , subject to the constraints that the MFI cannot lend more money than what it has and that the interest rates must be non-negative. When all the MFIs are able to do so, it corresponds to the well-known notion of *market clearing* in economics. Following is the optimization problem for MFI  $i$ , which we will refer to as the *MFI-side optimization problem*. (Note that this problem is basically a constraint satisfaction problem due to the constant objective function.)

---

<sup>2</sup>For simplicity, we assume that all the villages have the same diversification parameters.

<sup>3</sup>When we apply our model to real-world settings, we will see that in contrast to the other inputs,  $d_j$  and  $e_j$  are not explicitly mentioned in the data and therefore, need to be learned from the data. The machine learning scheme for that will be presented in Section 3.3.2.

$$\begin{aligned}
& \max_{r_i} && 1 \\
\text{subject to} & && r_i \left( T_i - \sum_{j \in V_i} x_{j,i} \right) = 0 \\
& && \sum_{j \in V_i} x_{j,i} \leq T_i \\
& && r_i \geq 0
\end{aligned} \tag{P_M}$$

On the other hand, each village  $j$  wishes to maximize the amount of diversified loans it can obtain, subject to the constraint that it is able to repay that loan. We will call the following the *village-side optimization problem*.

$$\begin{aligned}
& \max_{\mathbf{x}_j = (x_{j,i})_{i \in B_j}} && \sum_{i \in B_j} x_{j,i} + \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}} \\
\text{subject to} & && \sum_{i \in B_j} x_{j,i} (1 + r_i - e_j) \leq d_j \\
& && \mathbf{x}_j \geq 0
\end{aligned} \tag{P_V}$$

We will call the second term in the above objective function the *diversification term*. Note that although this term bears a similarity with the well-known entropic term, it can be mathematically different, since  $x_{j,i}$ 's can be larger than 1. Although the diversification parameter  $\lambda$  can be thought of as exogenous, we will assume that  $\lambda$  is small enough so that the first term in the objective function dominates the second (details will be given in Section 3.2.4). We will call the first constraint of  $(P_V)$  the *budget constraint*. It simply expresses that each village  $j$  is able to repay the loan with accrued interest, i.e.,  $\sum_{i \in B_j} x_{j,i} (1 + r_i) \leq g_j \left( \sum_{i \in B_j} x_{j,i} \right)$ . Here, the assumption is that the revenue generation function is linear.

For the above two-sided market, we will apply the solution concept of an *equilibrium point*, which is defined by an interest rate  $r_i^*$  for each MFI  $i$  and a vector  $\mathbf{x}_j^* = (x_{j,i}^*)_{i \in B_j}$  of loan allocations for each village  $j$  such that the following two conditions hold. First, given the allocations  $\mathbf{x}$ , each MFI  $i$  is optimizing  $(P_M)$ . Second, given the interest rates  $\mathbf{r}$ , each village  $j$  is optimizing  $(P_V)$ .

The parameter  $T_i$  for all MFI  $i$  and the diversification parameter  $\lambda$  are assumed to be exogenous inputs to this model. The revenue generation func-

tions  $g_j$  of the villages  $j$  are to be estimated from the data. The interest rates  $r_i$  of MFIs  $i$  and the allocations  $x_{j,i}$  for any village  $j$  and MFI  $i \in B_j$  are to be defined by an equilibrium point.

An important observation here is that the MFI-side optimization problem ( $P_M$ ) does not have any *direct* control over the allocations  $\mathbf{x}$  (in the sense that ( $P_M$ ) treats  $\mathbf{x}$  as exogenous). Similarly, the village-side optimization problem ( $P_V$ ) does not have any *direct* control over the interest rates  $\mathbf{r}$ . Of course, the interest rates  $\mathbf{r}$  and the allocations  $\mathbf{x}$  influence each other indirectly. In this setting, an equilibrium point is specified by  $(\mathbf{r}^*, \mathbf{x}^*)$  such that both ( $P_M$ ) and ( $P_V$ ) are optimized *simultaneously* with respect to  $(\mathbf{r}^*, \mathbf{x}^*)$ . This notion of an equilibrium point is rooted in classical economics literature on markets [3, 30].

## Justification of Modeling Aspects

Various aspects of our modeling choice have been inspired by the book of de Aghion and Morduch [29] and several empirical studies on microfinance systems [89, 101, 120]. We list some of these modeling aspects below.

**Objective of MFIs.** At first, it may seem unusual that although MFIs are banks, we are not modeling them as profit-maximizing agents. The perception that MFIs are making profits while serving the poor has been described as a “myth” in Chapter 1 of de Aghion and Morduch’s book [29]. In fact, the book devotes a whole chapter to bust this myth and establish that MFIs very much depend on subsidies for sustainability (see Chapter 9 of [29] and [89]). Therefore, our modeling of MFIs as not-for-profit organizations is aligned with their humanitarian goals as well as empirical evidence.

**Objective of Villages.** Typical customers of MFIs are low-income people engaged in small projects such as rice husking, weaving, crafting, livestock raising, agricultural cultivation, etc. and a large majority of these customers are women working at home (e.g., Grameen Bank, a leading MFI, has a 95% female customer base) [29]. Clearly, there is a distinction between customers borrowing from an MFI and those borrowing from commercial banks, since many of the latter are profit-oriented corporations. Therefore, in line with the empirical evidence [29], we have modeled the village side as non-corporate agents wishing to obtain loans to invest in small household projects, not as revenue-maximizing or profit-maximizing agents.

**Diversification of Loan Portfolios.** A valid question is why the village side does not simply want to maximize the total amount of loan. Clearly, a village can do so by simply seeking loan only from the MFIs offering the lowest interest rate. However, empirical studies suggest that interest rates are not the only determinant for the demands of the villages [101, 120]. In fact, many of the NGO MFIs in Bangladesh, which operate at relatively high

interest rates, attract customers by offering various facilities, such as large loan sizes, shorter waiting periods, and flexible repayment schemes [120]. We have added the diversification term in the village objective function to reflect this. Furthermore, this formulation being in line with the quantal response approach [80], a similar interpretation can be ascribed to it. For example, mathematical psychology literature suggests that human subjects are more likely to respond according to such an approach [75].

**Complete repayment of loans.** A hallmark of microfinance systems worldwide is very high repayment rates. For example, loan recovery rates of Grameen Bank is 99.46% and that of PKSF is 99.51% [92]. As explained earlier, this striking phenomenon is typically attributed to joint liability contracts. In line with such empirical evidence, we assume that the village side completely repays its loan.

### 3.2.2 Special Case: No Diversification of Loan Portfolios

It will be useful to first study the case of non-diversified loan portfolios, i.e.,  $\lambda = 0$ . In this case, the villages simply wish to maximize the amount of loan that they can borrow. Several properties of an equilibrium point can be derived for this special case. We will later exploit these to establish one of our theoretical results. The proofs in this section are by contradiction.

**Property 3.2.1.** *At any equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$ , every MFI  $i$ 's supply must match the demand for its loan, i.e.,  $\sum_{j \in V_i} x_{j,i}^* = T_i$ . Furthermore, every village  $j$  borrows only from those MFIs  $i \in B_j$  that offer the lowest interest rate. That is,  $\sum_{i \in B_j, r_i^* = r_{m_j}^*} x_{j,i}^* (1 + r_i^* - e_j) = d_j$  for  $m_j \in \operatorname{argmin}_{i \in B_j} r_i^*$ , and  $x_{j,k}^* = 0$  for any MFI  $k \notin m_j$ .*

*Proof.* Suppose that there is an MFI  $i$  such that at an equilibrium point we have  $\sum_{j \in V_i} x_{j,i}^* < T_i$ . Clearly, in this situation, MFI  $i$ 's constraint of  $r_i^* \left( T_i - \sum_{j \in V_i} x_{j,i}^* \right) = 0$  can only be satisfied if  $r_i^* = 0$ . However, if  $r_i^* = 0$  then for any village  $j \in V_i$ , the optimal demand  $x_{j,i}^*$  will be unbounded. This happens because each village  $j$  wants to maximize  $\sum_{i \in B_j} x_{j,i}$ , and with  $r_i^* = 0$  and  $e_j \geq 1$ , the term  $(1 + r_i^* - e_j)$  in the first constraint of  $(P_V)$  becomes  $\leq 0$ . Since  $d_j > 0$ , that constraint is satisfied for  $x_{j,i}^* = +\infty$ . But  $x_{j,i}^* = +\infty$  contradicts the constraint  $\sum_{j \in V_i} x_{j,i}^* \leq T_i$  in the MFI side  $(P_M)$ . Therefore, for every MFI  $i$ ,  $\sum_{j \in V_i} x_{j,i}^* = T_i$  must hold at any equilibrium point.<sup>4</sup>

---

<sup>4</sup>It is important to note that in the above argument, the village side has been allowed

For the second claim, in order to maximize its objective function  $\sum_{i \in B_j} x_{j,i}^*$ , every village  $j$  will be interested to borrow only from those MFIs that have the minimum interest rate  $r_{m_j}^*$ , where  $m_j \in \operatorname{argmin}_{i \in B_j} r_i^*$ . Furthermore, at any equilibrium point, each village's budget constraint, i.e., the first constraint in  $(P_V)$ , must hold with equality. Otherwise, suppose that the following strict inequality holds for some village  $j$  at an equilibrium point:  $\sum_{i \in B_j, r_i^* = r_{m_j}^*} x_{j,i}^* (1 + r_i^* - e_j) < d_j$ . Since this is a strict inequality, village  $j$  can still increase its objective function  $\sum_{i \in B_j} x_{j,i}^*$ . Therefore, village  $j$  is not maximizing its objective function, contradicting our assumption that this is an equilibrium point.  $\square$

We next present a lower bound on interest rates at an equilibrium point, which will be used in proving the main result of this section.

**Property 3.2.2.** *At any equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$ , for every MFI  $i$ ,  $r_i^* > \max_{j \in V_i} e_j - 1$ .*

*Proof.* Consider any MFI  $i$ . Let  $k \in \operatorname{argmax}_{j \in V_i} e_j$ . Suppose that  $r_i^* \leq e_k - 1$ . Then we get  $(1 + r_i^* - e_k) \leq 0$ . This allows  $x_{j,i}^*$  to go to  $+\infty$ , and that violates the constraint  $\sum_{j \in V_i} x_{j,i}^* \leq T_i$ , contradicting the assumption that this is an equilibrium point.<sup>5</sup>  $\square$

Following are two related results that preclude certain trivial allocations at an equilibrium point.

**Property 3.2.3.** *At any equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$ , for any village  $j$ , there exists an MFI  $i \in B_j$  such that  $x_{j,i}^* > 0$ .*

*Proof.* Suppose that for some village  $j$ , and for all  $i \in B_j$  we have  $x_{j,i}^* = 0$ . Since  $d_j > 0$ , the constraint  $\sum_{i \in B_j} x_{j,i}^* (1 + r_i^* - e_j) \leq d_j$  of  $(P_V)$  is satisfied, but village  $j$  is not maximizing  $\sum_{i \in B_j} x_{j,i}^*$ . This contradicts that  $(\mathbf{x}^*, \mathbf{r}^*)$  is an equilibrium point.  $\square$

---

to demand  $x_{j,i}^* = +\infty$  even though  $T_i$  is finite for MFI  $i$ . As mentioned earlier, the reason is that the MFI-side optimization problem  $(P_M)$  treats  $x_{j,i}^*$  as exogenous and does not have a direct control over it inside  $(P_M)$ . Moreover, the village-side optimization problem  $(P_V)$  for village  $j$  selects  $(x_{j,i}^*)_{i \in B_j}$  in order to maximize its objective function  $\sum_{i \in B_j} x_{j,i}^*$ , without considering the MFI-side constraint  $\sum_{j \in V_i} x_{j,i}^* \leq T_i$ . The contradiction is due to the necessary condition that *at any equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$* , all the constraints of both  $(P_M)$  and  $(P_V)$  must be satisfied.

<sup>5</sup>Once again, village  $j$ 's demand  $x_{j,i}^*$  is determined in  $(P_V)$  without trying to satisfy the constraints of  $(P_M)$ . However, at an equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$ , both the village-side and the MFI-side problems must be optimized simultaneously.

**Property 3.2.4.** *At any equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$ , for any MFI  $i$ , there exists a village  $j \in V_i$  such that  $x_{j,i}^* > 0$ .*

*Proof.* If there exists an MFI  $i$  such that for all villages  $j \in V_i$ ,  $x_{j,i}^* = 0$ , then this violates the first constraint of  $(P_M)$  in the following way. By Property 3.2.2,  $r_i^* > 0$ , and by our modeling assumption,  $T_i > 0$ . Therefore,  $r_i^* \left( T_i - \sum_{j \in V_i} x_{j,i}^* \right) > 0$ .  $\square$

### 3.2.3 Eisenberg-Gale Formulation

We now present an Eisenberg-Gale convex program formulation of a restricted case of our model where the diversification parameter  $\lambda = 0$  and all the villages  $j$ ,  $1 \leq j \leq m$ , have the same revenue generation function  $g_j(l) := d + el$ , where  $d > 0$  and  $e \geq 1$  are constants. We will first prove that this case is equivalent to the following Eisenberg-Gale convex program [33, 118], which will give us the existence of an equilibrium point and the uniqueness of the equilibrium interest rates as a corollary.

Let us explain our overall plan here. We will first write down the Eisenberg-Gale convex program  $(P_E)$  below. We will then make a connection between an equilibrium point  $(\mathbf{x}^*, \mathbf{r}^*)$  of a microfinance market and the variables of program  $(P_E)$ . In particular, we will define  $x_{j,i}^* \equiv z_{j,i}^*$  and express  $r_i^*$  in terms of certain dual variables of  $(P_E)$ . Once we do that, we will show that the equilibrium conditions of  $(P_M)$  and  $(P_V)$  for the above mentioned special case are equivalent to the Karush-Kuhn-Tucker (KKT) conditions of  $(P_E)$ . Let us begin by writing down the Eisenberg-Gale program  $(P_E)$ .<sup>6</sup>

$$\begin{aligned} \min_{\mathbf{z}} \quad & \sum_{j=1}^m -\log \sum_{i \in B_j} z_{j,i} \\ \text{subject to} \quad & \sum_{j \in V_i} z_{j,i} - T_i \leq 0, \quad 1 \leq i \leq n \\ & z_{j,i} \geq 0, \quad 1 \leq i \leq n, j \in V_i \end{aligned} \tag{P_E}$$

Following are the Karush-Kuhn-Tucker (KKT) conditions for  $(P_E)$ .

---

<sup>6</sup>This is not *exactly* the convex program that Eisenberg and Gale defined [33, p. 166], but rather a simple variant of that.



*Stationary condition:*

$$\nabla_z \left( \sum_{j=1}^m -\log \sum_{i \in B_j} z_{j,i} \right) + \sum_{i=1}^n \gamma_i \nabla_z \left( \sum_{j \in V_i} z_{j,i} - T_i \right) + \sum_{i=1}^n \sum_{j \in V_i} \mu_{j,i} \nabla_z (-z_{j,i}) = 0$$

Evaluating this at  $z_{j,i}^*$  for any  $i \in \{1, \dots, n\}$  and any  $j \in V_i$ , we obtain the following.<sup>7</sup>

$$-\frac{1}{\sum_{k \in B_j} z_{j,k}^*} + \gamma_i^* - \mu_{j,i}^* = 0 \quad (3.1)$$

*Primal feasibility:*

$$\begin{aligned} \sum_{j \in V_i} z_{j,i}^* - T_i &\leq 0, & 1 \leq i \leq n \\ z_{j,i}^* &\geq 0, & 1 \leq i \leq n, j \in V_i \end{aligned}$$

*Dual feasibility:*

$$\begin{aligned} \gamma_i^* &\geq 0, & 1 \leq i \leq n \\ \mu_{j,i}^* &\geq 0, & 1 \leq i \leq n, j \in V_i \end{aligned}$$

*Complementary slackness:*

$$\gamma_i^* \left( \sum_{j \in V_i} z_{j,i}^* - T_i \right) = 0, \quad 1 \leq i \leq n \quad (3.2)$$

$$\mu_{j,i}^* (-z_{j,i}^*) = 0, \quad 1 \leq i \leq n, j \in V_i \quad (3.3)$$

Note that if  $\gamma_i^* > 0$ , then (3.2) gives us the following.

$$\sum_{j \in V_i} z_{j,i}^* - T_i = 0, \quad 1 \leq i \leq n \quad (3.4)$$

2

Furthermore, if  $z_{j,i}^* > 0$  then (3.3) implies  $\mu_{j,i}^* = 0$ . In that case, we obtain the following from the stationary condition (3.1).

$$\gamma_i^* = \frac{1}{\sum_{k \in B_j} z_{j,k}^*} \quad (3.5)$$

---

<sup>7</sup>The quantities at an optimal solution are denoted by \*.

The following properties are obtained from the optimality condition of the above Eisenberg-Gale convex program ( $P_E$ ).

**Lemma 3.2.5.** *For any  $i$ , there exists a  $j \in V_i$  such that  $z_{j,i}^* > 0$ .*

*Proof.* Suppose that for some  $i$ , and for all  $j \in V_i$ ,  $z_{j,i}^* = 0$ . This contradicts the optimality condition, because  $T_i > 0$ , and  $\sum_{j \in V_i} z_{j,i}^* = 0$ . Thus,  $\sum_{j \in V_i} z_{j,i}^* - T_i < 0$ , and  $\sum_{j=1}^m -\log \sum_{i \in B_j} z_{j,i}^*$  can be further decreased by increasing the value of  $z_{j,i}^*$  for some  $j$ .  $\square$

Let us define  $I^*(j) \equiv \{i \mid z_{j,i}^* > 0\}$ . We will later see that this represents the set of MFIs from which a village  $j$  borrows at an equilibrium point.

**Lemma 3.2.6.** *For any  $j$ ,  $|I^*(j)| > 0$ .*

*Proof.* Suppose that  $z_{j,i}^* = 0$  for some  $j$  and for all  $i \in B_j$ . Rearranging the terms of (3.1), we have for any  $i \in B_j$ :

$$\gamma_i^* = \frac{1}{\sum_{k \in B_j} z_{j,k}^*} + \mu_{j,i}^*.$$

Since  $\mu_{j,i}^* \geq 0$  by the dual feasibility condition, we have  $\gamma_i^* = +\infty$  from the above expression. This contradicts the complementary slackness condition (3.2), because  $T_i > 0$  by our modeling assumption. Therefore, for any  $j$  and some  $i \in B_j$ ,  $z_{j,i}^* > 0$ , which completes the proof.  $\square$

Another way of proving Lemma 3.2.6 is to note that if  $z_{j,i}^* = 0$  for some  $j$  and for all  $i \in B_j$ , then the objective function of the Eisenberg-Gale program ( $P_E$ ) goes to  $+\infty$ . This cannot happen, because ( $P_E$ ) is minimizing the objective function, and the program is guaranteed to have a bounded optimal solution (for example, one bounded feasible solution is achieved by  $z_{j,i} = \frac{T_i}{|V_i|}$  for all  $i \in \{1, \dots, n\}$  and all  $j \in V_i$ ).

We can use Lemma 3.2.6 to rewrite (3.5) in terms of  $I^*(j)$ . For any  $j$  and any  $i^*(j) \in I^*(j)$ , the following holds.

$$\gamma_{i^*(j)}^* = \frac{1}{\sum_{k \in B_j} z_{j,k}^*} \quad (3.6)$$

To present an Eisenberg-Gale formulation of our market model given by ( $P_M$ ) and ( $P_V$ ), we define the following terms.

$$x_{j,i}^* \equiv z_{j,i}^*, \quad \text{for all } i \in \{1, \dots, n\} \text{ and all } j \in V_i \quad (3.7)$$

$$r_{i^*(j)}^* \equiv \gamma_{i^*(j)}^* d + e - 1, \quad \text{for all } j \in \{1, \dots, m\} \text{ and all } i^*(j) \in I^*(j) \quad (3.8)$$

Note that in (3.8) above,  $r_i^*$  has not been explicitly defined for all  $i \in 1, \dots, n$ . We first prove that  $\cup_{j \in \{1, \dots, m\}} I^*(j) = \{1, \dots, n\}$  (that is, for all  $i$  we have defined  $r_i^*$  above). We then prove that if  $i^*(j) = i^*(j')$ , where  $i^*(j) \in I^*(j)$  and  $i^*(j') \in I^*(j')$  for  $j \neq j'$ , then  $r_{i^*(j)}^* = r_{i^*(j')}^*$  (that is, if the same  $i$  appears in two different  $I^*(\cdot)$ , then the definition of  $r_i^*$  is consistent with respect to these two cases).

For the first claim, suppose that for some  $i$ ,  $r_i^*$  has not been defined. This implies that for all  $j$ ,  $i \notin I^*(j)$ . That is, for all  $j$ ,  $z_{j,i}^* = 0$ , which violates Lemma 3.2.5.

For the second claim, consider the definition of  $r_{i^*(j)}^*$ .

$$\begin{aligned} r_{i^*(j)}^* &= \gamma_{i^*(j)}^* d + e - 1 \\ &= \gamma_{i^*(j')}^* d + e - 1 && \text{[Since } i^*(j) = i^*(j')\text{]} \\ &= r_{i^*(j')}^* && \text{[By definition]} \end{aligned}$$

Next, we use the definitions (3.7) and (3.8) above to show that none of the villages has any left-over money.

$$\begin{aligned} d &= \frac{1 + r_{i^*(j)}^* - e}{\gamma_{i^*(j)}^*} && \text{[Rearranging (3.8)]} \\ &= (1 + r_{i^*(j)}^* - e) \sum_{k \in B_j} z_{j,k}^* && \text{[Using (3.6)]} \end{aligned} \quad (3.9)$$

Next, we show that for any  $i^*(j) \in I^*(j)$ ,

$$\gamma_{i^*(j)}^* = \min_{k \in B_j} \gamma_k^*.$$

By Equation (3.6), for any  $i^*(j) \in I^*(j)$ ,

$$\gamma_{i^*(j)}^* = \frac{1}{\sum_{k \in B_j} z_{j,k}^*}.$$

For any  $l \in B_j$ , we obtain from the stationary condition,

$$\gamma_l^* \geq \frac{1}{\sum_{k \in B_j} z_{j,k}^*}, \text{ since } \mu_{j,l}^* \geq 0.$$

Therefore, for any  $i^*(j) \in I^*(j)$ ,

$$\gamma_{i^*(j)}^* = \min_{k \in B_j} \gamma_k^*.$$

Thus, using the definition of  $r$ ,  $r_{i^*(j)}^* = \min_{k \in B_j} r_k^*$ . Furthermore, for any  $l \in B_j - I^*(j)$ ,  $z_{j,l}^* = 0$ . We obtain from (3.9),

$$\begin{aligned} d &= (1 + \min_{l \in B_j} r_l^* - e) \sum_{k \in B_j} z_{j,k}^* \\ &= \sum_{k \in B_j} z_{j,k}^* (1 + r_k^* - e). \end{aligned}$$

Using the definition of  $x_{j,k}^*$  from (3.7),

$$d = \sum_{k \in B_j} x_{j,k}^* (1 + r_k^* - e).$$

Furthermore, by Lemma 3.2.5, for any MFI  $i$ , there exists a village  $j \in V_i$  such that  $z_{j,i}^* > 0$ . Thus, we get  $\mu_{j,i}^* = 0$ . The stationary condition gives us

$$\gamma_i^* = \frac{1}{\sum_{k \in B_j} z_{j,k}^*} > 0.$$

That is, for each MFI  $i$ ,  $\gamma_i^* > 0$ . Therefore,  $r_i^* > 0$  by (3.8). Also, (3.4) holds for  $\gamma_i^* > 0$ . Again, using the definition of  $x_{j,i}^*$  from (3.7), the other equilibrium condition for our model can be obtained from (3.4):

$$T_i - \sum_{j \in V_i} x_{j,i}^* = 0.$$

Thus, we have the following theorem and corollary.

**Theorem 3.2.7.** *The special case of microfinance markets with identical villages and no loan portfolio diversification, has an equivalent Eisenberg-Gale formulation.*

**Corollary 3.2.8.** *For the above special case, there exists an equilibrium point with unique interest rates [33] and a combinatorial polynomial-time algorithm to compute it [118].*

Theorem 3.2.7 allows us to make a connection between a more restricted case of our model and the linear Fisher model. When  $\lambda = 0$ , and all the

villages have an identical revenue generation function and all the MFIs have the same total amount of money, our model is indeed a graphical linear Fisher model where all the utility coefficients are set to 1. The reader is referred to the convex program 5.1 [118] and the Eisenberg-Gale formulation ( $P_E$ ) above to verify this.

### 3.2.4 Equilibrium Properties of General Case

We are now back to the general case of the problem, formally specified by mathematical programs ( $P_M$ ) and ( $P_V$ ) in Section 3.2.1. We begin with an expository discussion on the village objective function, which can be rewritten as  $\sum_{i \in B_j} x_{j,i} - \lambda \sum_{i \in B_j} x_{j,i} \log x_{j,i}$ . While the first term wants to maximize the total amount of loan, the second term wants, in colloquial terms, “not to put all the eggs in one basket.” For this reason, we name it the *diversification term*. The extent of this diversification is controlled by the exogenous parameter  $\lambda \geq 0$ . However, if  $\lambda$  is sufficiently small, then the first term dominates the second term, i.e.,  $\sum_{i \in B_j} x_{j,i} \geq \lambda \sum_{i \in B_j} x_{j,i} \log x_{j,i}$ . This roughly says that it is more important to the villages to obtain as much loan as possible than to diversify its loan portfolio. We deem this as a desirable property of the model. For theoretical purposes, it would suffice to assume the following.

**Assumption 3.2.1.**

$$0 \leq \lambda \leq \frac{1}{2 + \log T_{max}}$$

where  $T_{max} \equiv \max_i T_i$  and w.l.o.g.,  $T_i > 1$  for all  $i$ .

The following equilibrium properties will be used in the next section.

**Property 3.2.9.** *The first constraint of the village-side optimization program ( $P_V$ ) must be tight at any equilibrium point.*

*Proof.* Suppose that this is not the case, i.e., at an optimal solution  $\mathbf{x}_j^*$  for some village  $j$ ,  $\sum_{i \in B_j} x_{j,i}^* (1 + r_i - e_j) < d_j$ . We will show that village  $j$  can improve its objective function by slightly increasing  $x_{j,i}^*$  for any  $i \in B_j$  while maintaining the constraint. The derivative of the village-side objective function w.r.t.  $x_{j,i}$  is

$$1 - \lambda \log x_{j,i}^* - \lambda$$

which is positive (by Assumption 3.2.1 and equilibrium condition  $x_{j,i}^* \leq T_i$ ).  $\square$

We define  $e_{max}^i \equiv \max_{j \in V_i} e_j$  and  $d_{max}^i \equiv \max_{j \in V_i} d_j$  and obtain the following bounds on interest rates.

**Property 3.2.10.** *At any equilibrium point, for each MFI  $i$ ,  $e_{max}^i - 1 < r_i^* \leq \frac{|V_i|d_{max}^i}{T_i} + e_{max}^i - 1$ .*

*Proof.* Proof of  $e_{max}^i - 1 < r_i^*$  is similar to the proof of Property 3.2.2. Although compared to Property 3.2.2, we have a different objective function here, the proof of Property 3.2.9 shows that increasing  $x_{j,i}$  also increases the village objective function.

For the proof of the upper bound, the total amount of loan that villages in  $V_i$  can seek from MFI  $i$  is at most  $\sum_{j \in V_i} \frac{d_j}{1+r_i^* - e_j}$  (this bound is obtained using the first constraint in the village-side optimization program ( $P_V$ ), when each village in  $V_i$  seeks loan *only* from MFI  $i$ ). We have the following at an equilibrium point.

$$\begin{aligned} T_i &\leq \sum_{j \in V_i} \frac{d_j}{1+r_i^* - e_j} \\ &\leq d_{max}^i \sum_{j \in V_i} \frac{1}{1+r_i^* - e_j} \\ &\leq d_{max}^i \frac{|V_i|}{1+r_i^* - e_{max}^i} \end{aligned}$$

Rewriting this, we obtain  $r_i^* \leq \frac{|V_i|d_{max}^i}{T_i} + e_{max}^i - 1$ . □

### 3.3 Computational Scheme

In this section, we will primarily focus on two computational problems—learning the parameters of the model from data and designing an algorithm to compute an equilibrium point. Although in practice, we first solve the learning problem before going on to computing an equilibrium point, for clarity of presentation, we will reverse the order here.

#### 3.3.1 Computing an Equilibrium Point

We will now give a constructive proof of the existence of an equilibrium point in the microfinance market defined by ( $P_M$ ) and ( $P_V$ ) in Section 3.2.1. The

input to the problem is specified by the diversification parameter  $\lambda > 0$  (the case of  $\lambda = 0$  has been discussed earlier), the revenue generation parameters  $e_j$  and  $d_j$  of each village  $j$ , and the total disbursement amount  $T_i$  of each MFI  $i$ . The goal is to compute an equilibrium point consisting of an interest rate  $r_i^*$  for each MFI  $i$  and a vector of loan allocations  $\mathbf{x}_j^*$  for each village  $j$ . We will also prove that an equilibrium point always exists. But first, let us give a brief outline of the equilibrium computation scheme in Algorithm 1. The ensuing technical results will be presented in the context of this outline.

---

**Algorithm 1** Outline of Equilibrium Computation

---

- 1: For each MFI  $i$ , initialize  $r_i$  to  $e_{max}^i - 1$ .
  - 2: For each village  $j$ , compute its best response  $\mathbf{x}_j$ .
  - 3: **repeat**
  - 4:     **for all** MFI  $i$  **do**
  - 5:         **while**  $T_i \neq \sum_{j \in V_i} x_{j,i}$  **do**
  - 6:             Change  $r_i$  as described later.
  - 7:             For each village  $j \in V_i$ , update its best response  $\mathbf{x}_j$  reflecting the change in  $r_i$ .
  - 8:         **end while**
  - 9:     **end for**
  - 10: **until** no change in  $r_i$  occurs for any  $i$
- 

Before going on to the details of how to change  $r_i$  in Line 6 of Algorithm 1, we will characterize the best response of the villages used in Line 7.

**Lemma 3.3.1. (Village's Best Response)** *Given the interest rates of all the MFIs, the following is the unique best response of any village  $j$  to any MFI  $i \in B_j$ :*

$$x_{j,i}^* = \exp\left(\frac{1 - \lambda - \alpha_j^*(1 + r_i - e_j)}{\lambda}\right) \quad (3.10)$$

where  $\alpha_j^* \geq 0$  is the unique solution to

$$\sum_{i \in B_j} \exp\left(\frac{1 - \lambda - \alpha_j^*(1 + r_i - e_j)}{\lambda}\right) (1 + r_i - e_j) = d_j. \quad (3.11)$$

*Proof.* Following is the Lagrangian of the village-side optimization program ( $P_V$ ) for village  $j$ :

$$L(\mathbf{x}_j, \alpha_j) = - \sum_{i \in B_j} x_{j,i} - \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}} + \alpha_j \left( \sum_{i \in B_j} x_{j,i} (1 + r_i - e_j) - d_j \right).$$

At an optimal solution, we have  $\frac{\delta L}{\delta x_{j,i}} = 0$  for any  $i \in B_j$ . This is expanded below:

$$\begin{aligned} -1 - \lambda \log \frac{1}{x_{j,i}^*} + \lambda x_{j,i}^* \frac{1}{x_{j,i}^*} + \alpha_j (1 + r_i - e_j) &= 0 \\ \Leftrightarrow x_{j,i}^* &= \exp \left( \frac{1 - \lambda - \alpha_j (1 + r_i - e_j)}{\lambda} \right). \end{aligned}$$

By Property 3.2.9,  $\sum_{i \in B_j} x_{j,i}^* (1 + r_i - e_j) = d_j$ . Substituting the expression for  $x_{j,i}^*$  we obtain the second equation claimed in the statement. Moreover,  $\alpha_j^*$  must be unique; otherwise, by the above expression for  $x_{j,i}^*$ , we would have multiplicity in the best response of village  $j$ , which is precluded by the convex optimization ( $P_V$ ).  $\square$

In the above characterization of the village best response, as soon as the interest rate  $r_i$  of some MFI  $i$  changes in Line 6 of Algorithm 1, both the best response allocation  $x_{j,i}^*$  and the Lagrange multiplier  $\alpha_j^*$  change in Line 7, for any village  $j \in V_i$ . Next, we show the direction of these changes.

**Lemma 3.3.2.** *Whenever  $r_i$  increases (decreases) in Line 6,  $x_{j,i}$  must decrease (increase) for every village  $j \in V_i$  in Line 7 of Algorithm 1.*

*Proof.* We prove the case of  $r_i$  increasing. The other case can be proved in the same way. First, observe that we cannot simply invoke Equation (3.10) to prove the statement, because  $\alpha_j^*$  has also changed once  $r_i$  has changed and the direction of change of  $\alpha_j^*$  is not immediately clear from Equation (3.11).

Here, we treat the terms  $r_i$ ,  $x_{j,i}^*$ , and  $\alpha_j^*$  as *names* of variables instantiated with specific values at each iteration of Lines 6 and 7 of Algorithm 1. Suppose that the value of  $r_i$  has been increased in Line 6. Suppose, for a contradiction, that in response to this increase, some village  $j \in V_i$  has either increased its value of  $x_{j,i}^*$  or kept it unchanged in Line 7. By Property 3.2.9, the first constraint of the village-side program ( $P_V$ ) is tight at any optimal solution, including village  $j$ 's previous best response in Line 7 (i.e., the old values of



$\mathbf{x}_j$  just before the current update in Line 7). Therefore, in the current best response, village  $j$  must decrease the value of  $x_{j,k}^*$  for some  $k \in B_j$  (otherwise that constraint cannot be satisfied, because  $r_i$  has increased). Rewriting the expression of  $x_{j,i}^*$  given in Lemma 3.3.1 in terms of  $\alpha_j^*$ , we obtain the first equation below. The second equation follows similarly.

$$\alpha_j^* = \frac{1 - \lambda - \lambda \log x_{j,i}^*}{1 + r_i - e_j} \quad (3.12)$$

$$\alpha_j^* = \frac{1 - \lambda - \lambda \log x_{j,k}^*}{1 + r_k - e_j} \quad (3.13)$$

By Equation (3.12), our assumption that  $x_{j,i}^*$  has increased in response to the increase of  $r_i$  implies that the value of  $\alpha_j^*$  has decreased from its previous one. Therefore, by Equation (3.13), the value of  $x_{j,k}^*$  must increase, which gives us a contradiction (note that  $r_k$  has not been changed, i.e., its value remains the same as the one during village  $j$ 's previous best response). Therefore, whenever  $r_i$  increases in Line 6,  $x_{j,i}^*$  must decrease, for all  $j \in V_i$ .  $\square$

The next lemma is a cornerstone of our theoretical results. Here, we use the term *turn of an MFI* to refer to the iterative execution of Line 6, wherein an MFI tries to set its interest rate to make supply equal demand. At the end of its turn, an MFI has successfully set its interest rate to achieve this objective.

**Lemma 3.3.3. (Strategic Complementarity)** *Suppose that an MFI  $i$  has increased its interest rate at the end of its turn. Thereafter, it cannot be the best response of any other MFI  $k$  to lower its interest rate when its turn comes in the algorithm.*

*Proof.* Consider a village  $j \in V_i$ . By Lemma 3.3.2, when an MFI  $i$  increases its interest rate in Line 6, village  $j$  must decrease  $x_{j,i}^*$  in Line 7. Considering Equation (3.12), it may at first seem possible that the value of  $\alpha_j^*$  can increase, decrease, or even remain the same, depending on how much  $x_{j,i}^*$  has decreased. However, we will next show that  $\alpha_j^*$  cannot increase. For this, we define  $\beta_j^* \equiv \frac{\alpha_j^*}{\lambda}$  and  $\rho_{j,i} \equiv 1 + r_i - e_j$  and rewrite Equation (3.11) as follows.

$$\begin{aligned} & \sum_{i \in B_j} \exp\left(\frac{1-\lambda}{\lambda}\right) \exp(-\beta_j^* \rho_{j,i}) \rho_{j,i} = d_j \\ \Leftrightarrow & \sum_{i \in B_j} \frac{\rho_{j,i}}{\exp(\beta_j^* \rho_{j,i})} = d_j \exp\left(\frac{-1+\lambda}{\lambda}\right) \end{aligned}$$

Here, the right hand side is constant, since  $\lambda$  and  $d_j$  are both constants. Consider the left hand side. It suffices to show that if we increase  $\rho_{j,i}$  (i.e., increase  $r_i$ ) by *any amount*, but keep  $\beta_j^*$  unchanged, then the left hand side must decrease.<sup>8</sup> In this case, only one term of the sum on the left hand side changes:  $\frac{\rho_{j,i}}{\exp(\beta_j^* \rho_{j,i})}$ . We show that the derivative of this term w.r.t.  $\rho_{j,i}$  is non-positive.

$$\begin{aligned} & \frac{1}{\exp(\beta_j^* \rho_{j,i})} - \frac{\beta_j^* \rho_{j,i}}{\exp(\beta_j^* \rho_{j,i})} \leq 0 \\ \Leftrightarrow & \rho_{j,i} \beta_j^* \geq 1 \\ \Leftrightarrow & (1 + r_i - e_j) \frac{\alpha_j^*}{\lambda} \geq 1 \\ \Leftrightarrow & 1 - \lambda - \lambda \log x_{j,i}^* \geq \lambda, \text{ by Equation (3.12)} \\ \Leftrightarrow & \lambda \leq \frac{1}{2 + \log x_{j,i}^*} \end{aligned}$$

which holds by Assumption 3.2.1. Therefore,  $\alpha_j$  cannot increase when  $r_i$  increases.

Since  $\alpha_j$  can only decrease when  $r_i$  increases, using Equation (3.13) we obtain that in Line 7 of the algorithm, village  $j$  cannot decrease  $x_{j,k}^*$  for any  $k \neq i \in B_j$ . Thus, when its next turn comes, MFI  $k$  can only find a rise in demand for its loans, which can only exceed  $T_k$ , since at the end of every turn, an MFI successfully sets its interest rate so that the demand for its loan equals its supply. Therefore, by Lemma 3.3.2, decreasing its interest rate cannot be MFI  $k$ 's best response.  $\square$

In essence, Lemma 3.3.2 is a result of *strategic substitutability* [32] between the MFI and the village sides, while Lemma 3.3.3 is a result of *strategic complementarity* [23] among the MFIs. We will see that our algorithm exploits these two properties as we fill in the details of Lines 6 and 7 next.

---

<sup>8</sup>Note that increasing  $\beta_j^*$  will only further decrease the left hand side.

### Line 6: MFI's Best Response

By Lemma 3.3.2, the total demand for MFI  $i$ 's loan monotonically decreases with the increase of  $r_i$ . Therefore, a simple search, such as a binary search, between the upper and the lower bounds of  $r_i$  as stated in Property 3.2.10, can efficiently find the “right” value of  $r_i$  that makes supply equal demand for  $i$ . For example, in the first iteration of the while loop, in Line 6,  $r_i$  is set to the midpoint  $r_i^m = \frac{r_i^l + r_i^h}{2}$  between its lower bound  $r_i^l$  and upper bounds  $r_i^h$ . Then the best response of the villages are computed in the next line. If still  $T_i \neq \sum_{j \in V_i} x_{j,i}$  then in the next iteration, in Line 6,  $r_i$  is set to either  $\frac{r_i^l + r_i^m}{2}$  or  $\frac{r_i^m + r_i^h}{2}$  depending on whether  $T_i > \sum_{j \in V_i} x_{j,i}$  or the opposite, respectively. The search progresses in this way until  $T_i = \sum_{j \in V_i} x_{j,i}$ . As an implementation note, to circumvent issues of numerical precision, we can adopt the notion of  $\epsilon$ -equilibrium point, where the market  $\epsilon$ -clears (i.e., the absolute value of the difference between supply and demand for each MFI  $i$  is below  $\epsilon$ ) and each village plays its  $\epsilon$ -best response (i.e., it cannot improve its objective function more than  $\epsilon$  by changing its current response). Having said that, all of our results hold for  $\epsilon = 0$ .

### Line 7: Village's Best Response

We use Lemma 3.3.1 to compute each village  $j$ 's best response  $x_{j,i}^*$  to MFIs  $i \in B_j$ . However, Equation (3.10) requires computation of  $\alpha_j^*$ , the solution to Equation (3.11). We can exploit the convexity of the left hand side of Equation (3.11) to design a simple search algorithm to find  $\alpha_j^*$  up to a desired numerical accuracy.

Next, we make the following statement about our constructive proof of the existence of an equilibrium point.

**Theorem 3.3.4.** *There always exists an equilibrium point in a microfinance market specified by programs  $(P_M)$  and  $(P_V)$ .*

*Proof.* Algorithm 1 begins with initial values of interest rates arbitrarily close to their lower bound established in Property 3.2.10. Thereafter, by Lemma 3.3.3, these interest rates can only increase, and by Lemma 3.3.1, every village has a unique best response to these interest rates. Now, the interest rates are upper bounded by Property 3.2.10. Therefore, by the well-known monotone convergence theorem, the process of incrementing the interest rates must come to an end. And that point of termination must be an equilibrium point.  $\square$

### 3.3.2 Learning the Parameters of the Model

We take an optimization approach to learning the parameters of the model. The inputs to the parameter learning problem are the spatial structure of the market (specified by  $B_j$  for all village  $j$ , or equivalently, by  $V_i$  for all MFI  $i$ ), the observed loan allocations  $\tilde{x}_{j,i}$  for all village  $j$  and all MFI  $i \in B_j$ , the observed interest rates  $\tilde{r}_i$  and total supply  $T_i$  for all MFI  $i$ . The objective of the learning scheme is to instantiate parameters  $e_j$  and  $d_j$  of the revenue generation function of each village  $j$ , as well as to compute allocations  $\mathbf{x}^*$  and interest rates  $\mathbf{r}^*$  that satisfy equilibrium conditions. In essence, the goal is to learn the parameters of the model so that an equilibrium point closely approximates the observed data.<sup>9</sup> Using index  $i$  for MFIs and  $j$  for villages, the optimization program is formally defined below.

$$\min_{\mathbf{e}, \mathbf{d}, \mathbf{r}} \sum_i \sum_{j \in V_i} (x_{j,i}^* - \tilde{x}_{j,i})^2 + C \sum_i (r_i^* - \tilde{r}_i)^2$$

such that

for all  $j$ ,

$$\mathbf{x}_j^* \in \arg \max_{\mathbf{x}_j} \sum_{i \in B_j} x_{j,i} + \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}}$$

$$\text{s. t. } \sum_{i \in B_j} x_{j,i} (1 + r_i^* - e_j) \leq d_j$$

$$\mathbf{x}_j \geq 0$$

(3.14)

$$\mathbf{e}_j \geq 1, \mathbf{d}_j \geq 0$$

$$\sum_{j \in V_i} x_{j,i}^* = T_i, \text{ for all } i$$

$$r_i \geq e_j - 1, \text{ for all } i \text{ and all } j \in V_i$$

Above is a nested (bi-level) optimization program. The term  $C$  in the objective function is a constant. In the interior optimization program, we always make sure that  $\mathbf{x}^*$  are best responses of the villages, with respect to the model parameters and interest rates  $\mathbf{r}^*$ . In fact, we exploit the results of Lemma 3.3.1 to compute  $\mathbf{x}^*$  more efficiently. That is, instead of searching for  $x_{j,i}$ 's that optimize the interior objective function, it suffices to search for Lagrange multipliers  $\alpha_j$  in a much smaller search space and then apply

---

<sup>9</sup>Note that we do not assume the observed data to correspond to an equilibrium point, but rather want to instantiate our model such that it would lead to an equilibrium point that is “close” to the observed data.

Equation (3.10) to compute  $\mathbf{x}^*$ . We used the interior-point algorithm of Matlab’s large-scale optimization package to solve the above program (the interior optimization program has been solved using the results of Lemma 3.3.1).

Initialization plays a big role in solving this problem fast, especially in large instances (e.g., the instance with data from Bangladesh that has thousands of constraints and variables). If we initialize the parameters arbitrarily, then the interior point algorithm spends an enormous amount of time searching for a feasible solution. Fortunately, we can avoid this issue by computing a feasible solution first. For this, with arbitrary values of  $\mathbf{e}$  and  $\mathbf{d}$  as inputs, we run Algorithm 1 and compute an initial equilibrium point with respect to the inputs of  $\mathbf{e}$  and  $\mathbf{d}$ . Subsequently, the learning procedure updates  $\mathbf{e}$  and  $\mathbf{d}$  compute an optimal solution. Using such an initial feasible solution to the above optimization problem, we observed a much faster convergence.

In the next section, we will show that the above learning procedure does not overfit the real-world data. We will also highlight the issue of equilibrium selection for parameter estimation.

## 3.4 Empirical Study

Our empirical study is based on microfinance data of Bangladesh and Bolivia. The reason we have chosen these two countries is that over time, microfinance programs in these two countries have behaved very much differently with respect to competition and interest rates [101].

### 3.4.1 Case Study: Bolivia

#### Data

We have obtained microfinance data of Bolivia from several sources, such as ASOFIN,<sup>10</sup> the apex body of MFIs in Bolivia, and the Central Bank of Bolivia.<sup>11</sup> We were only able to collect somewhat coarse, region-level data. The data, dated June 2011, consists of eight MFIs operating in 10 regions. These MFIs (and their interest rates) are: Bancosol (21.54%), Banco Los Andes (19.39%), Banco FIE (20.49%), Prodem (23.55%), Eco Futuro (29.25%), Fortaleza (21.22%), Fassil (22.38%), and Agro Capital (21%). The number of edges in the bipartite network is 65, out of a maximum possible 80.

---

<sup>10</sup><http://www.asofinbolivia.com>

<sup>11</sup><http://www.bcb.gob.bo/>

## Learning the Parameters of the Model

Given the exogenous parameter  $\lambda$ , the learning scheme above estimates the parameters  $e_j$  and  $d_j$  such that an equilibrium point of the game is a close approximation of the observed data. Let us first explain how we choose the exogenous diversification parameter  $\lambda$ .

Figure 3.1 shows how the objective function of the optimization program varies as a function of  $\lambda$ . It shows that for a range of smaller values of  $\lambda$ , the objective function value of the learning program is consistently small (note that the optimization routine wants to minimize the objective function). As  $\lambda$  grows, the objective function value oscillates a lot and is sometimes very high.

Also shown in Figure 3.1 is how the interest rates become dissimilar as  $\lambda$  is varied. For that, we first define a *negative entropy* term,  $C \sum_i \frac{r_i}{Z} \log \frac{Z}{r_i}$ , where  $Z = \sum_i r_i$ , where  $C$  is a constant set to 100. Here, the negative entropy quantifies the similarity among the interest rates. That is, high negative entropy means the interest rates among the MFIs are more similar. As we can see, at relatively low values of  $\lambda$ , the interest rates are similar to each other and as  $\lambda$  becomes high, they become very much dissimilar at some points.




Figure 3.1: *Optimal objective function values and “negative entropy” of interest rates as  $\lambda$  varies. This shows that the objective function value becomes large (which is undesirable) as  $\lambda$  grows.*

We choose  $\lambda = 0.05$ , because at this level of  $\lambda$ , the objective function value of the learning optimization is low as well as stable and the interest rates are also allowed to be relatively dissimilar. As we will show later, dissimilarities among the interest rates of the MFIs are very often observed in the real-world data.

The learned values of the parameters  $e_j$  and  $d_j$  for villages  $j$  in the Bolivia market capture the variation among the villages with respect to the revenue generation function. Although the rate  $e_j$  of revenue generation varies only from 1.001 to 1.234 among the villages, the variation in  $d_j$  is much greater.

The individual loan allocations learned from data closely approximate the observed allocations. In fact, the average relative deviation between these two allocations is only 4.41% (relative deviation is calculated by  $\frac{\text{abs}(x_{j,i} - \tilde{x}_{j,i})}{\sum_i \tilde{x}_{j,i}}$ ). Figure 3.2 shows this. The 45° line is the locus of equality between these two allocations.




Figure 3.2: *Learned allocations vs. observed allocations. This shows that the learned allocations closely approximate the allocations in the observed data.*

The learned model matches the total loan allocations of the MFIs due to the constraint of the program. As shown in Figure 3.9, the learned interest rates are, however, slightly different from the observed rates.

### Issues of Bias and Variance.

Our dataset consists of a single sample. As a result, the traditional approach of performing cross validation using hold-out sets or plotting learning curves by varying the number of samples would not work in our setting. To investigate whether our model overfits the data, we have applied the following procedure of systematically introducing noise to the observed data sample. In the case of overfitting, increasing the level of noise would lead the equilibrium outcome to be significantly different from the observed data.

We use a parameter  $\nu$  to control the level of noise. For any fixed  $\nu$ , we derive noisy samples by modifying each non-zero observed allocation  $\tilde{x}_{j,i}$  by adding to it a random noise (under certain noise models to be described later). We denote the resulting noisy allocation by  $x_{j,i}^\nu$  and treat the newly constructed noisy dataset as a training set and the observed dataset as test. We learn the parameters of the model using the training set. We then find an equilibrium allocation  $\mathbf{x}^*$  using Algorithm 1.<sup>12</sup> We compute the following mean relative

---

<sup>12</sup>Note that this equilibrium allocation would have remained the same all the time had we treated the the observed dataset as the training set.

deviation as the test error:  $\frac{1}{n} \sum_i \frac{1}{|V_i|} \sum_j \frac{\text{abs}(x_{j,i}^* - \tilde{x}_{j,i})}{T_i}$ . The training error is computed similarly (by replacing  $\tilde{x}_{j,i}$  by  $x_{j,i}^\nu$ ). For each noise level  $\nu$ , we perform the whole procedure a number of times and calculate the average error.

**Gaussian Noise Model.** In this model, we obtain  $x_{j,i}^\nu$  by adding to each non-zero observed allocation  $\tilde{x}_{j,i}$  a Gaussian random noise of mean 0 and standard deviation  $\nu\sigma(i)$ , where  $\sigma(i)$  is the standard deviation of the allocations of MFI  $i$  across all villages in which it operates. For the Bolivia dataset, we find that varying the noise level  $\nu$  between 0 and 1 and taking the average over 25 trials, both the training and test errors are below 5.83% and are close to each other (i.e., within the 95% confidence interval of each other). The learning curve, shown in Figure 3.3, does not suggest overfitting.




Figure 3.3: *Learning curve for the Bolivia dataset under Gaussian random noise (vertical bars denote 95% CI). The learning curve does not suggest overfitting.*

**Dirichlet Noise Model.** In this noise model, we derive noisy allocations while keeping the total amount of loan disbursed by each MFI the same as its observed total amount. We follow the commonly used procedure of deriving a Dirichlet distribution from a gamma distribution [43, Ch. 18]. We control the noise (i.e., variance) of the Dirichlet distribution using the parameter  $\nu$  in the following way. For each MFI  $i$ ,  $\mathbf{x}_i^\nu = T(i) \times \text{Dir}(\nu\tilde{\mathbf{x}}_i)$ .<sup>13</sup> As the  $\nu > 0$  increases, the variance of the distribution  $\text{Dir}(\nu\tilde{\mathbf{x}}_i)$  decreases. Varying  $\nu$  from  $2^{-5}$  (high variance) to  $2^{15}$  (low variance) and taking the average over 50 trials at each  $\nu$ , we found that the training and the test errors are within the confidence intervals of each other across the whole spectrum of noise levels. The maximum test error of 8.83% occurs at  $\nu = 2^{-5}$  where we also get the maximum offset of the 95% confidence interval, which is 0.66%. Once again, the learning curve, shown in Figure 3.4, does not suggest overfitting.

<sup>13</sup>Slightly abusing the notation, the vector  $\tilde{\mathbf{x}}_i$  corresponds to non-zero observed allocations only.



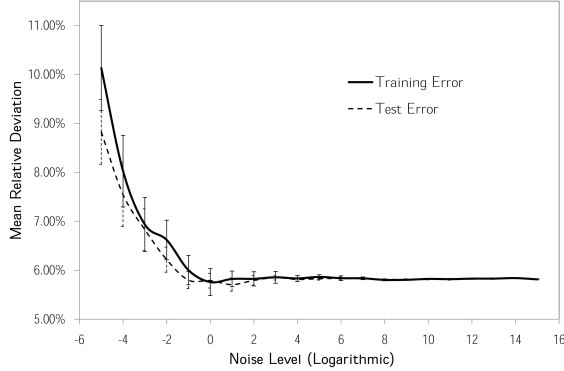


Figure 3.4: *Learning curve for the Bolivia dataset under the Dirichlet noise model. Logarithm of the noise parameter  $\nu$  is shown on the x-axis (vertical bars denote 95% CI). The learning curve in this noise model does not indicate overfitting as well.*

### Equilibrium Selection.

In practice, equilibrium selection is an important issue. In general, we cannot rule out the possibility of multiplicity of equilibria. In such cases, our learning scheme biases its search for an equilibrium point that most closely explains the data. One important question is: does the equilibrium point that we compute change drastically when noise is added to our data? In other words, how robust is our scheme? To answer this, we extend the above experimental procedure using the following bootstrapping scheme.<sup>14</sup>

Suppose that for each noise level  $\nu$ , we have  $t$  trials (i.e.,  $t$  noisy training sets, each derived from the observed dataset using a particular noise model with the given parameter  $\nu$ ). For each noise level  $\nu$ , we iterate the following procedure  $M$  times. At each iteration  $k$ , we uniformly sample  $t$  times (with replacement, of course) from the  $t$  noisy training sets and then compute the following relative mean equilibrium allocations:  $\hat{\mu}_{j,i} = \frac{1}{t} \sum_{l=1}^t \frac{x_{j,i}^{*(l)}}{T_i}$  (here,  $x_{j,i}^{*(l)}$  denotes equilibrium allocation for the  $l$ -th training set). Within the same  $k$ -th iteration, we compute the following average deviation from mean:  $\hat{\delta}(k) = \frac{1}{n} \sum_i \frac{1}{|V_i|} \sum_j \frac{1}{t} \sum_l \text{abs}(x_{j,i}^{*(l)}/T_i - \hat{\mu}_{j,i})$ . This quantity signifies the average distance of the equilibria of the sampled examples from the mean equilibrium. Now, for each value of  $\nu$ , we average this distance measure over these  $M$  iterations. We perform this bootstrapping procedure for various values of  $\nu$ .

<sup>14</sup>In the previous experiment on the issue of overfitting, the focus was on the distance between an equilibrium point and the data. Now, our focus is on the distance between different equilibrium points when noise is added to the data.

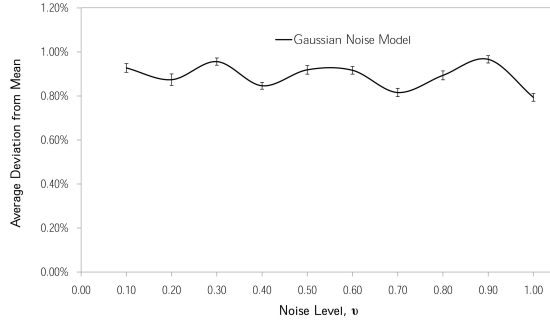


Figure 3.5: Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Gaussian noise model (vertical bars denote 95% CI). It shows that the equilibrium point computed is robust with respect to noise.

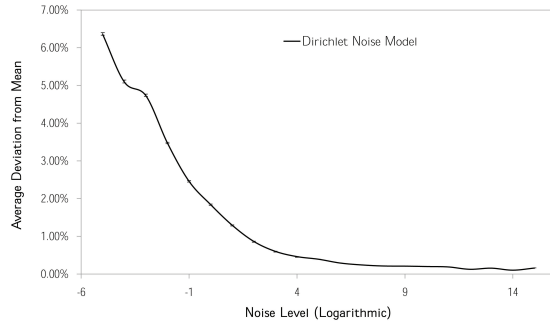


Figure 3.6: Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Dirichlet noise model. Logarithm of the noise parameter  $\nu$  is shown on the x-axis. Vertical bars denote 95% CI. It also shows the robustness of the computed equilibrium point although the noise model is different—Dirichlet.

For the Bolivia dataset, under the Gaussian noise model (described above) and using  $t = 25$  and  $M = 100$ , we found that this average distance varies from 0.79% to 0.96%, with the offset of the 95% CI ranging from 0.015% to 0.026% for varied noise levels  $0 < \nu \leq 1$ . On the other hand, for the Dirichlet noise model and using  $t = 50$  and  $M = 100$ , the maximum average distance is 6.35%, which happens at a very high variance parameterized by  $\nu = 2^{-5}$ . The minimum average distance of 0.10% happens at low variance with  $\nu = 2^{14}$ . The offset of the 95% CI ranges from 0.001% to 0.05% across all the noise levels considered. The plots for these two noise models are shown in Figures 3.5 and 3.6. Moreover, under both noise models, the equilibrium interest rates do not deviate much from the mean either. These suggest that an equilibrium point does not change much when noise is introduced to the data and that our

scheme is robust with respect to noise in the real world data.

### Equilibrium Computation

We discuss equilibrium computation on the model learned with  $\lambda = 0.05$ . First, we would like to remind the reader about the point we made regarding equilibrium selection. In practice, we have observed that the equilibrium computed by Algorithm 1 converges to the learned values of  $\mathbf{x}$  and  $\mathbf{r}$ , *even if we start with different initial values*. For example, Figure 3.7 shows the case of MFI Bancosol’s convergence to the same equilibrium interest rate despite different initialization (other interest rates were also differently initialized). This equilibrium interest rate is the same as the learned one. Not only that, as Figure 3.8 shows individual loan allocations were also almost the same.

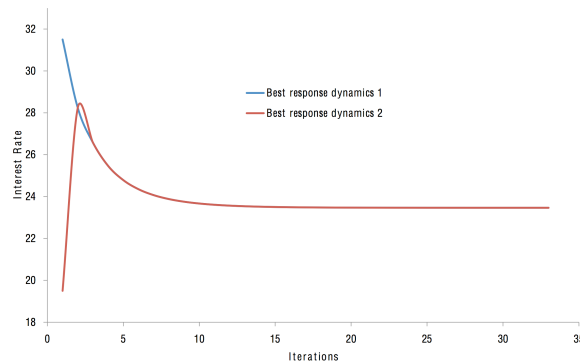


Figure 3.7: *Two best response dynamics of MFI Bancosol with different initialization. Both of these converged to the same solution.*

Finally, Figure 3.9 shows a comparison among the observed, learned, and equilibrium interest rates.

### 3.4.2 Case Study: Bangladesh

#### Data

We have obtained microfinance data, dated December 2005, from Palli Karma Sahayak Foundation (PKSF), which is the apex body of NGO MFIs in Bangladesh. There are seven major MFIs (or collection of MFIs) operating in 464 *upazillas* or collection of villages. The data can be simplified as a 464-by-7 matrix where an element in location  $(j, i)$  denotes the number of borrowers that MFI  $i$  has in village  $j$ . The bipartite network-structure induced by this data is very dense, consisting of 3096 edges out of a maximum possible 3248.




Figure 3.8: *Learned allocations vs. equilibrium allocations. The two allocations being almost identical, shows that the learning algorithm was able to capture an equilibrium point using the inner part of the nested optimization program.*

The seven major MFIs or bodies of MFIs (and their flat interest rates) are BRAC (15%), ASA (15%), PKSF partner organizations (12.5%), Grameen Bank(10%), BRDB (8%), Other government organizations (8%), and Other MFIs (12.5%) [101, 120].

### Learning the Parameters of the Model

Due to the size of Bangladesh data, we are posed with the problem of solving a nonlinear optimization problem of the order of thousands of variables and constraints. As discussed above, the interior point algorithm is initialized with a feasible solution, which makes computation much faster. Still, solving the problem takes time in the order of hours, compared to minutes for the Bolivia case.

Similar to the Bolivia case, the learned parameters  $e_j$  (rate of revenue generation) and  $d_j$  (revenue from other sources) show variation among the villages  $j$ . This is particularly the case with the estimated parameter  $d_j$ , while the estimation of  $e_j$  varies around 1.07 for all the villages. A more detailed analysis of the estimated parameters (for example, their correlation with access to resources such as rivers) is left for future work.

We also obtain a close approximation of observed individual allocations in the learned model (see Figure 3.10. For example, the average deviation is only 5.54% when  $\lambda = 0.05$ . The market clears in the learned model, and as shown in Figure 3.11, the learned interest rates are close to the actual ones, except for the government MFIs numbered 5 and 6, which are known to be operating inefficiently, i.e., with much lower interest rates (8%) than that required for




Figure 3.9: *Comparison among observed, learned, and equilibrium interest rates. The equilibrium interest rates and the observed interest rates do not completely match, which is fine as we do not assume the data to be an equilibrium point.*

sustainability without subsidies [29, 120].




Figure 3.10: *Learned allocations vs. observed allocations. Although they are not exactly the same, the learned allocations do approximate the observed ones.*

### Equilibrium Computation

Similar to the Bolivia case, we have observed that the best response dynamics of Algorithm 1 quickly converges to the allocations and interest rates of the learned model. Figures 3.12 shows the similarity between learned and equilibrium allocations.




Figure 3.11: Comparison among observed, learned, and equilibrium interest rates. Again, we do not assume the data to be an equilibrium point. Therefore, the difference between observed interest rates and equilibrium interest rates suggests an interesting aspect. That is, some banks are charging less than they should and some are charging more than they should. The first aspect is well-established in the literature with respect to government-owned banks. There is also practical evidence for the second aspect, which will be elaborated when we study interest rate caps.

### 3.5 Policy Experiments

The end goal of modeling microfinance economy is to be able to help policy makers in microfinance sector make decisions, as well as to evaluate possible interventions in the market. We take the following approach to studying the effects of interventions and policy decisions. For a specific intervention policy, e.g., removal of government-owned MFIs, we first learn the parameters of the model and then compute an equilibrium point, both in the original setting (before removal of any MFI). Using the parameters learned, we compute a new equilibrium point after the removal of the government-owned MFIs. Finally, we study changes in these two equilibria (before and after removal) in order to predict the effect of such an intervention. As we demonstrate below, our model can be used in a variety of settings regarding interventions and policy decisions.

*Role of subsidies.* It has been well documented that the sustainability of MFIs very much depend on subsidies [29, 89]. We can ask a related question: how does giving subsidies to an MFI affect the market? More concretely, how do equilibrium interest rates change due to subsidies? To answer this, we first learn the parameters of the model and compute an equilibrium point *before* the injection of new subsidies into an MFI. We then add the subsidies to the MFI's total amount of loans to be disbursed and compute a new equilibrium point.




Figure 3.12: *Learned allocations vs. equilibrium allocations. The two allocations are very similar.*




Figure 3.13: *Equilibrium interest rates before and after adding subsidies to Eco Futuro (MFI 5). Subsidies help bring down the interest rate of not only Eco Futuro but also the other MFIs.*

For instance, in the Bolivia case, one of the MFIs named Eco Futuro exhibits very high interest rates both in observed data and at an equilibrium point. A quick investigation of the data reveals that Eco Futuro is connected to all the villages, but has very little total loan to be disbursed compared to the leading MFI Bancosol (which is approximately one-fourth of that of Bancosol). As shown in Figure 3.13, if we inject further subsidies into Eco Futuro (MFI 5) to make its total loan amount equal to Bancosol’s (MFI 1), not only do these two MFIs have the same (but lower than before) equilibrium interest rates, it also drives down the interest rates of the other MFIs.

*Changes in interest rates.* Our model computes lower equilibrium interest rate (around 12%) for ASA than its observed interest rate (15%). It is interesting to note that in late 2005, ASA lowered its interest rate from 15% to

12.5%, which is close to what our model predicts at an equilibrium point.<sup>15</sup>

*Interest rates ceiling.* PKSF recently capped the interest rates of its partner organization to 12.5% [101], and more recently, the country's Microfinance Regulatory Authority has also imposed a ceiling on interest rate at around 13.5% flat.<sup>16</sup> Such evidence on interest rate ceiling is consistent with the outcome of our model, since in our model, 13.4975% is the highest interest rate at an equilibrium point.

*Government-owned MFIs.* It is well-documented in literature that many of the government-owned MFIs do not have the goal of meeting their operating costs [120]. Our model shows that an intervention by removing government-owned MFIs from the market would result in an increase of equilibrium interest rates by approximately 0.5% for every other MFI. This increase is not that large partly because the government-owned MFIs supply very little amount of loan to the market compared to the other MFIs. Yet, it does suggest that less competition leads to higher interest rates, which is consistent with empirical findings [101].

*Adding new branches.* Our model can be used to predict how the market would be affected if the underlying network structure is changed due to an MFI's opening of new branches. For instance, suppose that MFI Fassil in Bolivia expands its business to all villages (by adding six new branches). It may at first seem that due to the increase in competition, equilibrium interest rates would go down. However, since Fassil's total amount of loans does not change, the new connections and the ensuing increase in demand actually increases equilibrium interest rates of all MFIs. In other words, supply remained the same while demand increased. Similar studies can be done for removing existing branches.

*Other types of intervention.* Through our model, we can ask more interesting questions such as would an interest rate ceiling be still respected after the removal of certain MFIs from the market? Surprisingly, according to our discussion above, the answer is yes if we were to remove government-owned MFIs. Similarly, we can ask what would happen if a major MFI gets entirely shut down? We can also evaluate effects of subsidies from the donor's perspective (e.g., which MFIs should a donor select and how should the donor distribute its grants among these MFIs in order to achieve some goal).

---

<sup>15</sup><http://www.adb.org/documents/policies/microfinance/microfinance0303.asp?p=microfnc>.

<sup>16</sup><http://www.microfinancegateway.org/p/site/m/template.rc/1.1.10946/>



## 3.6 Conclusion

In this chapter, we studied causal strategic inference in the setting of network-structured microfinance markets. The ultimate objective for modeling microfinance markets, learning the parameters of the model from real world data, and designing algorithms for computing equilibrium points was to study policy-level questions. In fact, one of the major challenges in economic systems is to understand and predict the effects of policy-making without the possibility or capability of evaluating the policy in practice.

Our study of microfinance markets can be extended in several ways. First, the parameters that we learned can be analyzed further using real world evidence. For example, the learned revenue generation function of a village can be analyzed in correlation with its access to resources, such as nearby rivers. Also, the temporal aspects of such a model have been left open here.

Our general approach can be extended to many other systems. These include smart grid systems, which we will discuss a little in the next chapter. We should mention here that a game-theoretic approach to modeling markets and estimating the parameters of the model has been taken before and is being actively pursued now in econometrics. However, one difference between our approach and the econometrics approach is that ours is more of an algorithmic approach than an analytic approach, which allows us to use the classical models of abstract economies that offer little or no analytic insight. On the other hand, most of the game-theoretic models in econometrics are rooted in the discrete choice model [77, 78], which makes an analytic approach more amenable.

# Chapter 4

## Conclusion

In this dissertation, we have studied causal strategic inference in two contrasting application domains. Our first study focused on social networks, where we posed several interesting questions, such as identifying the most influential nodes in a network, as causal strategic inference questions. Our second study was on microfinance markets, where we studied cause and effect questions without the privilege of conducting trial-and-error experiments.

Despite their contrasting application areas, these two studies bear a common signature that can be seen in other domains as well. In a variety of complex systems, this common property is manifested in the strategic interdependence among the entities of the system. That is, the choice of an action by an entity depends on the choice of actions by the other entities. The general framework presented in this dissertation can be applied to many such scenarios. Several future directions in this theme are outlined below.

### **Minimal Targeted Interventions to Reduce Smoking or Obesity**

It is a common perception that many of our behavioral aspects, such as smoking, have an inherent element of social influence in them. In fact, one of the instructions in American Cancer Society's guide to quitting smoking is: "Ask family and friends who still smoke not to smoke around you."<sup>1</sup>

Medical sociologists have long been interested in investigating a scientific connection between social influence and such behavioral aspects. In recent times, some remarkable progress has been made on this issue. Using the Framingham Heart Study<sup>2</sup> data ranging over several decades, Christakis and Fowler showed that there is indeed scientific evidence of social influence stretching be-

---

<sup>1</sup><http://www.cancer.org/healthy/stayawayfromtobacco/guidetoquittingsmoking/guide-to-quitting-smoking-how-to-quit>

<sup>2</sup><http://www.framinghamheartstudy.org>

yond smoking and touching many other phenomena, such as obesity and even happiness [24, 25, 40].

One interesting question in these settings is minimal targeted intervention: Can we reduce the level of smoking or obesity in a community by selecting *a small number* of individuals and persuading them to quit smoking or eat healthy? Identifying such a small set of individuals who can influence others and thereby make a positive impact on the whole network is a natural application of the most influential individuals problem we studied in Chapter 2. The main challenge in this new setting would be computing or counting Nash equilibria in *very large-scale* games.

### **Causal Strategic Inference in House of Representatives**

In our study of influence among the senators in the U.S. Congress, we treated all the legislation issues the same way. However, the type of legislation issue is often an important factor for decision-making, as suggested in political science literature in the context of the U.S. House of Representatives [86]. Developing computational models to complement the political science literature on it is an interesting avenue. Causal strategic inference in this setting would, of course, involve larger games than the ones we considered in Chapter 2.

One possible way of incorporating legislation issues within the influence games framework would be to consider the Bayesian games setting, where legislation issues could be incorporated into the type information of the game. Two interesting avenues in this setting would be to learn such a game model from voting data and to perform causal strategic inference on this model, which can now be conditioned on the type of legislation issue. It would also be interesting to investigate the interplay between scalability (in terms of the number of players) and the types of legislation issues. For instance, would the information about legislation issues help us deal with large games?

### **Using Game Theory in Smart Grids**

With the advent of smart electricity grids, a variety of exciting, new problems have surfaced in AI, many of which have been nicely reported in a recent article by Ramchurn et al. [102]. One of the key problems is demand-side management, where the goal is to achieve a balance between the supply and the demand. A major challenge in achieving this goal is posed by the bidirectional nature of the links in a smart grid. That is, any household buyer may also act as a seller. In this scenario, “developing simulation and prediction tools to allow the systemwide consequences of deploying pricing mechanisms and energy management agents to be assessed by grid operators and suppliers” [102] is a

grand challenge.

Similar to our study of causal strategic inference in microfinance markets in Chapter 3, modeling smart grid markets to predict the outcome of the system with respect to a pricing scheme would be an interesting future direction. Using such a model of smart grid markets, we can answer various policy-level intervention questions.

Another interesting direction would be to investigate how the selfish optimization of each individual's own electricity usage affects the "social" or collective outcomes, such as load-balancing. Formally known as the *price of anarchy*, it would indicate how inefficient a system could be in the worst case. Therefore, it would allow a quantitative comparison of different smart grid systems. Such a study also has the potential to impact future designs of smart grid systems.

Causal strategic inference questions like the above and many others shape the outlook of this dissertation.

# Bibliography

- [1] D. A. Akerberg and G. Gowrisankaran. Quantifying equilibrium network externalities in the ach banking industry. *The RAND journal of economics*, 37(3):738–761, 2006.
- [2] R. Albert and A.-L. Barabási. Statistical mechanics of complex networks. *Rev. Mod. Phys.*, 74(1):47–97, Jan 2002. doi: 10.1103/RevModPhys.74.47.
- [3] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
- [4] A. Augereau, S. Greenstein, and M. Rysman. Coordination versus differentiation in a standards war: 56K modems. *The Rand Journal of Economics*, 37(4):887–909, 2006.
- [5] M. Babaioff, N. Nisan, and E. Pavlov. Mechanisms for a spatially distributed market. In *Proceedings of the 5th ACM conference on Electronic commerce, EC '04*, pages 9–20, New York, NY, USA, 2004. ACM. ISBN 1-58113-771-0. doi: 10.1145/988772.988776. URL <http://doi.acm.org/10.1145/988772.988776>.
- [6] L. Backstrom and J. Kleinberg. Network bucket testing. In *Proceedings of the 20th international conference on World wide web*, pages 615–624. ACM, 2011.
- [7] P. Bajari, H. Hong, and D. Nekipelov. Game theory and econometrics: A survey of some recent research. Technical report, University of Minnesota, Department of Economics.
- [8] P. Bajari, H. Hong, J. Krainer, and D. Nekipelov. Estimating static models of strategic interactions. *Journal of Business & Economic Statistics*, 28(4), 2010.

- [9] P. Bajari, H. Hong, and S. P. Ryan. Identification and estimation of a discrete game of complete information. *Econometrica*, 78(5):1529–1568, 2010.
- [10] M.-F. Balcan, A. Blum, and Y. Mansour. Circumventing the price of anarchy: Leading dynamics to good behavior. *SIAM Journal on Computing*, 42(1):230–264, 2013.
- [11] C. Ballester, A. Calv-Armengol, and Y. Zenou. Who’s who in crime network. wanted the key player. Working Paper Series 617, Research Institute of Industrial Economics, Mar 2004. URL <http://ideas.repec.org/p/hhs/iuiwop/0617.html>.
- [12] C. Ballester, A. Calv-Armengol, and Y. Zenou. Who’s who in networks. wanted: The key player. *Econometrica*, 74(5):1403–1417, 09 2006. URL <http://ideas.repec.org/a/ecm/emetrp/v74y2006i5p1403-1417.html>.
- [13] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson. The diffusion of microfinance. Technical report, National Bureau of Economic Research, 2012.
- [14] R. A. Berk. A gaming approach to crowd behavior. *American Sociological Review*, 39(3):pp. 355–373, 1974. ISSN 00031224. URL <http://www.jstor.org/stable/2094295>.
- [15] S. Berry, J. Levinsohn, and A. Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.
- [16] S. T. Berry. Estimation of a model of entry in the airline industry. *Econometrica: Journal of the Econometric Society*, pages 889–917, 1992.
- [17] P. A. Bjorn and Q. H. Vuong. Simultaneous equations models for dummy endogenous variables: A game theoretic formulation with an application to labor force participation. Technical Report 527, California Institute of Technology, Division of the Humanities and Social Sciences, 1984.
- [18] L. E. Blume, D. Easley, J. Kleinberg, and É. Tardos. Trading networks with price-setting agents. *Games and Economic Behavior*, 67(1):36 – 50, 2009. ISSN 0899-8256. doi: 10.1016/j.geb.2008.12.002. URL <http://www.sciencedirect.com/science/article/pii/S0899825608002182>.

- [19] H. G. Blumer. Collective behavior. In R. E. Park, editor, *Principles of Sociology*, pages 219–288. Barnes & Noble, New York, 1939.
- [20] H. G. Blumer. *Review of Sociology: Analysis of a Decade*. J. Wiley, 1975. URL <http://books.google.com/books?id=ILFUMAEACAAJ>.
- [21] T. F. Bresnahan and P. C. Reiss. Entry in monopoly market. *The Review of Economic Studies*, 57(4):531–553, 1990.
- [22] T. F. Bresnahan and P. C. Reiss. Empirical models of discrete games. *Journal of Econometrics*, 48(1):57–81, 1991.
- [23] J. I. Bulow, J. D. Geanakoplos, and P. D. Klemperer. Multimarket oligopoly: Strategic substitutes and complements. *The Journal of Political Economy*, 93(3):488–511, 1985. ISSN 00223808. URL <http://www.jstor.org/stable/1832005>.
- [24] N. A. Christakis and J. H. Fowler. The Spread of Obesity in a Large Social Network over 32 Years. *N Engl J Med*, 357(4):370–379, 2007. doi: 10.1056/NEJMsa066082. URL <http://content.nejm.org/cgi/content/abstract/357/4/370>.
- [25] N. A. Christakis and J. H. Fowler. The Collective Dynamics of Smoking in a Large Social Network. *N Engl J Med*, 358(21):2249–2258, 2008. doi: 10.1056/NEJMsa0706154. URL <http://content.nejm.org/cgi/content/abstract/358/21/2249>.
- [26] M. S.-Y. Chwe. Structure and strategy in collective action. *American Journal of Sociology*, 105(1):pp. 128–156, 1999. ISSN 00029602. URL <http://www.jstor.org/stable/10.1086/210269>.
- [27] F. Ciliberto and E. Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791–1828, 2009.
- [28] T. Dawber. *The Framingham study: The epidemiology of atherosclerotic disease*. Cambridge: Harvard Univeristy Press, 1980.
- [29] B. de Aghion and J. Morduch. *The economics of microfinance*. MIT Press, 2005.
- [30] G. Debreu. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences*, 38(10):886–893, 1952.

- [31] P. Domingos and M. Richardson. Mining the network value of customers. In *Proceedings of the Seventh International Conference on Knowledge Discovery and Data Mining*, pages 57–66, San Francisco, CA, 2001. ACM Press.
- [32] P. Dubey, O. Haimanko, and A. Zapechelnyuk. Strategic complements and substitutes, and potential games. *Games and Economic Behavior*, 54(1):77 – 94, 2006. ISSN 0899-8256. doi: DOI:10.1016/j.geb.2004.10.007. URL <http://www.sciencedirect.com/science/article/B6WFW-4DVBGS9-1/2/a5aa255deb01ed64d1d2bcf850420a31>.
- [33] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *Annals Math. Stat.*, 30:165–168, 1959.
- [34] E. Even-Dar and A. Shapira. A note on maximizing the spread of influence in social networks. In *Workshop on Internet and Network Economics (WINE)*, 2007.
- [35] E. Even-dar, M. Kearns, and S. Suri. A network formation game for bipartite exchange economies. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*, 2007.
- [36] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC '04: Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 604–612, New York, NY, USA, 2004. ACM. ISBN 1-58113-852-0. doi: <http://doi.acm.org/10.1145/1007352.1007445>.
- [37] I. Fisher. *Mathematical Investigations in the Theory of Value and Prices*. Yale University, 1892. [books.google.com/books?id=djIoAAAAYAAJ](http://books.google.com/books?id=djIoAAAAYAAJ).
- [38] R. A. Fisher. *The Design of Experiments*. Oliver and Boyd, 1935. URL <http://books.google.com/books?id=-EsNAQAAIAAJ>.
- [39] R. A. Fisher. *The design of experiments*. Hafner Publishing Company, 1966. URL <http://books.google.com/books?id=909qAAAAMAAJ>.
- [40] J. H. Fowler and N. A. Christakis. Dynamic spread of happiness in a large social network: longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ*, 337, 12 2008. doi: 10.1136/bmj.a2338.
- [41] J. H. Fowler and N. A. Christakis. Dynamic spread of happiness in a large social network: Longitudinal analysis over 20 years in the Framingham Heart Study. *BMJ*, 337(dec04\_2):a2338–, 2008. doi: 10.1136/



- bmj.a2338. URL [http://www.bmj.com/cgi/content/abstract/337/dec04\\_2/a2338](http://www.bmj.com/cgi/content/abstract/337/dec04_2/a2338).
- [42] D. M. Gale and S. Kariv. Financial networks. *American Economic Review*, 97(2):99–103, 2007. doi: 10.1257/aer.97.2.99. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.97.2.99>.
- [43] A. Gelman, J. Carlin, H. Stern, and D. Rubin. *Bayesian Data Analysis*. Chapman & Hall/CRC, 2003. ISBN 9781584883883. URL <http://books.google.com/books?id=TNYhmkXQSjAC>.
- [44] M. Ghatak and T. Guinnane. The economics of lending with joint liability: Theory and practice. *Journal of Development Economics*, 60: 195–228, 1999.
- [45] E. Goode. *Collective Behavior*. Saunders College Pub., 1992. ISBN 9780155000339. URL <http://books.google.com/books?id=XEsgAAAACAAJ>.
- [46] A. Gopnik, C. Glymour, D. M. Sobel, L. E. Schulz, T. Kushnir, and D. Danks. A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, 111(1):3–32, 2004. doi: 10.1037/0033-295X.111.1.3.
- [47] G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: Hard and easy games. *J. Artif. Int. Res.*, 24(1):357–406, 2005.
- [48] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, pages 424–438, 1969.
- [49] M. Granovetter. Threshold models of collective behavior. *The American Journal of Sociology*, 83(6):1420–1443, 1978. ISSN 00029602. URL <http://www.jstor.org/stable/2778111>.
- [50] M. Granovetter and R. Soong. Threshold models of diffusion and collective behavior. *The Journal of Mathematical Sociology*, 9(3):165–179, 1983. doi: 10.1080/0022250X.1983.9989941. URL <http://www.tandfonline.com/doi/abs/10.1080/0022250X.1983.9989941>.
- [51] G. Heal and H. Kunreuther. You only die once: Managing discrete interdependent risks. Technical report, August 2003. NBER Working Paper No. W9885. Available at SSRN: <http://ssrn.com/abstract=430599>.

- [52] G. Heal and H. Kunreuther. Supermodularity and tipping. Working Paper 12281, National Bureau of Economic Research, June 2006. URL <http://www.nber.org/papers/w12281>.
- [53] G. Heal and H. Kunreuther. Modeling interdependent risks. *Risk Analysis*, 27:621–634, July 2007. doi: 10.1111/j.1539-6924.2007.00904.x. URL <http://dx.doi.org/10.1111/j.1539-6924.2007.00904.x>.
- [54] J. J. Heckman. Dummy endogenous variables in a simultaneous equation system. *Econometrica*, 46(4):931–959, 1978.
- [55] J. Honorio and L. Ortiz. Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *ArXiv e-prints*, June 2012.
- [56] K. Hoover. *Causality in Macroeconomics*. Cambridge University Press, 2001. ISBN 9780521002882. URL <http://books.google.com/books?id=4pGZLX071jkC>.
- [57] J. Ishii. Compatibility, competition, and investment in network industries: Atm networks in the banking industry. *Unpublished working paper*, 2008.
- [58] M. O. Jackson. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA, 2008. ISBN 0691134405, 9780691134406.
- [59] M. O. Jackson. *Social and economic networks*. Princeton University Press, 2010.
- [60] E. B. Janovskaja. Equilibrium situations in multi-matrix games. *Litovsk. Mat. Sb.*, 8:381–384, 1968. ISSN 0132-2818.
- [61] D. S. Johnson. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences*, 9(3):256–278, December 1974.
- [62] S. M. Kakade, M. Kearns, and L. E. Ortiz. Graphical economics. In *In Proceedings of the 17th Annual Conference on Learning Theory (COLT)*, pages 17–32. Springer, 2004.
- [63] S. M. Kakade, M. Kearns, L. E. Ortiz, R. Pemantle, and S. Suri. Economic properties of social networks. In *Advances in Neural Information Processing Systems 17*, pages 633–640. MIT Press, 2005.

- [64] R. Karp. Reducibility among combinatorial problems. In R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, 1972.
- [65] G. Karypis and V. Kumar. Metis - unstructured graph partitioning and sparse matrix ordering system, version 2.0. Technical report, 1995.
- [66] M. Kearns, M. Littman, and S. Singh. Graphical models for game theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, pages 253–260, 2001.
- [67] J. Kleinberg. Cascading behavior in networks: Algorithmic and economic issues. In N. Nisan, T. Roughgarden, Éva Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 24, pages 613–632. Cambridge University Press, 2007.
- [68] R. E. Kranton and D. F. Minehart. A theory of buyer-seller networks. *American Economic Review*, 91(3):485–508, 2001. doi: 10.1257/aer.91.3.485. URL <http://www.aeaweb.org/articles.php?doi=10.1257/aer.91.3.485>.
- [69] H. Kunreuther and G. Heal. Interdependent security. *Journal of Risk and Uncertainty*, 26(2-3):231–249, March 2003. doi: 10.1023/A:1024119208153. URL <http://www.springerlink.com/content/xr72631q5718216q/fulltext.pdf>.
- [70] H. C. Kunreuther and E. O. Michel-Kerjan. Assessing, managing and benefiting from global interdependent risks: The case of terrorism and natural disasters. Technical report, August 2007. Prepared for the 2007 CREATE Symposium. Available at <http://opim.wharton.upenn.edu/risk/library/AssessingRisks-2007.pdf>.
- [71] G. Le Bon. *The Crowd: A Study of the Popular Mind*. Macmillan, 1897. URL <http://books.google.com/books?id=ip03UHMK5w8C>.
- [72] M. E. Levine. Airline competition in deregulated markets: Theory, firm strategy, and public policy. *Yale J. on Reg.*, 4:393–503, 1987.
- [73] D. Little. *On the Reliability of Economic Models: Essays in the Philosophy of Economics*. Recent Economic Thought Series. Kluwer Academic Publishers, 1995. ISBN 9780792394945. URL <http://books.google.com/books?id=NqcY21jzDEwC>.

- [74] D. A. Locher. *Collective Behavior*. Prentice Hall, 2001. ISBN 9780130886682. URL <http://books.google.com/books?id=AAPYAAAAMAAJ>.
- [75] R. D. Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley New York, 1959.
- [76] G. T. Marx and D. McAdam. *Collective behavior and social movements: Process and structure*. Prentice Hall Foundations of Modern Sociology Series. Prentice Hall, 1994. URL <http://books.google.com/books?id=uBjaAAAAIAAJ>.
- [77] D. McFadden. Conditional logit analysis of qualitative choice behavior. In P. Zarembka, editor, *Frontiers of Econometrics*, pages 105–142. Academic Press, New York, NY, 1974.
- [78] D. McFadden. Econometric models of probabilistic choice. In C. Manski and D. McFadden, editors, *Structural analysis of discrete data*, pages 198–272. MIT Press, Cambridge, 1981.
- [79] D. McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, pages 995–1026, 1989.
- [80] R. McKelvey and T. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.
- [81] R. D. McKelvey and A. McLennan. Computation of equilibria in finite games. In *Handbook of Computational Economics*, volume I, pages 87–142. 1996.
- [82] C. McPhail. *The Myth of the Madding Crowd*. Social Institutions and Social Change. A. de Gruyter, 1991. ISBN 9780202303758. URL <http://books.google.com/books?id=-G9RudakgXsC>.
- [83] C. McPhail. *Crowd Behavior*. Oxford, England, 2007.
- [84] P. Milgrom and J. Roberts. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–1277, 1990. ISSN 00129682. URL <http://www.jstor.org/stable/2938316>.
- [85] D. L. Miller. *Introduction to Collective Behavior and Collective Action*. Waveland Press, 2000. ISBN 9781577661054. URL <http://books.google.com/books?id=qDQAAAAACAAJ>.

- [86] W. E. Miller and D. E. Stokes. Constituency influence in congress. *American Political Science Review*, 57:45–56, 2 1963. ISSN 1537-5943. doi: 10.2307/1952717. URL [http://journals.cambridge.org/article\\_S0003055400000113](http://journals.cambridge.org/article_S0003055400000113).
- [87] D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [88] R. L. Montgomery. *Introduction to the Sociology of Missions*. Praeger, 1999. ISBN 9780275966911. URL <http://books.google.com/books?id=C2St9SpgX20C>.
- [89] J. Morduch. The role of subsidies in microfinance: evidence from the grameen bank. *Journal of Development Economics*, 60:229–248, 1999.
- [90] S. Morris. Contagion. *The Review of Economic Studies*, 67(1):57–78, 2000. ISSN 00346527. URL <http://www.jstor.org/stable/2567028>.
- [91] J. F. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950.
- [92] I. of Microfinance (InM). *Bangladesh Microfinance Statistics*. The University Press Limited (UPL), Dhaka 1000, Bangladesh, 2009. URL [www.inm.org.bd/publication.htm](http://www.inm.org.bd/publication.htm).
- [93] L. E. Ortiz and M. Kearns. Nash propagation for loopy graphical games. In *Neural Information Processing Systems (NIPS)*, 2002. URL <http://www.cs.sunysb.edu/~leortiz/papers/nashprop.pdf>.
- [94] R. E. Park and E. W. Burgess. *Introduction to the science of sociology*. The University of Chicago Press, 1921. URL <http://books.google.com/books?id=saZBAAAAIAAJ>.
- [95] J. Pearl. From bayesian networks to causal networks. *Proceedings of the Adaptive Computing and Information Processing Seminar*, pages 25–27, 1994.
- [96] J. Pearl. Graphical models for probabilistic and causal reasoning. *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, 1: 367–389, 1998.
- [97] J. Pearl. *Causality: Models, Reasoning, and Inference*. University Press, 2000. ISBN 9780521773621. URL [http://books.google.com/books?id=wnGU\\_Tsw3BQC](http://books.google.com/books?id=wnGU_Tsw3BQC).

- [98] J. Pearl and T. Verma. A theory of inferred causation. 1991.
- [99] J. Pearl and T. Verma. A statistical semantics for causation. *Statistics and Computing*, 2(2):91–95, 1992.
- [100] J. Pearl and T. Verma. A theory of inferred causation. *Studies in Logic and the Foundations of Mathematics*, 134:789–811, 1995.
- [101] D. Porteous. Competition and microcredit interest rates. *Focus Note*, 33, 2006.
- [102] S. D. Ramchurn, P. Vytelingum, A. Rogers, and N. R. Jennings. Putting the 'smarts' into the smart grid: a grand challenge for artificial intelligence. *Commun. ACM*, 55(4):86–97, Apr. 2012. ISSN 0001-0782. doi: 10.1145/2133806.2133825. URL <http://doi.acm.org/10.1145/2133806.2133825>.
- [103] R. Raz and S. Safra. A sub-constant error-probability low-degree-test, and a sub-constant error-probability PCP characterization of NP. In *Proceedings of the 29th ACM Symposium on Theory of Computing (STOC)*, pages 475–484, 1997.
- [104] T. Roughgarden. *Selfish routing*. PhD thesis, Cornell University, 2002.
- [105] S. J. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Pearson Education, 2003. ISBN 0137903952.
- [106] J. Rust. Estimation of dynamic structural models, problems and prospects: discrete decision processes. In *Advances in econometrics: Sixth world congress*, volume 2, pages 119–170. Cambridge Univ Pr, 1996.
- [107] A. D. Sarma, D. Chakrabarty, and S. Gollapudi. Public advertisement broker markets. In *Proceedings of the 3rd international conference on Internet and network economics*, WINE'07, pages 558–563, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 3-540-77104-2, 978-3-540-77104-3. URL <http://dl.acm.org/citation.cfm?id=1781894.1781967>.
- [108] T. Schelling. *Micromotives and Macrobehavior*. Norton, 1978.
- [109] T. C. Schelling. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2):143 – 186, 1971. ISSN 0022250X. URL <https://libproxy.cc.stonybrook.edu/login?url=http://search.ebscohost.com/login.aspx?direct=true&db=sih&AN=13985129&site=ehost-live&scope=site>.

- [110] P. Schmidt. Constraints on the parameters in simultaneous tobit and probit models. *Structural analysis of discrete data with econometric applications*, pages 422–434, 1981.
- [111] Y. Sharma and D. P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. *Games and Economic Behavior*, 67(1):174–190, 2009.
- [112] C. Tilly. *From mobilization to revolution*. Random House, 1978. ISBN 9780394349411. URL [http://books.google.com/books?id=\\_aYUAQAIAAJ](http://books.google.com/books?id=_aYUAQAIAAJ).
- [113] D. M. Topkis. Equilibrium points in nonzero-sum  $n$ -person submodular games. *SIAM Journal on Control and Optimization*, 17(6):773–787, 1979. doi: 10.1137/0317054. URL <http://link.aip.org/link/?SJC/17/773/1>.
- [114] P. Toulis and E. Kao. Estimation of causal peer influence effects. In *Proceedings of international conference on Machine learning (ICML)*, New York, NY, USA, 2013. ACM.
- [115] R. H. Turner and L. M. Killian. *Collective Behavior*. Prentice-Hall Sociology Series. Prentice-Hall, 1957. URL <http://books.google.com/books?id=qtCVAAAAIAAJ>.
- [116] J. Ugander, L. Backstrom, C. Marlow, and J. Kleinberg. Structural diversity in social contagion. *Proceedings of the National Academy of Sciences*, 109(16):5962–5966, 2012. doi: 10.1073/pnas.1116502109. URL <http://www.pnas.org/content/109/16/5962.abstract>.
- [117] J. Ugander, B. Karrer, L. Backstrom, and J. M. Kleinberg. Graph cluster randomization: network exposure to multiple universes. *CoRR*, abs/1305.6979, 2013.
- [118] V. V. Vazirani. Combinatorial algorithms for market equilibria. In N. Nisan, T. Roughgarden, Éva Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 5, pages 103–134. Cambridge University Press, 2007.
- [119] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.
- [120] D. Wright and D. Alamgir. Microcredit interest rates in bangladesh ‘capping v competition’. *Donors Local Consultative Group on Finance*, March 2004.

- [121] D. K. X. Gine, P. Jakiela and J. Morduch. Microfinance games. Discussion paper 936, Yale University Economic Growth Center, 2006.