Causal Strategic Inference in a Game-Theoretic Model of Multiplayer Networked Microfinance Markets

MOHAMMAD T. IRFAN, Bowdoin College
LUIS E. ORTIZ, University of Michigan - Dearborn

Performing interventions is a major challenge in economic policy-making. We present *causal strategic inference* as a framework for conducting interventions and apply it to large, networked microfinance economies. The basic solution platform consists of modeling a microfinance market as a networked economy, learning the model using single-sample real-world microfinance data, and designing algorithms for various causal questions. For a special case of our model, we show that an equilibrium point always exists and that the equilibrium interest rates are unique. For the general case, we give a constructive proof of the existence of an equilibrium point. Our empirical study is based on microfinance data from Bangladesh and Bolivia, which we use to first learn our models. We show that causal strategic inference can assist policy-makers by evaluating the outcomes of various types of interventions, such as removing a loss-making bank from the market, imposing an interest-rate cap, and subsidizing banks.

CCS Concepts: •**Applied computing** → **Economics;** •**Computing methodologies** → *Artificial intelligence; Machine learning;*

Additional Key Words and Phrases: Game theory and computational economics, causal inference, artificial intelligence, markets, microfinance economy, economic networks

ACM Reference Format:
Mohammad T. Irfan and Luis E. Ortiz. ????. Causal Strategic Inference in a Game-Theoretic Model of Multiplayer Networked Microfinance Markets. *ACM Trans. Econ. Comp.* 0, 0, Article 00 (0000), 60 pages.
DOI: 0000001.0000001

An extended abstract of this paper appeared in the conference proceedings of NIPS 2014, held in Montreal, Quebec [Mohammad T. Irfan and Luis E. Ortiz. Causal Strategic Inference in Networked Microfinance Economies. In, *Advances in Neural Information Processing Systems (NIPS) 27,* Z. Ghahramani, M. Welling, C. Cortes, N.D. Lawrence, and K.Q. Weinberger, Editors, Pages 1161–1169, 2014. Curran Associates, Inc.]. Parts of the work also appears in the first author's Ph.D. Dissertation [Irfan 2013, Ch. 3], completed under the guidance of the second author. Pieces of the work were presented at a game-theory conference without published proceedings [Causal Inference in Game-Theoretic Settings with Applications to Microfinance Markets. In *The 26th International Conference on Game Theory*, part of the *Stony Brook Game Theory Summer Festival 2015*, July 2015]; and in other events and workshops also without published proceedings [A Game-Theoretic Model of Microfinance Markets (Poster). In *New York Computer Science and Economics (NYCE) Day,* September 2011; and Causal Strategic Inference in Networked Microfinance Economies (Poster), In *Workshop on Transactional Machine Learning and E-Commerce, Neural Information Processing Systems (NIPS 2014),* Dec. 12, 2014, Montreal, Quebec, Canada]. The work was supported in part by National Science Foundation's Faculty Early Career Development Program (CAREER) Award IIS-1643006 (transferred from IIS-1054541).
Author's addresses: M. T. Irfan, Department of Computer Science, Bowdoin College; L. E. Ortiz, Department of Computer and Information Science, University of Michigan - Dearborn.

DOI: 0000001.0000001

1. INTRODUCTION

"Money, says the proverb, makes money. When you have got a little, it is often easy to get more. The great difficulty is to get that little." — Adam Smith

Microfinance programs are typically not-for-profit programs with a humanitarian mission of poverty alleviation and women empowerment. We begin with a one-paragraph historical account of microfinance movement, followed by a brief discussion on several distinctive characteristics of the market. The goal is to help the reader appreciate the potential societal impact of this line of work as well as to give the reader an insight into our model.

A chain of unsuspecting events led to today's worldwide microfinance movement. We only need to go as far back as the 1970s to get to its root. The place of its inception was a small village in Bangladesh and the person credited for the idea is Dr. Muhammad Yunus, an Economics professor at that time. 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 experiment 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 [X. Gine and Morduch 2006]. 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."[1]

We next discuss the key characteristics of microfinance markets that will help us understand some of the peculiarities of this system and will also serve as justifications for our modeling decisions.

First, there are two sides in a microfinance market: *microfinance institutions (MFIs)* and borrowers. The MFIs are almost always not-for-profit "banks" that obtain funding from donor agencies at low interest rates and provide loans to mostly poor clients to invest on household projects like husking, weaving, crafting, livestock raising, agricultural cultivation [de Aghion and Morduch 2005]. In Section 2, we model MFIs as not-for-profit humanitarian organizations with the objective of "market clearance." We model borrowers at an aggregate level as non-industrial or non-corporate agents, as evident from their activities. To piece together these two sides, we need to discuss some peculiar aspects of microfinance programs.

A puzzling element in the success of microfinance programs is that while commercial banks dealing with well-off customers struggle to recover loans, MFIs operate without taking any collateral and yet experience very low default rates. (For the benefit of the reader, Appendix A provides a table with a summary of the different acronyms used in this article.) 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. Although we do not model individuals in this work, group lending with joint-liability contract is key

[1] http://www.nobelprize.org/nobel_prizes/peace/laureates/2006/

to our model. In particular, it accounts for the modeling assumption of full repayment of loans (more on it later).

Individual components of microfinance systems have been subjected to various theoretical and empirical studies. For example, the *adverse selection problem* has been analyzed by Ghatak and Guinnane [Ghatak and Guinnane 1999]. 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 [de Aghion and Morduch 2005]. 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 [Ghatak and Guinnane 1999]. An empirical validation of this result can be found in a World Bank study conducted in Peru [X. Gine and Morduch 2006], 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, then like-minded players are paired together. Such an assortative matching induces risk aversion. Although we do not model partnership-formation among individuals, we do assume that assortative matching helps mitigate the adverse selection problem. In other words, we assume that loans end up in "safe hands" who invest these loans in projects that lead to repayment.

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). Clearly, moral hazard is a deterrent to repayment of 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 [Ghatak and Guinnane 1999]. 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 line with these studies, we will assume that assortative matching and joint-liability contracts would mitigate the moral hazard problem. 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 idealistic. However, besides the studies cited above, practical evidence also suggests very high repayment rates. For example, Grameen Bank's loan recovery rate is 99.46% [Institute of Microfinance (InM) 2009].

Several empirical studies investigated whether competition among MFIs has any substantial effect on interest rates. As reported in [Porteous 2006], 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. Furthermore, 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 [Wright and Alamgir 2004]. Although these studies did not mathematically model the system, their findings are useful in our model. For example, in our model, interest rate is not the only factor that attracts borrowers. In particular, borrowers in a village diversify their aggregate loan portfolio. We also address the question of competition vs. interest rates [Porteous 2006], among other questions, *through our model*.

1.1. Brief Overview of Our Model, Approach, and Questions Addressed

We present a two-sided networked economy model, 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. Each village can only interact with the MFIs having branches there. The MFIs want to set their interest rates to achieve *market-clearance*. The objective of each village is to borrow the *maximum amount of diversified loans* (defined in Section 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.

This paper proposes a computational game-theoretic framework grounded in artificial intelligence (AI) and and machine learning within computer science (CS) for studying causality in strategic scenarios that often appear in economic settings. We call this framework *causal strategic inference*. Within the context of this paper, we use this framework to mathematically *model* microfinance markets as a networked system, *learn* the model from real-world data, and design *algorithms* to predict the effects of interventions. After providing a formal mathematical definition of a microfinance market, we study some of its equilibrium and computational properties. We provide algorithms to perform causal strategic inference to answer a variety of questions about the potential consequences of interventions in the system (i.e., counterfactual queries). We answer these based on the resulting *joint predictions* of villages' loan allocations and banks' interest rates after performing such interventions.

1.2. Our Contributions

Mathematically, we model real-world microfinance markets. Computationally, we complement it with an AI approach to learning the model and designing algorithms for it. Our main contributions are the following.

(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 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 [Bulow et al. 1985] is inherent in our model.
(3) We use real-world microfinance data from Bangladesh and Bolivia to learn the respective models. Our learning 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 a loss-making MFI is removed from the market, or if an MFI decides to close some of its branches?

We next give a very brief description of causal strategic inference. For details, we refer the reader to Appendix C. We also refer the reader interested in how our work fits within the existing literature in econometrics, game theory, CS, AI, and machine learning to Appendix D for a thorough comparative study.

1.3. Causal Strategic Inference

A *game* in *non-cooperative* game theory is mathematically defined 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 [von Neumann and Morgenstern 1947; Fudenberg and Tirole 1991]. Here, a joint-action is specified by an action for each player. A central solution concept in non-cooperative game theory is *Nash equilibrium* [Nash 1951]. We can define a *pure-strategy Nash equilibrium (PSNE)* as a joint-action of the players such that every player *plays* the *best response* to the other players' actions simultaneously. [2] Here, a best-response of a player to the the other players' actions is defined by an action that maximizes its payoff 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.

As mentioned in Section 1.1, our goal is to study one type of causal queries known as interventions in game-theoretic settings. Interventions are generally carried out by "surgeries" [Pearl 2000]. We detail the notion of surgeries in game theoretic settings in Appendix C. One common type of surgery we use here is changing the structure of the game. That is, how will a PSNE change if we make certain changes to a game-theoretic model? For example, what would happen to a microfinance market if loss-making government-owned MFIs are removed? How would the market react if subsidies are injected into certain MFIs? Such intervention questions are addressed in Section 6 using our model of microfinance markets.

We study the aforementioned causal questions under the heading of *causal strategic inference*. We should mention here that we learn our model using real-world data only once, before doing any intervention. After performing interventions, we answer causal questions via PSNE computation. Often times, PSNE computation before and after an intervention helps us understand the effect of the intervention.

We review similar studies of causal strategic inference in econometrics in Appendix D. One major difference between our approach and that of many econometrics studies is that the latter approach is in large part analytic, whereas ours is algorithmic. Yes, 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 [Bajari et al. 2010a]. In contrast, the main focus of this paper is an approach rooted in algorithms, computation, AI, and machine learning ideas. Several other differences between econometrics and our approach are described in Appendix D.3.

1.4. Brief Review of Literature in Microfinance and Econometrics

The *effectiveness* of microfinance has been extensively studied in econometrics. Closest to ours in terms of modeling is Salim [2013]'s investigation of whether the objective function of MFIs is to maximize profit or to alleviate poverty. The author uses branch placement by two of the most prominent MFIs in Bangladesh to show that their objective functions are in line with poverty alleviation. The majority of the other studies focuses on the implications and effectiveness of microfinance and its various components. For example, Mahjabeen [2008] studies the welfare and distributional implications of Grameen Bank in Bangladesh. The author shows that Grameen Bank has a positive impact on increasing income and welfare and reducing poverty at a household level. The social welfare claims of microfinance have been constantly under the microscope. Khandker [2005] uses panel data from Bangladesh to show that microfinance indeed helps poverty alleviation. In a more recent study, Becchetti and Castriota [2011] show that the post-tsunami equity injection into a Sri Lankan MFI significantly helped the recovery of clients heavily impacted by the tsunami. In contrast to these econometrics studies, we do not evaluate the effectiveness of microfinance. Instead, we model a microfinance market with the goal of performing interventions.

[2]There is a more general solution concept known as *mixed-strategy Nash equilibrium (MSNE)*, in which each player independently chooses a probability distribution for playing one of its actions so that every player maximizes the *expected* payoff simultaneously [Fudenberg and Tirole 1991].

Going beyond the specific setting of microfinance, econometrics is rich with studies on modeling markets and economies, estimating the model parameters, and studying causality using the models [Heckman and Vytlacil 2007; Abbring and Heckman 2007]. One of the prime differences between the main vein of econometrics studies and this work is that here we have a strictly game-theoretic setting without any element of statistical modeling. [3] We make this distinction clear by comparing and contrasting our model with a strawman structural equation model. We present the strawman model in Section 4.1, learn it using data from Bangladesh in Section 5.3, and perform interventions using it in Section 6.

Another major difference is that unlike most of the econometrics literature, we do not consider the observed data to be an equilibrium. We give details in Section 3.2. In fact, there is some evidence (e.g., an interest rate cap that matches our prediction) that suggests that the observed data may not correspond to an equilibrium (see Section 6).

Within econometrics, market entry decisions are viewed as strategic and are naturally modeled game theoretically. In Appendix D.2, we review the relevant literature, beginning from the early equilibrium model by Bjorn and Vuong [1984] to some of the recent developments, including Ciliberto and Tamer [2009]'s model of airline markets. We also compare and contrast our approach with market entry models in Appendix D.3. In sum, our computational approach is rooted in algorithmic theory, where we do not have any closed-form solutions for the quantities we are interested in. Furthermore, the majority of the work on entry market models uses random utility models [McFadden 1974; 1981], whereas ours is inspired by the classical Arrow-Debreu competitive economy model [Arrow and Debreu 1954a] rooted in applied general equilibrium (AGE) [Kehoe et al. 2005] [4]. Also, as mentioned in the previous paragraph, we do not assume the observed data to be an equilibrium.

Although our model can be classified as a networked economy model within the AGE framework, there are various differences between the AGE literature and our work. First, in most of the AGE literature, including the seminal work of Arrow and Debreu [1954a], prices are homogeneous in the first degree meaning that the constraints do not change if we scale prices by the same factor. As we will see in Section 2.2, that does not hold in our case. This has implications beyond just the proof of equilibrium existence. There are other differences as well, which we outline in Appendix D.4. We also present a strawman AGE model in Section 4.2 and compare and contrast it with our model in Section 6.

1.5. Overview of the Article

We begin by formally presenting our proposed model of microfinance markets, and establishing some of its more analytic properties, in Section 2. In Section 3, we discuss computational aspects of our work for both inference and learning. In Section 4, we design two strawman models and compare them with our model. In Section 5, we present an empirical study of our general approach on two real-world microfinance markets: Bolivia and Bangladesh. In Section 6, we illustrate how our proposed approach may be applied for policy making, using those two real-world markets as particular instances. In Section 7, we conclude the article with a brief summary of our work and a few remarks, including opportunities for future work.

[3] We thank one of the anonymous TEAC reviewers for pointing this.

[4] Quoting Mas-Colell et al. [1995], "From a substantive viewpoint, general equilibrium theory has a more specific meaning: It is a theory of the determination of equilibrium prices and quantities in a system of perfectly competitive markets. This theory is often referred to as the Walrasian theory of markets." Therefore, although our work is inspired by AGE models for determining equilibrium prices and quantities, it is not an AGE model from this substantive viewpoint.

2. OUR 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.

2.1. Preliminary Notation and Definitions

There are n MFIs and m villages. First, for simplicity of presentation, we denote $[n] \equiv \{1, 2, \ldots, n\}$, and similarly for $[m]$. We use the following notation heavily in the remainder of the article.

— $V_i \equiv$ *set of villages where MFI* i *operates*
— $B_j \equiv$ *set of MFIs that operate in village* j
— $T_i \equiv$ *finite total amount of loan available to MFI* i *to be disbursed*
— $r_i \equiv$ *flat interest rate at which MFI* i *gives loans*
— $x_{j,i} \equiv$ *amount of loan borrowed by village* j *from MFI* i

The following function also plays a key role in our model.

Definition 2.1. **(Revenue-Generation Function of a Village)** We denote by

$$g_j(l) \equiv d_j + e_j\, l \tag{1}$$

the *revenue generation function of village* j, which is a function of the *loan amount* $l \geq 0$, and defined in terms of two parameters corresponding to each village j:

(1) the *initial endowment* $d_j > 0$ (i.e., each village has other sources of income [de Aghion and Morduch 2005, Ch. 1.3]) and
(2) the *rate of revenue generation* $e_j \geq 1$.

For each MFI i, we assume that $T_i > 0$. Finally, the villages have a *diversification parameter* $\lambda \geq 0$ that quantifies how much they want their loan portfolios to be diversified.[5] We delay further discussion of the diversification parameter λ until Sections 2.2, 2.3, and 5. But as a preview, we choose λ following a systematic cross-validation procedure. Roughly, we use this parameter as our way to account for, or in some general sense “rationalize,” real-world behavior some may find “irrational” at a glance: i.e., village’s obtaining loans from banks with interest rates higher than the minimum over all those banks (locally) available to them.

We provide a notation legend in Appendix B for the benefit of the reader.

2.2. Problem Formulation

The following is a formal mathematical definition of our microfinance model as a type of two-sided network economy.

Definition 2.2. (Microfinance Market Model Parameters) A *two-sided networked microfinance model*, or a *microfinance market* for short, is defined by the tuple $(n, m, \mathbf{T}, \mathbf{V}, \mathbf{B}, \mathbf{e}, \mathbf{d}, \lambda)$, where $\mathbf{T} \equiv (T_i)_{i \in [n]}$, $\mathbf{V} \equiv (V_i)_{i \in [n]}$, $\mathbf{B} \equiv (B_j)_{j \in [m]}$, $\mathbf{e} \equiv (e_j)_{j \in [m]}$, $\mathbf{d} \equiv (d_j)_{j \in [m]}$, and each component is as defined in the previous subsection (Section 2.1).

[5] For simplicity, we assume that all the villages have the same diversification parameters.

We note that when we apply our model to real-world settings, we will see that the parameters d_j's and e_j's, in contrast to the other input parameters defining our microfinance market model (Definition 2.2), are not explicitly mentioned in the data and therefore, need to be learned from the data. We present the machine learning scheme for that in Section 3.2. We should also note that implicit in the definition is that our model has an underlying *undirected* network structure: that is, we have $j \in V_i \iff i \in B_j$.

We now explain the semantics of our model.

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 it has available, 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 clearance* in economics. The 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 (CSP)* because of the constant objective function.

Definition 2.3. **(Optimization Problem P_M^i of Each MFI i)**

$$\max_{r_i} \quad 1 \tag{2}$$

$$\text{subject to} \quad r_i\left(T_i - \sum_{j \in V_i} x_{j,i}\right) = 0 \tag{3}$$

$$\sum_{j \in V_i} x_{j,i} \leq T_i \tag{4}$$

$$r_i \geq 0$$

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*.

Definition 2.4. **(Optimization Problem P_V^j for Each Village j)**

$$\max_{\mathbf{x}_j = (x_{j,i})_{i \in B_j}} \quad \sum_{i \in B_j} x_{j,i} + \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}} \tag{5}$$

$$\text{subject to} \quad \sum_{i \in B_j} x_{j,i}(1 + r_i - e_j) \leq d_j \tag{6}$$

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

We will call the second term in the objective function given in Equation 5 of Definition 2.4 the *diversification term*. Note that although this term bears a similarity with the well-known entropic term (i.e., Shannon's entropy), it can be mathematically different. This is because the value of the $x_{j,i}$'s can be larger than 1. Although we can think of the diversification parameter λ as exogenous, we will assume that λ is small enough so that the first term in the objective function dominates the second. We return to this last point and provide more details in Section 2.5. We will call the first constraint of P_V^j given in Equation 6 of Definition 2.4 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. (We refer the reader to Definition 2.1.)

For the two-sided market defined by each (local, distributed) optimization problem P_M^i or P_V^j for each MFI i or village j, respectively, we will apply the solution concept of a *(general) equilibrium point* $(\mathbf{r}^*, \mathbf{x}^*)$, which is a *global solution concept* defined here by *(individually optimal) interest rates* r_i^* for each MFI i and vectors $\mathbf{x}_j^* \equiv (x_{j,i}^*)_{i \in B_j}$ of *(individually optimal) loan allocations* for each village j such that the following two conditions hold. First, given the allocations $\mathbf{x}$, each MFI i is individually optimal (i.e., for each individual MFI i, we have that r_i^* is an optimal solution to P_M^i, as given in Definition 2.3, given $(\mathbf{x}_j^*)_{j \in V_i}$). Second, given the interest rates $\mathbf{r}^*$, each village j is individually optimal (i.e., for each individual village j, we have that $\mathbf{x}_j^*$ is an optimal solution to P_V^j, as given in Definition 2.4, given $(r_i^*)_{i \in B_j}$). A formal mathematical definition follows.

Definition 2.5. **(Equilibrium Point of a Microfinance Market)** We call the tuple $(\mathbf{r}^*, \mathbf{x}^*)$ a *general or market equilibrium point*, or simply an *equilibrium*, *of a microfinance market* $(n, m, \mathbf{T}, \mathbf{V}, \mathbf{B}, \mathbf{e}, \mathbf{d}, \lambda)$ if, simultaneously,

(1) (r_i^* is a *solution* of P_M^i for all i) for each MFI $i \in [n]$, we have $r_i^* \geq 0$ and the instantiations of the variables in Equations 3 and 4 based on the equilibrium point satisfy the respective conditions given in those equations; and
(2) ($\mathbf{x}_j^*$ is a *solution* of P_V^j for all j) for each village $j \in [m]$, we have $\mathbf{x}_j^* \geq 0$ (pointwise),
$$\mathbf{x}_j^* \in \arg\max_{\mathbf{x}_j = (x_{j,i})_{i \in B_j}} \sum_{i \in B_j} x_{j,i} + \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}} \ ,$$
and the instantiations of the variables in Equation 6 (budget constraint) based on the equilibrium point satisfy the respective condition given in that equation.

Said differently, no MFI i and no village j has *any incentive to unilaterally deviate* from their equilibrium prescribed by r_i^* and $\mathbf{x}_j^*$, respectively, because none would gain anything *strictly better* given the loan demands $(\mathbf{x}_j^*)_{j \in V_i}$ for all the villages that MFI i serves and interest rates $(r_i^*)_{i \in B_j}$ for all the MFIs serving village j, respectively, at the *equilibrium*.

The parameter T_i for all MFI i and the diversification parameter λ are assumed to be exogenous inputs to this model. The revenue generation functions 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^i (Definition 2.3) does not have any *direct* control over the allocations $\mathbf{x}$ (in the sense that each P_M^i treats $\mathbf{x}$ as exogenous). Similarly, each village-side optimization problem P_V^j (Definition 2.4) 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, as formally stated in Definition 2.5, an equilibrium is specified by $(\mathbf{r}^*, \mathbf{x}^*)$ such that *all* the optimizations $(P_M^i)_{i \in [n]}$ and $(P_V^j)_{j \in [m]}$ are solved *simultaneously* with respect to $(\mathbf{r}^*, \mathbf{x}^*)$. This notion of an equilibrium point is rooted in the classical mathematical-economics literature on markets [Debreu 1952; Arrow and Debreu 1954b].

2.3. Justification of Modeling Aspects

Various aspects of our modeling choice have been inspired by the book of de Aghion and Morduch [2005] and several empirical studies on microfinance systems [Morduch 1999; Wright and Alamgir 2004; Porteous 2006]. We list some of these modeling aspects below. In the discussion, we often refer the reader back to some of our comments in the Introduction, and several concepts presented there.

Objective of MFIs. As noted in Introduction, we model MFIs as not-for-profit organizations. This may seem unusual at first, especially considering MFIs are banks. 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 2005's book. In fact, the book devotes a whole chapter to bust this myth and establish that MFIs very much depend on subsidies for sustainability. We refer the reader to Chapter 9 of de Aghion and Morduch [2005] and to Morduch [1999]. Therefore, our modeling of MFIs as not-for-profit organizations is aligned with their humanitarian goals as well as a considerable amount of significant empirical evidence.

Objective of Villages. As mentioned in Introduction, typical customers of MFIs are low-income people engaged in small household projects such as rice husking, weaving, crafting, etc. A large majority of these customers are women working at home (e.g., Grameen Bank, a leading MFI, has a 95% female customer base) [de Aghion and Morduch 2005]. Clearly, there is a distinction between customers borrowing from an MFI and those borrowing from commercial banks, because many of the latter are for-profit corporations. Therefore, in line with the empirical evidence [de Aghion and Morduch 2005], 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. Clearly, one would expect a village to seek loan only from the MFIs offering the lowest interest rate. However, real-world data in Section 5 reflects that it is often the case that the same village borrows from multiple MFIs at different interest rates. For example, Bagerhat Sadar in Bangladesh borrows from all major MFIs, including BRAC, Grameen Bank, and BRDB, at varying interest rates. In addition, empirical studies suggest that interest rates are not the only determinant for the demands of the villages [Wright and Alamgir 2004; Porteous 2006] (see Introduction for further discussion). In fact, many of the *non-governmental organization (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 [Wright and Alamgir 2004]. 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 [McKelvey and Palfrey 1995], 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 [Luce 1959].

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% [Institute of Microfinance (InM) 2009]. As explained earlier in Introduction, this striking phenomenon is typically attributed to joint liability contracts, which mitigates moral hazard, and to assortative matching, which mitigates adverse selection. In line with such empirical evidence, we assume that the village side completely repays its loan.

2.4. Special Case: No Diversification of Loan Portfolios

It is useful to first study the case of non-diversified loan portfolios, i.e., $\lambda = 0$.

Definition 2.6. We say that a microfinance market $(n, m, \mathbf{T}, \mathbf{V}, \mathbf{B}, \mathbf{e}, \mathbf{d}, \lambda)$ is *without loan-portfolio diversification* if $\lambda = 0$.

We start with this case because, from a purely mathematical and computational perspective, starting our study with the simplest case may allow us to gain some insights into the problem. In this case, the villages simply wish to maximize the amount of loan that they can borrow. Several properties of an equilibrium point (Definition 2.5) can be derived for this special case. We will later exploit these to establish one of our theoretical results. We refer the reader to Appendix E for all proofs. The following property states that at any equilibrium point, the villages will only borrow at the lowest interest rate and the MFIs will not have any excess demand or supply.

PROPERTY 2.7. *At any equilibrium point* $(\mathbf{r}^*, \mathbf{x}^*)$*, every MFI* i*'s supply must match the demand for its loan, i.e.,* $\sum_{j \in V_i} x^*_{j,i} = T_i$ *(Equation 3). Furthermore, every village* j *borrows only from those MFIs* $i \in B_j$ *that offer the lowest interest rate. That is,* $\sum_{i \in B_j, r^*_i = r^*_{m_j}} x^*_{j,i}(1 + r^*_i - e_j) = d_j$ *(Equation 6) for* $m_j \in \operatorname{argmin}_{i \in B_j} r^*_i$*, and* $x^*_{j,k} = 0$ *for any MFI* $k \notin m_j$*.*

We refer the reader to Appendix E.1 for a formal proof of this property.

The following properties are related in that they jointly preclude certain trivial allocations at an equilibrium point. In particular, the next property, proved in Appendix E.3, says that each village will receive some non-zero amount of loan by spending its initial endowment.

PROPERTY 2.8. *At any equilibrium point* $(\mathbf{r}^*, \mathbf{x}^*)$*, each village* j *has non-zero loan demands* $x^*_{j,i}$ *from some MFI* i*: formally, for any village* j*, there exists an MFI* $i \in B_j$ *such that* $x^*_{j,i} > 0$*.*

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. The proof is in Appendix E.2.

PROPERTY 2.9. *At any equilibrium point* $(\mathbf{r}^*, \mathbf{x}^*)$*, for every MFI* i*,* $r^*_i > \max_{j \in V_i} e_j - 1$*.*

The following property, proved in Appendix E.4, is similar to Property 2.8, only that it states the MFI side of the equilibrium situation.

PROPERTY 2.10. *At any equilibrium point* $(\mathbf{r}^*, \mathbf{x}^*)$*, every MFI* i *has some village* $j \in V_i$ *with non-zero loan demand* $x^*_{j,i}$ *from* i*: formally, for any MFI* i*, there exists a village* $j \in V_i$ *such that* $x^*_{j,i} > 0$*.*

2.4.1. Eisenberg-Gale Formulation for the Case of No Diversification of Loan Portfolios and Village-independent Revenue Generating Functions. 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, have the same revenue generation function g_j (Definition 2.1). [6]

Definition 2.11. We say that a microfinance market $(n, m, \mathbf{T}, \mathbf{V}, \mathbf{B}, \mathbf{e}, \mathbf{d}, \lambda)$ has *village-independent revenue generating functions* if $\mathbf{e}$ and $\mathbf{d}$ are such that, for all villages $j \in [m]$, we have $e_j = e$ and $d_j = d$, for some real-valued numbers e and d, independent of j.

We prove that computing a market equilibrium in this case is equivalent to solving the following Eisenberg-Gale convex program [Eisenberg and Gale 1959; Vazirani 2007]. [7]

[6] The reader should not confuse e here for the "natural exponential" constant.

[7] This is not *exactly* the convex program that Eisenberg and Gale defined [Eisenberg and Gale 1959, p. 166], but rather a simple variant adapted to our context.

Definition 2.12. **(Eisenberg-Gale Program P_E for Microfinance Markets without Loan-Portfolio Diversification and with Village-Independent Revenue Generating Functions)**

$$\min_{\mathbf{z}} \quad \sum_{j=1}^{m} - \log \sum_{i \in B_j} z_{j,i} \tag{7}$$

$$\text{subject to} \quad \sum_{j \in V_i} z_{j,i} - T_i \leq 0, 1 \leq i \leq n \tag{8}$$

$$z_{j,i} \geq 0, 1 \leq i \leq n,\ j \in V_i$$

The main idea of the proof is to make a concrete connection between an equilibrium point $(\mathbf{r}^*, \mathbf{x}^*)$ of a microfinance market and the variables $\mathbf{z}$ of P_E (Definition 2.12). In particular, we 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 show that the equilibrium conditions (Definition 2.5) for the aforementioned special case of no diversification of loan-portfolios (Definition 2.6 and village-independent revenue generating functions (Definition 2.11) are equivalent to the *Karush-Kuhn-Tucker (KKT)* conditions of P_E.

The Eisenberg-Gale formulation immediately implies the existence of an equilibrium point, and the uniqueness of the equilibrium interest rates as a corollary. The proof of the following theorem is in Appendix E.5.

THEOREM 2.13. *The special case of microfinance markets with village-independent revenue generating functions and no loan portfolio diversification has an equivalent Eisenberg-Gale formulation given in Definition 2.12.*

The following corollary is an immediate computational consequence of the convex-programming properties of the Eisenberg-Gale formulation.

COROLLARY 2.14. *For the above special case, there exists an equilibrium point with unique interest rates [Eisenberg and Gale 1959] and a combinatorial polynomial-time algorithm to compute it [Vazirani 2007].*

Theorem 2.13 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 [Vazirani 2007] and the Eisenberg-Gale formulation P_E (Definition 2.12) above to verify this. To emphasize the network structure of microfinance markets, the running time of the convex program is polynomial in the number of variables of the program, which in this case is the number of *edges* of the bipartite graph induced by the V_i's (or equivalently, the B_j's).

2.5. Equilibrium Properties of General Case

We are now back to the general case of the problem, formally specified by mathematical programs $(P^i_M)_{i \in [m]}$ (Definition 2.3) and P^j_V (Definition 2.4) in Section 2.2. We begin with an expository discussion on the village objective function (Equation 5), 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 λ 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 given the behavior exhibited in the real-world data in which villages may have demand for loans from banks other than the one with the minimum interest rates among all those locally available to the village, a point we return to in Section 5. We view our approach to dealing with apparent diversification is a simple first step. While we do recognize that there may be several other ways to approach this issue, we did not explore it in this work. For theoretical purposes, it would suffice to assume the following.

ASSUMPTION 2.1. *The diversification parameter λ satisfies the following condition:*

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

where $T_{max} \equiv \max_i T_i$ and, without loss of generality, $T_i > 1$ for all i.

The following equilibrium properties will be used in the next section (Section 3). The proof is in Appendix E.6.

PROPERTY 2.15. *The first constraint of the village-side optimization programs $(P_V^j)_{j\in[m]}$ (Definition 2.4) must be tight at any equilibrium point.*

Let us define $e^i_{max} \equiv \max_{j \in V_i} e_j$ and $d^i_{max} \equiv \max_{j\in V_i} d_j$, where e_j is village j's rate of revenue generation and d_j is j's initial endowment. We can obtain the following bounds on interest rates.

PROPERTY 2.16. *At any equilibrium point, for each MFI i, $e^i_{max} - 1 < r_i^* \le \frac{|V_i| d^i_{max}}{T_i} + e^i_{max} - 1$.*

The proof is in Appendix E.7.

3. COMPUTATIONAL SCHEME

In this section, we will primarily focus on two computational problems:

(1) learning the model from data and
(2) 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.

3.1. Computing an Equilibrium Point

We will now give a constructive proof of the existence of an equilibrium point of a microfinance market $(n, m, \mathbf{T}, \mathbf{V}, \mathbf{B}, \mathbf{e}, \mathbf{d}, \lambda)$, *given* as input, with respect to the programs $(P_M^i)_{i\in[n]}$ (Definition 2.3) and $(P_V^j)_{j\in[m]}$ (Definition 2.4), stated in Section 2.2, formally defining the game-theoretic model. Here we assume that $\lambda > 0$; we discussed the case of $\lambda = 0$ in the previous section (Section 2.4). 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.

As preview of the technical discussion of the algorithm, we note that it might be interpreted as a type of sequential *partial* best-response dynamics. The villages' demands always correspond to *full* best responses to the current interest rates of the *local* banks to which each village has access. The banks' interest rates are typically only *partial* best responses, roughly in the direction of the best-response gradient, given the *local* demands of the villages that each bank serves. While neither sequential best-response dynamics nor best-response-gradient dynamics are guaranteed to converge in general, we show here that our sequence of partial best-response dynamics always converges to an equilibrium point of the microfinance market. (We note that in practice the equilibrium may potentially be approximate because of inherent numerical imprecisions induced by the computer representation of real-valued numbers.) The ensuing technical results will be presented in the context of the outline given in Algorithm 1.

ALGORITHM 1: Outline of Equilibrium Computation

1: For each MFI i, initialize its interest rate r_i to $e^i_{max} - 1$.
2: For each village j, compute its best-response loan demands $\mathbf{x}_j$ for the current interest rates r_i for all banks $i \in B_j$ serving the village j.
3: **repeat**
4:    **for all** MFI i **do**
5:       **while** $T_i \neq \sum_{j \in V_i} x_{j,i}$ **do**
6:          Update interest rate r_i: *intuitively* this operation seeks to set the *local* interest rate in a *direction* that would *reduce the (local) difference* $T_i - \sum_{j \in V_i} x_{j,i}$ between the *local* total supply of loans T_i and the current *local* loan demands $\sum_{j \in V_i} x_{j,i}$ (Details provided in Section 3.1.1).
7:          For each village $j \in V_i$, update its best response $\mathbf{x}_j$ to reflect the update in r_i (Details provided in Section 3.1.2).
8:       **end while**
9:    **end for**
10: **until** no change in r_i occurs for any i

It turns out to be more convenient for the presentation to first characterize the best response of the villages used in Line 7 of Algorithm 1, before specifying the details of how to update r_i in Line 6 at each round of the algorithm.

LEMMA 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 $i \in B_j$:*

$$x^*_{j,i} = \exp\left(\frac{1 - \lambda - \alpha^*_j(1 + r_i - e_j)}{\lambda}\right) \tag{9}$$

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

$$\sum_{i \in B_j} \exp\left(\frac{1 - \lambda - \alpha^*_j(1 + r_i - e_j)}{\lambda}\right)(1 + r_i - e_j) = d_j \ . \tag{10}$$

The proof is in Appendix E.8.

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 α^*_j change in Line 7, for any village $j \in V_i$. Next, we show the direction of these changes.

LEMMA 3.2. **(Strategic Substitutability)** *Whenever r_i strictly increases (decreases) in Line 6, $x_{j,i}$ must strictly decrease (increase) for every village $j \in V_i$ in Line 7 of Algorithm 1.*

The proof is in Appendix E.9.

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

The proof is in Appendix E.10.

In essence, Lemma 3.2 is a result of *strategic substitutability* [Dubey et al. 2006] between the MFI and the village sides, while Lemma 3.3 is a result of *strategic complementarity* [Bulow et al. 1985] 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.

3.1.1. Line 6: MFI's Best Response. By Lemma 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 2.16, can efficiently find the "right" value of r_i that makes supply equal demand for MFI i. For example, in the first iteration of the while loop, in Line 6, the value of 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, the value of 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 *ϵ-equilibrium point*, where the market ϵ-clears (i.e., the absolute value of the difference between supply and demand for each MFI i is below ϵ) and each village plays its ϵ-best response (i.e., it cannot improve its objective function more than ϵ by changing its current response). Having said that, all of our results hold for $\epsilon = 0$.

3.1.2. Line 7: Village's Best Response. We use Lemma 3.1 to compute each village j's best response $x_{j,i}^*$ to MFIs $i \in B_j$. However, Equation 9 requires computation of α_j^*, the solution to Equation 10. We can exploit the convexity of the left-hand side of Equation 10 to design a simple search algorithm to find α_j^* up to a desired numerical accuracy.

Next, we make the following statement about our constructive proof of the existence of an equilibrium point. The proof is in Appendix E.11.

THEOREM 3.4. *There always exists an equilibrium point in a microfinance market specified by all the optimization programs $(P_M^i)_{i \in [n]}$ (Definition 2.3) and $(P_V^j)_{j \in [m]}$ (Definition 2.4).*

We would like to emphasize that Algorithm 1 exploits the network structure among the villages and the banks. In particular, Line 7 acts locally on the edges incident on villages $j \in V_i$. Although the induced bipartite graphs of the microfinance markets we consider in the experiments (Section 5) happen to be relatively dense, there is no reason to expect that this will always be the case in every microfinance system, particularly large ones, either now or in the future.

3.2. Learning the Model from Behavioral Data

We take a machine learning based optimization approach to learn the model from data on interest rates and loan allocations. The inputs to the machine learning problem are

(1) the spatial structure of the market (specified by B_j for each village j, or equivalently, by V_i for each MFI i),
(2) the *observed* loan allocations $\tilde{x}_{j,i}$ for each village j and each MFI $i \in B_j$,
(3) the *observed* interest rates $\tilde{r}_i$, and
(4) the total supply T_i for each 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 model so that *an equilibrium point closely approximates the observed data*. Note that we do not assume that the observed data corresponds 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 behavioral data. Using index i for MFIs and j for villages, the optimization program is formally defined below.

$$\min_{\mathbf{e},\mathbf{d},\mathbf{r}} \sum_i \sum_{j \in V_i} (x^*_{j,i} - \tilde{x}_{j,i})^2 + C \sum_i (r^*_i - \tilde{r}_i)^2 \tag{11}$$

such that

$$\begin{aligned}
&\text{for all } j,\\
&\quad \mathbf{x}^*_j \in \arg\max_{\mathbf{x}_j} \sum_{i \in B_j} x_{j,i} + \lambda \sum_{i \in B_j} x_{j,i} \log \frac{1}{x_{j,i}} \qquad (12)\\
&\qquad\quad \text{s. t.} \sum_{i \in B_j} x_{j,i}(1 + r^*_i - e_j) \le d_j\\
&\qquad\qquad\quad \mathbf{x}_j \ge 0\\
&\quad \mathbf{e}_j \ge 1,\ \mathbf{d}_j \ge 0\\
&\sum_{j \in V_i} x^*_{j,i} = T_i,\ \text{for all } i\\
&r_i \ge e_j - 1,\ \text{for all } i \text{ and all } j \in V_i
\end{aligned}$$

Equation 11 above defines a *nested (bi-level) optimization program* [Bard 2013]. The factor C in the objective function is a constant, and it can be viewed as a weighting parameter to reconcile the differing scales of the two terms in the objective function. To emphasize how the computational representation corresponding to the formulation of the program exploits the network structure of the microfinance market, note that the expressions in the objective functions and the constraints involve only local aspects of each bank or each village (i.e., via the B_j's ot the V_i's).

In the interior optimization program (Equation 12), 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 3.1 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 α_j in a much smaller search space and then apply Equation 9 to compute $\mathbf{x}^*$. We used the interior-point algorithm of MATLAB's large-scale optimization package to solve the above program (Equation 11); at that point, the interior optimization program (Equation 12) would had been solved using the results of Lemma 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 problem 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 order to compare and contrast our approach with common approaches rooted in econometrics and applied general equilibrium (AGE), we next present two elementary models.

4. STRAWMAN MODELS

As detailed in Appendix C, causality has been studied in a variety of disciplines. Close to our study here are structural equation models widely used in econometrics [Heckman and Vytlacil 2007; Abbring and Heckman 2007] and general equilibrium models used in a variety of economic settings [Kehoe et al. 2005] (please see Footnote 4 on the meaning of general equilibrium and why our model is not exactly a general equilibrium model). We next present very simple structural equation and general equilibrium models. The goal is not to devise new research alternatives, but to position our work within well-known causal models for market settings. We further use these two simple models for compare and contrast when we present policy experiments in Section 6.

4.1. Strawman Structural Equation Model

Following is a simple structural equation model for a two-sided microfinance market. Below, all the noise terms ζ and η are respectively independent.

For each MFI i:

$$\begin{aligned} r_i &= \sum_{j \in V_i} w_{j,i} x_{j,i} + b_i + \zeta_i \\ \zeta_i &\sim \mathcal{N}(0, \sigma_i^2) \end{aligned} \tag{13}$$

For each village j and for each MFI $i \in B_j$:

$$\begin{aligned} x_{j,i} &= \sum_{k \in B_j} v_{j,i,k} r_k + a_{j,i} + \eta_{j,i} \\ \eta_{j,i} &\sim \mathcal{N}(0, \nu_{j,i}^2) \end{aligned} \tag{14}$$

We use linear regression to learn $\mathbf{w}$, $\mathbf{b}$, $\mathbf{v}$, and $\mathbf{a}$ above. The challenge with linear regression in our setting is that we only have one sample of data. This is because the loan allocations across the villages and the interest rates of the MFIs together give us one snapshot of the market. [8] As can be readily seen, without multiple samples the problem becomes singular. In order to synthesize multiple samples we introduce Gaussian noise into the data. This helps to the extent of removing the problem of singularity and thereby solving the linear regression. Beyond that, the learned $w_{j,i}$ and $v_{j,i,k}$ parameters hover around 0 and the b_i and $a_{j,i}$ parameters hover around the average (over noisy samples) interest rates and loan allocations, respectively. [9] We further discuss learning using data from Bangladesh in Section 5.3.

Once the $\mathbf{w}$, $\mathbf{b}$, $\mathbf{v}$, and $\mathbf{a}$ parameters have been learned using linear regression, we solve the following optimization using Matlab's optimization package to compute an equilibrium point. For predicting the effect of an intervention, such as closing a bank, we do not re-learn the model; we use the same model learned before in the following optimization to compute a new equilibrium point. We discuss the results in Section 6.

[8] We thank an anonymous TEAC reviewer for pointing out that one could posit distributional assumptions similar to auctions where valuations can be estimated even though bids can be thought of as a snapshot. This is an interesting direction, since we do not consider stochasticity in this work.

[9] Although we leave the details, it can be shown that with enough samples, the $\mathbf{w}$ and $\mathbf{v}$ values will converge to 0 and the $\mathbf{b}$ and $\mathbf{a}$ values will converge to the observed interest rates and loan allocations, respectively.

$$
\begin{aligned}
\min_{\mathbf{r},\mathbf{x}} \sum_i \left(r_i - \left(\sum_{j \in V_i} w_{j,i} x_{j,i} + b_i \right) \right)^2 + \sum_j \sum_{i \in B_j} \left(x_{j,i} - \left(\sum_{k \in B_j} v_{j,i,k} r_k + a_{j,i} \right) \right)^2 \quad (15)\\
\text{subject to} \quad \sum_{j \in V_i} x_{j,i} \leq T_i \\
\mathbf{r} \geq \mathbf{0} \\
\mathbf{x} \geq \mathbf{0}
\end{aligned}
$$

In sum, the simple structural equation model, when learned using noisy versions of a single sample, cannot capture the interdependence between interest rates and loan allocations. This has a direct effect on studying interventions (Section 6). This model, albeit simple, distinguishes our game-theoretic approach from structural equation model approach. It can be shown that an extension of this strawman model using Tinbergen's linear-quadratic structural equation model [Tinbergen 1956; Heckman et al. 2003] gives similar outcomes.

4.2. Strawman Applied General Equilibrium (AGE) Model

We consider another simple model broadly within applied general equilibrium (AGE) models [Kehoe et al. 2005] and more narrowly within the Arrow-Debreu competitive economy setting [Arrow and Debreu 1954a]. Among numerous utility functions traditionally applied, we choose the Cobb-Douglas function (the logarithmic version of it), since it allows each village to obtain loans from all MFIs operating in that village and thereby capturing the diversification of loan allocations that we observe in data. Other widely used utility functions, such as the Leontief function, allows a village to borrow only from the lowest interest rate MFI(s) operating in that village. Next, we present the village and MFI sides of our simple AGE model.

For each village j:

$$
\begin{aligned}
\max_{\mathbf{x}_j} \quad & \sum_{i \in B_j} \tau_{j,i} \ln x_{j,i} \\
\text{subject to} \quad & \sum_{i \in B_j} x_{j,i}(1 + r_i) \leq d_j \qquad (16)\\
& x_{j,i} \geq 0, \quad \forall i \in B_j
\end{aligned}
$$

For each MFI i:

$$
\sum_{j \in V_i} x_{j,i} - T_i = 0 \qquad (17)
$$

In this simple model, $\tau_{j,i}$ is the Cobb-Douglas function parameter corresponding to village j's preference for MFI i. Note that unlike Definition 2.4, this model does not have a revenue generation function. This is because in classical AGE models, such as Arrow-Debreu competitive economy [Arrow and Debreu 1954a] or Fisher market [Fisher 1892], the endowment parameter d_j is fixed and cannot be changed. The MFI side of this model enforces the market clearance condition. In the MFI side, it is implicit that $x_{j,i}$'s are functions of r_i's.

Taking the Lagrangian of the village side (16), using the first oder condition of optimality, and after some algebra to substitute the Lagrangian variables, we get the following best response equation for each village j.

$$x_{j,i} = \frac{\tau_{j,i} d_j}{\left(\sum_{k \in B_j} \tau_{j,k}\right)(1 + r_i)} \tag{18}$$

Note that we are able to express loan allocations $x_{j,i}$ in terms of model parameters $\tau_{j,i}$'s, d_j, and interest rate r_i. Also, $x_{j,i}$ decreases as r_i increases, which captures what we would expect in real world. We next use Equation (18) to substitute the $x_{j,i}$ terms in the MFI side Equation (17).

$$r_i = \frac{1}{T_i} \sum_{j \in V_i} \frac{\tau_{j,i} d_j}{\sum_{k \in B_j} \tau_{j,k}} - 1 \tag{19}$$

In order to learn the model, we first use the observed values of $x_{j,i}$'s and r_i's from the data into the the first constraint of the village side (16) and estimate d_j's. Note that this constraint will be an equality at optimal. We then estimate $\tau_{j,i}$ using Equation (18). For this, note that we for any i, $\tau_{j,i}$ is divided by the same quantity $\sum_{k \in B_j} \tau_{j,k}$. Therefore, for each i, we first estimate $\tau_{j,i}$ as $x_{j,i}(1 + r_i)/d_j$. We then normalize $tau_{j,i}$ for each i. The model learned using data from Bangladesh is discussed in Section 5.3.

It is evident from Equation (19) that r_i is determined by MFI i's total loan T_i, preference parameters $\tau_{j,i}$'s, and endowment of villages d_j's. Interestingly, r_i does not depend on loan allocations x. Furthermore, we can substitute r_i in Equation (18) and express the $x_{j,i}$ terms independent of r_i's. Even though Equations (18) and (19) lead to a clean, analytic solution for an equilibrium point, they have profound implications on interventions. We discuss it further in Section 6.

5. EMPIRICAL STUDY

We base our empirical study on microfinance data of Bangladesh and Bolivia. The reason we chose these two countries is that over time, microfinance programs in these two countries have behaved very differently with respect to competition and interest rates [Porteous 2006].

5.1. Case Study: Bolivia

Here, we discuss our empirical study in the context of Bolivia's microfinance market.

5.1.1. Data. We have obtained microfinance data of Bolivia from several sources, such as ASOFIN,[10] the apex body of MFIs in Bolivia, and the Central Bank of Bolivia.[11] 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.

5.1.2. Learning the Model. Given the exogenous parameter λ, the learning scheme above estimates the parameters e_j and d_j such that an equilibrium point of the game is a *close approximation to the observed behavioral data.* Let us first explain how we choose the exogenous diversification parameter λ.

Fig. 1 shows how the objective function of the optimization program (Equation 11) varies as a function of λ. It shows that for a range of smaller values of λ, the objective

[10] http://www.asofinbolivia.com
[11] http://www.bcb.gob.bo/

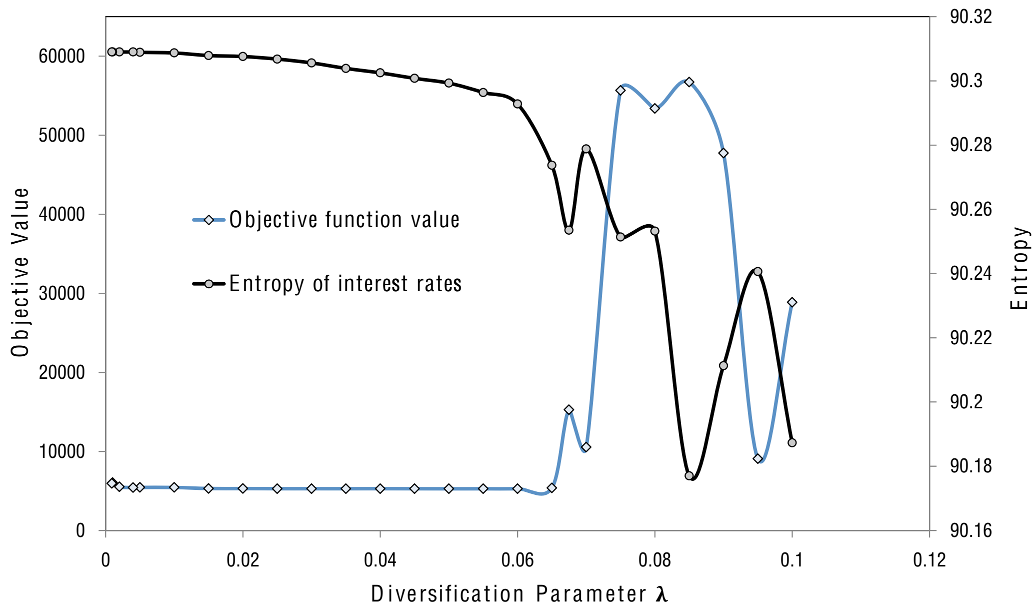

Fig. 1. **Optimal objective function values and entropy of interest rates as λ varies for the Bolivia case.** The plot shows that the objective function value becomes large (which is undesirable) as λ grows.

function value of the learning program is consistently small. (Recall that the optimization routine defined in Equation 11 wants to *minimize* the objective function.) As λ grows, the objective function value oscillates a lot and is sometimes very high.

Fig. 1 also shows how the interest rates become dissimilar as we vary λ. For that, we first define a *entropy* term, $\kappa \sum_i \frac{r_i}{Z} \log \frac{Z}{r_i}$, where $Z = \sum_i r_i$, where κ is a constant set to 100. Here, the entropy quantifies the similarity among the interest rates. That is, higher entropy means the interest rates among the MFIs are more similar. As we can see in the figure, for relatively low values of λ, the interest rates are similar to each other; and as λ becomes too high, they become highly dissimilar at some points.

We chose $\lambda = 0.05$, because at this level of λ, 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 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 larger.

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 $\frac{\text{abs}(x_{j,i} - \tilde{x}_{j,i})}{\sum_i \tilde{x}_{j,i}}$). Fig. 2 shows this. The 45^o line is the locus of equality between these two allocations.

The learned model matches the total loan allocations of the MFIs due to the constraint of the program. As shown later in Fig. 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

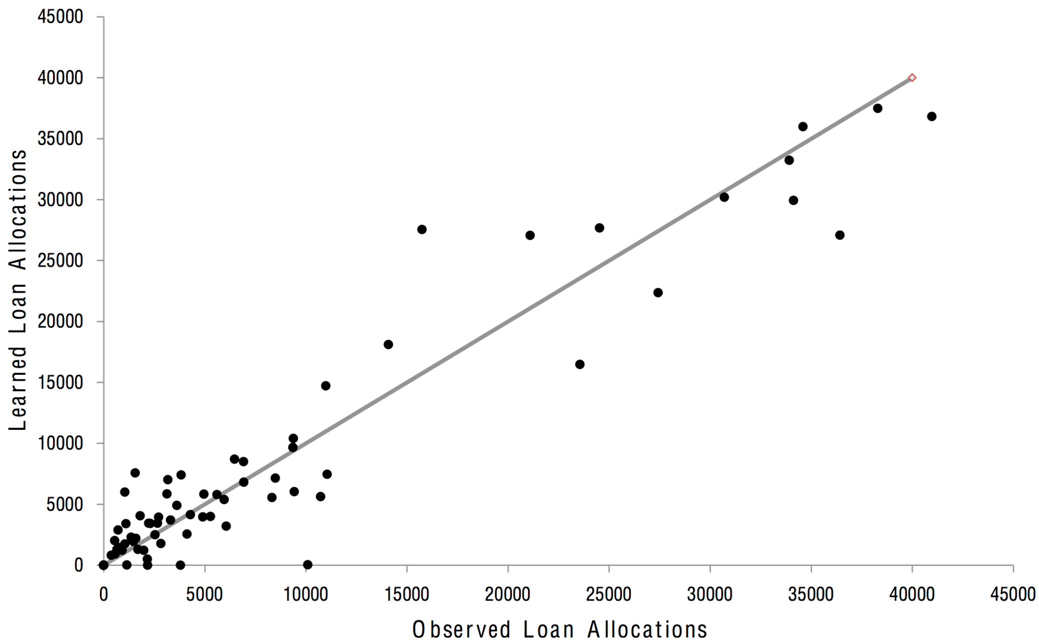

Fig. 2. **Learned allocations vs. observed allocations for the Bolivia dataset.** The plot shows that the learned allocations closely approximate the allocations in the observed data.

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 ν to control the level of noise. For any fixed ν, 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 $x^{\nu}_{j,i}$ and treat the newly constructed noisy dataset as a training set and the observed dataset as test. We learn the model using the training set. We then find an equilibrium allocation $\mathbf{x}^*$ using Algorithm 1. We compute the following *average relative 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 $\tilde{x}_{j,i}$ with $x^{\nu}_{j,i}$. For each noise level ν, we perform the whole procedure a number of times, as we soon describe in the upcoming paragraphs, and calculate the average error.

Gaussian Noise Model. In this model, we obtain $x^{\nu}_{j,i}$ 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 ν 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 Fig. 3, 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.

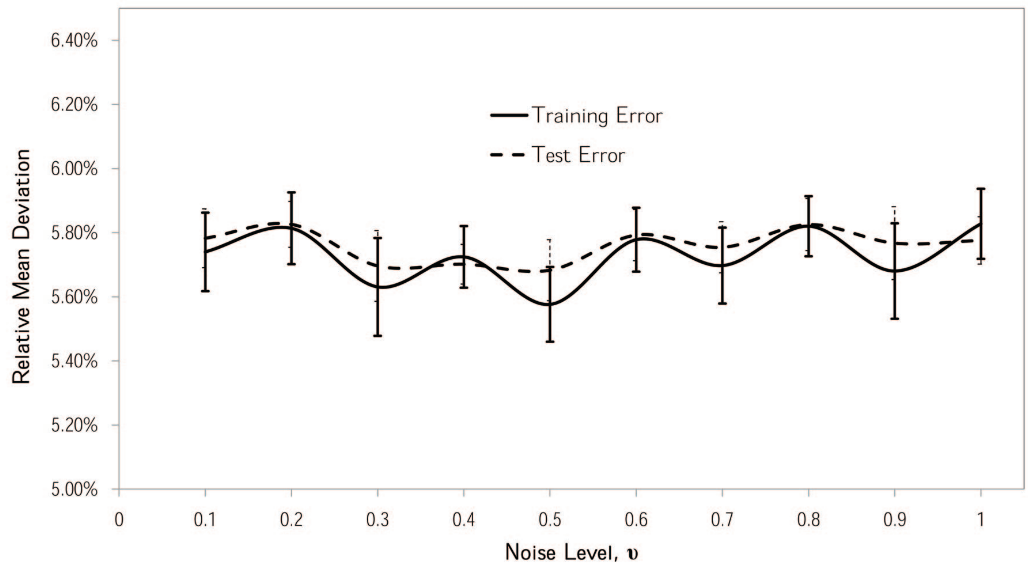

Fig. 3. **Learning curve for the Bolivia dataset under Gaussian random noise.** The learning curve does not suggest overfitting. The vertical bars denote individual 95% CI.

We follow the commonly used procedure of deriving a Dirichlet distribution from a gamma distribution [Gelman et al. 2003, Ch. 18]. We control the noise (i.e., variance) of the Dirichlet distribution using the parameter ν in the following way. For each MFI i, $\mathbf{x}_i^\nu = T(i) \times \mathrm{Dir}(\nu\tilde{\mathbf{x}}_i)$.[12] As the $\nu > 0$ increases, the variance of the distribution $\mathrm{Dir}(\nu\tilde{\mathbf{x}}_i)$ decreases. Varying ν from 2^{-5} (high variance) to 2^{15} (low variance) and taking the average over 50 trials at each ν, 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 Fig. 4, does not suggest overfitting.

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 question, we extend the above experimental procedure using the following bootstrapping scheme.[13]

Suppose that for each noise level ν, 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 ν). For each noise level ν, we iterate the following procedure M times. At each iteration k, we uniformly sample t times (with replacement) 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).

[12] Slightly abusing the notation, the vector $\tilde{\mathbf{x}}_i$ corresponds to non-zero observed allocations only.
[13] In the previous experiment about potential overfitting, the focus was on the distance between an equilibrium point and the data. Now, our focus is on the distance among different equilibrium points when we infuse noise to the data.

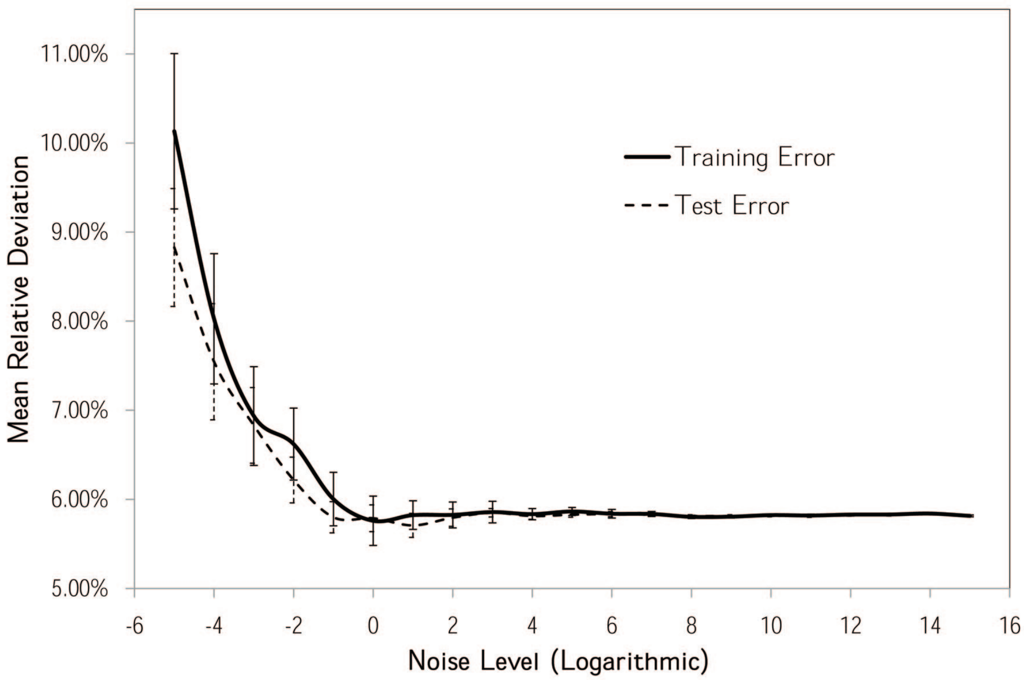

Fig. 4. **Learning curve for the Bolivia dataset under the Dirichlet noise model.** The plot shows the logarithm of the noise parameter ν on the x-axis. (The vertical bars denote individual 95% CI). Just like in Fig. 3, the learning curve in this noise model does not indicate overfitting either.

Within the same k-th iteration, we compute the following *average absolute deviation from the mean*: $\hat{\delta}(k) = \frac{1}{n}\sum_i \frac{1}{|V_i|}\sum_j \frac{1}{t}\sum_l \mathrm{abs}(x_{j,i}^{*(l)}/T_i - \hat{\mu}_{j,i})$. This quantity signifies the average distance of the equilibria of the sampled examples from the mean equilibrium. Now, for each value of ν, we average this distance measure over these M iterations. We perform this bootstrapping procedure for various values of ν.

For the Bolivia dataset, under the Gaussian noise model (described in the previous paragraph) 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. Figs. 5 and 6 show the plots for these two noise models. 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.

5.1.3. 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 equilbrium selection. In practice, we have observed that the equilibrium computed by Algorithm 1 converges to the learned values of x and r, *even if we start with different initial values*. For example, Fig. 7 shows the case of MFI Bancosol's convergence to the same equilibrium interest rate despite different initialization (other

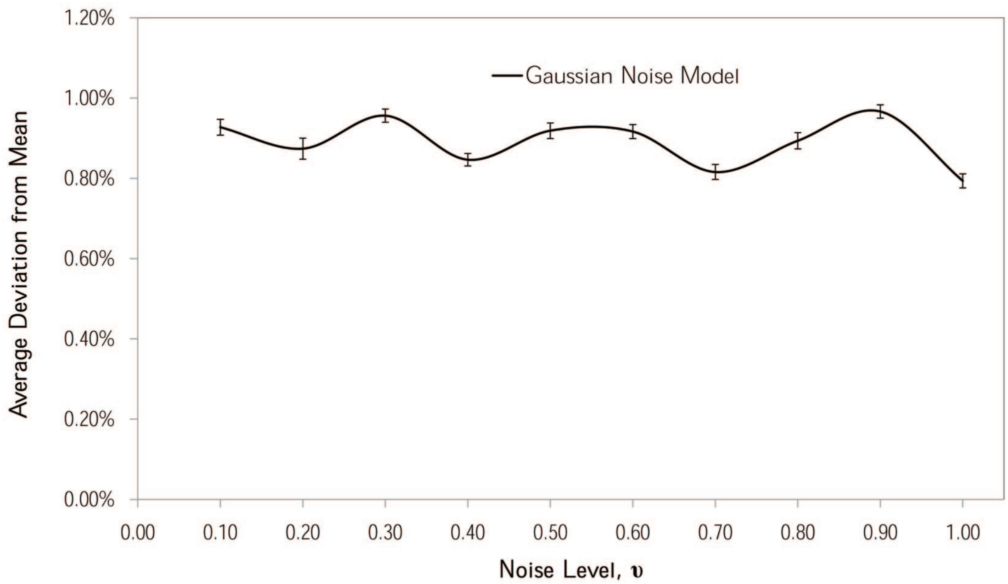

Fig. 5. **Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Gaussian noise model.** The plot shows that the equilibrium point computed is robust with respect to noise. The vertical bars denote individual 95% CI.

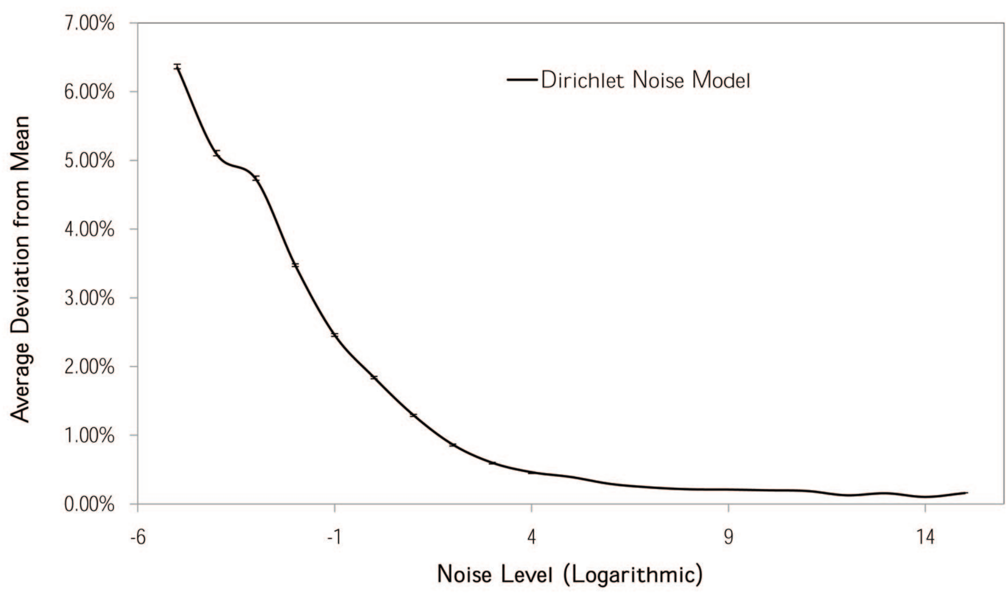

Fig. 6. **Average deviation of the equilibrium points from the mean for the Bolivia dataset under the Dirichlet noise model.** The plot shows the logarithm of the noise parameter ν on the x-axis. The vertical bars denote individual 95% CI. Just like in Fig. 5, the plot also shows the robustness of the computed equilibrium point although the noise model is different—Dirichlet.

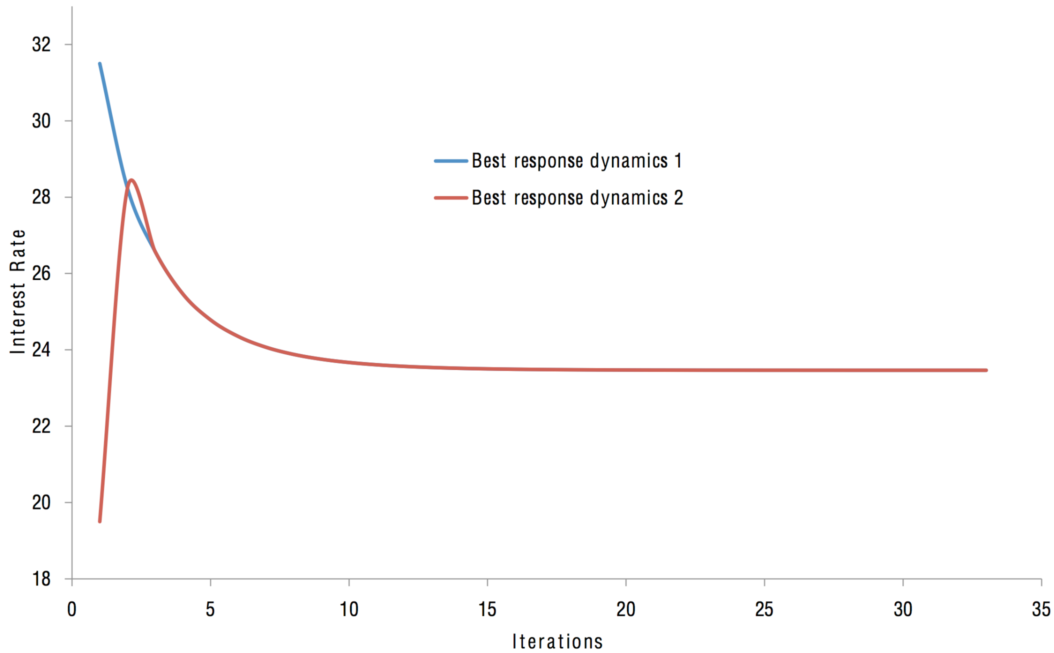

Fig. 7. **Two best-response dynamics of Bolivia's MFI Bancosol with different initialization.** Both of these trajectories converged to the same solution.

interest rates were also differently initialized). This equilibrium interest rate is the same as the learned one. Not only that, as Fig. 8 shows, individual loan allocations were also almost the same.

Finally, Fig. 9 shows a comparison among the observed, learned, and equilibrium interest rates.

5.2. Case Study: Bangladesh

Here, we discuss our empirical study on Bangladesh's microfinance market.

5.2.1. 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 collections of MFIs) operating in 464 *upazillas* or collections 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.

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%) [Porteous 2006; Wright and Alamgir 2004]. [14]

5.2.2. Learning the Model. Because of the size of Bangladesh data, we faced the problem of solving a nonlinear optimization problem of the order of thousands of variables

[14]Quoting Wright and Alamgir [2004] on the interest rates of the government-owned MFIs: "These vary from 10% charged [...] to 16% charged by the Bangladesh Academy for Rural Development (BARD), and 15-20% charged by different BRDB programmes." We take these rates as APRs, which are roughly double of flat rates.

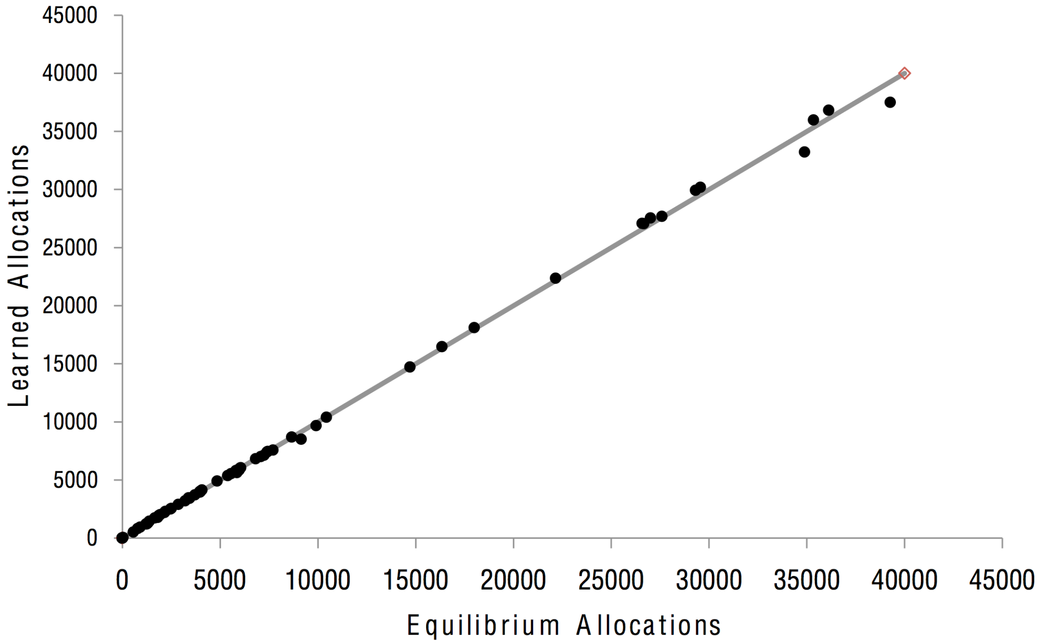

Fig. 8. **Learned allocations vs. equilibrium allocations.** The fact that two allocations are almost identical provides empirical evidence that the learning algorithm was able to capture an equilibrium point using the inner part of the nested optimization program (Equation 12).

and constraints. As discussed above, the interior point algorithm is initialized with a feasible solution, which makes computation a lot faster. Still, solving the problem takes time on the order of hours on a modern quad-core desktop computer, 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) are non-uniform across the villages j. The variation is more pronounced for the d_j values, whereas the e_j values are 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 10. For example, the average relative deviation is only 5.54% when $\lambda = 0.05$. The market clears in the learned model, and as Figure 11 shows, 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 sustainability without subsidies [de Aghion and Morduch 2005; Wright and Alamgir 2004].

5.2.3. 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. Fig. 12 shows the similarity between learned and equilibrium allocations in this case.

5.3. Strawman Models Using Bangladesh Data

For the simple structural equation model presented in Section 4.1, we use standard deviations of 0.003 for interest rates and 0.03 for loan allocations to generate $10,000$ noisy samples (Gaussian) from the single Bangladesh data sample that we have. We

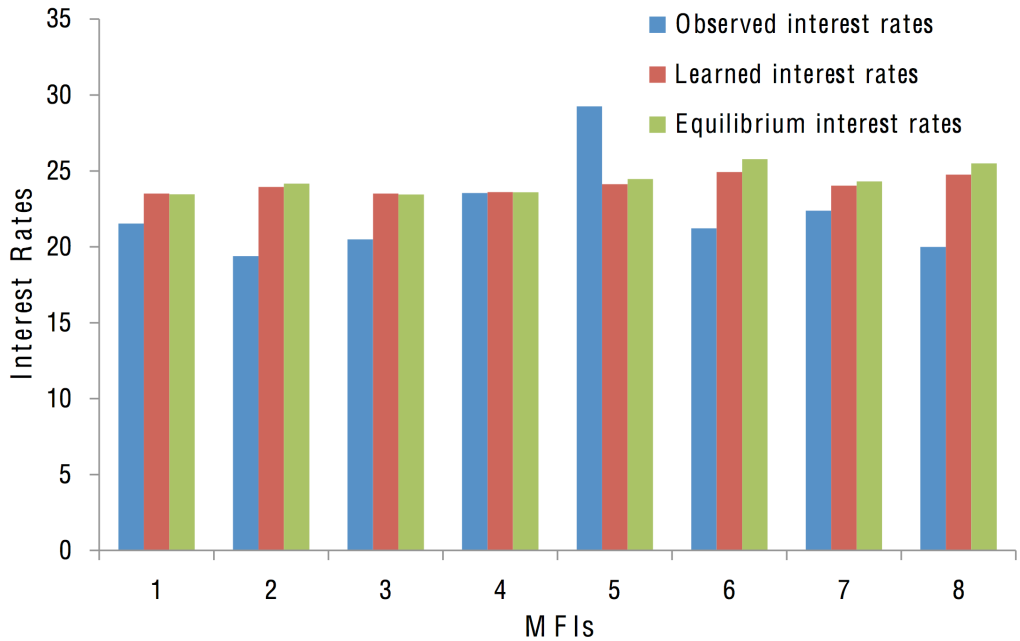

Fig. 9. **Comparison among observed, learned, and equilibrium interest rates for the Bolivia dataset.** As the bar-plot shows, the equilibrium interest rates and the observed interest rates do not completely match, which is ok, because for our model and analysis, we do not assume the data corresponds to an equilibrium point.

then use these samples to perform linear regression on Equations (13) and (14) to learn w, b, v, and a. As mentioned in Section 4.1, the learned $w_{j,i}$ and $v_{j,i,k}$ parameters are very close to 0 and the b_i and $a_{j,i}$ parameters are very close to the observed interest rates and loan allocations, respectively. After the model is learned, we compute an equilibrium point by solving the optimization problem stated in Section 4.1. We discuss our results in the next section.

In contrast to the structural equation model, the AGE model presented in Section 4.2 is computationally much less burdensome. This is largely due to the analytical solutions given by Equations (18) and (19). Once we learned the model using Bangladesh data, we found that the estimated preference parameters $\tau_{j,i}$'s reflect the demand for loans from an MFI. For example, the average τ is highest for Grameen Bank (0.2281) and lowest for government MFIs (0.0154). In the next section, we discuss the inadequacies of this strawman AGE model for examining interventions. [15]

6. POLICY EXPERIMENTS

Our end goal of modeling a 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 model and then compute an equilibrium point, both in the original setting (before removal of any MFI). Using the model learned earlier, we compute a new equilibrium point after the removal of the government-owned MFIs.

[15] We have also learned the AGE model using Bolivia data for some of the interventions.

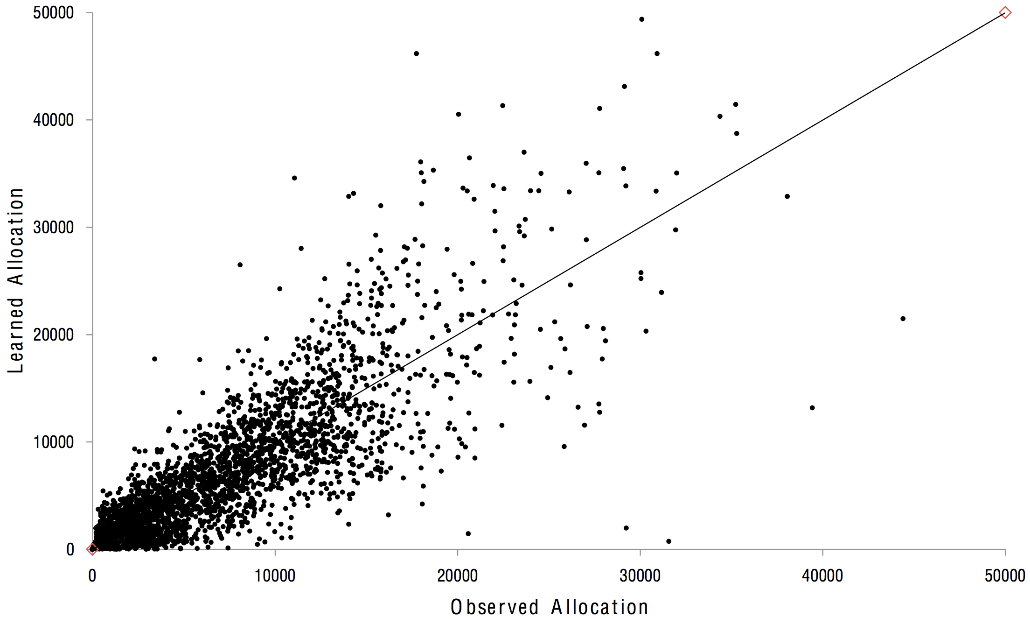

Fig. 10. **Learned allocations vs. observed allocations for the Bangladesh dataset.** Although the plot shows that the learned and observed allocations are not *exactly* the same, the learned allocations do approximate the observed ones.

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. We also draw a distinction between our game-theoretic approach and the traditional structural equation and applied general equilibrium (AGE) approaches using the two strawman models that we presented in Section 4.

6.1. The Role of Subsidies

It has been well documented that the sustainability of MFIs very much depends on subsidies [de Aghion and Morduch 2005; Morduch 1999]. 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 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. For instance, in the Bolivia case, one of the MFIs named Eco Futuro (MFI 5) exhibits very high interest rates both in observed data and equilbrium. A quick look at 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 (MFI 1). As Fig. 13 shows, 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.

To compare with the strawman models, note again that we do not re-learn the model after subsidizing an MFI. For the strawman structural equation model (Section 4.1), this means that the w and v parameters remain close to 0 and the b and a parameters remain close to the observed interest rates and loan allocations, respectively (as they

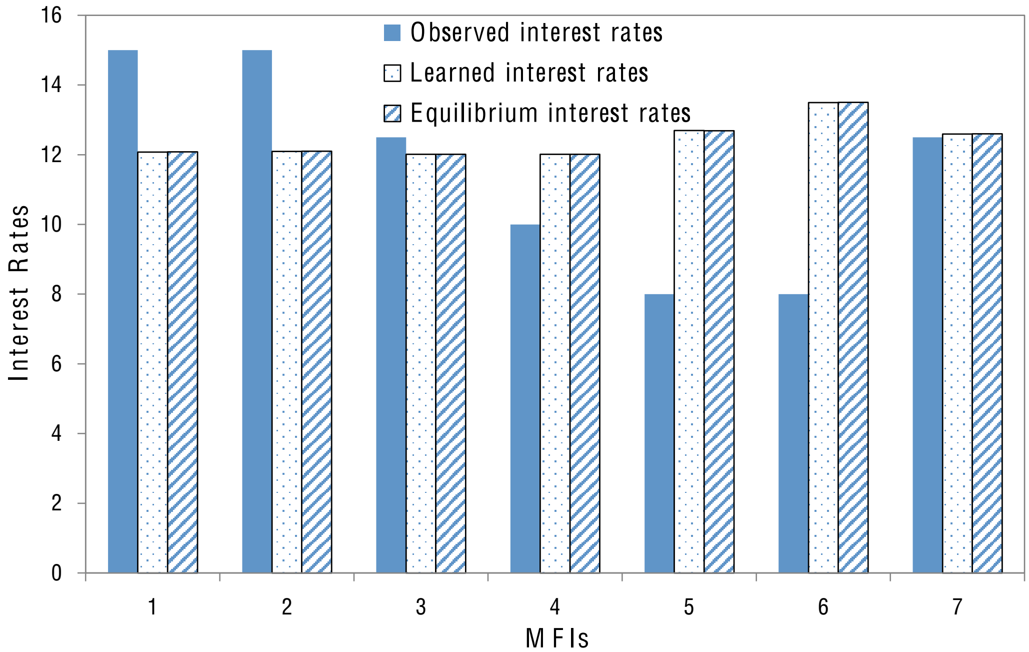

Fig. 11. **Comparison among observed, learned, and equilibrium interest rates for the Bangladesh dataset.** Again, we do not assume that the data corresponds to an equilibrium point. Therefore, the difference between observed interest rates and equilibrium interest rates suggests an interesting aspect. That is, some banks may be charging less than they should and some may be charging more than they should. The first aspect is well-established in the literature with respect to government-owned banks [de Aghion and Morduch 2005]. There is also practical evidence for the second aspect [Porteous 2006], which we will elaborate when we study interest-rate caps in Section 6.3.

were before the intervention). As a result, when we compute an equilibrium point by increasing T_i for the subsidized MFI and then solving the optimization problem (15) in Section 4.1, the equilibrium allocations and interest rates do not change in response to this subsidization. To see exactly why they do not change, note that the optimal values of $\mathbf{x_{j,i}}$ and r_i in (15) are close to $a_{j,i}$ and b_i, respectively, while w's and v's are close to 0, even when we increase T_i.

Using the strawman AGE model (Section 4.2), increasing T_i does decrease r_i and increase $x_{j,i}$ in accordance with Equations (19) and (18). However, when we subsidize Eco Futuro (MFI 5) and make its total loan equal to Banco Sol's, $1/T_5$ becomes so low that r_5 actually becomes an invalid, negative number (-63.13%). Improving this strawman model to account for this is out of scope here. Furthermore, unlike our model, subsidizing Eco Futuro does not have any effect on the interest rates of the other MFIs.

6.2. 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% [Bagazonzya et al. 2010; Porteous 2006], which is close to what our model predicts at an equilbrium point, as Fig. 11 shows.

In contrast, the equilibrium interest rates computed by the strawman structural equation model (i.e., solution to the optimization problem (15)) are almost the same as the observed interest rates. The same is true for the strawman AGE model.

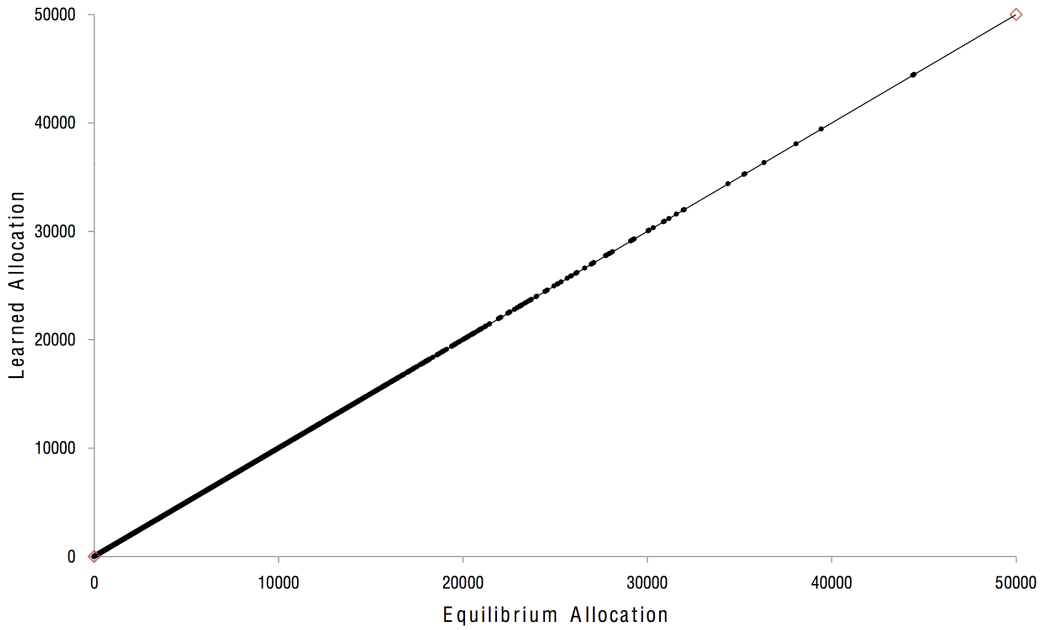

Fig. 12. **Learned allocations vs. equilibrium allocations for the Bangladesh dataset.** The plot shows that the two types of allocations are very similar.

6.3. Cap on Interest-rates

PKSF recently capped the interest rates of its partner organization to 12.5% [Porteous 2006], and more recently, the country's Microfinance Regulatory Authority (MRA) has also imposed a ceiling on interest rate at around 13.5% flat.[16] Such evidence on interest rate ceiling is consistent with the outcome of our model. As Fig. 11 shows, our model predicts that the highest interest rate at an equilibrium point is 13.4975%.

In contrast, as mentioned in Section 6.2, the strawman structural equation model predicts equilibrium interest rates to be the same as the observed interest rates. When we perform an intervention on the structural equation model by setting an interest rate ceiling of 13.5%, we find that equilibrium interest rates for BRAC and ASA go down from 15% to 13.5%, but unsurprisingly this does not affect the interest rate of any other MFI whose rates were already below 13.5%. Again, due to the learned parameters w and v being close to 0, the interdependence among the interest rates (and allocations) is lost. Similarly, due to Equation (19) of the AGE model, lowering the interest rate of one MFI does not affect the rates of the other MFIs.

6.4. Shutting Down Governement-Owned MFIs

It is well-documented in the literature that many of the government-owned MFIs do not have the goal of meeting their operating costs [Wright and Alamgir 2004]. 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 [Porteous 2006].

[16] https://www.ft.com/content/fd16a1f0-ecea-11df-9912-00144feab49a. Alternative: https://goo.gl/j4JJgZ.

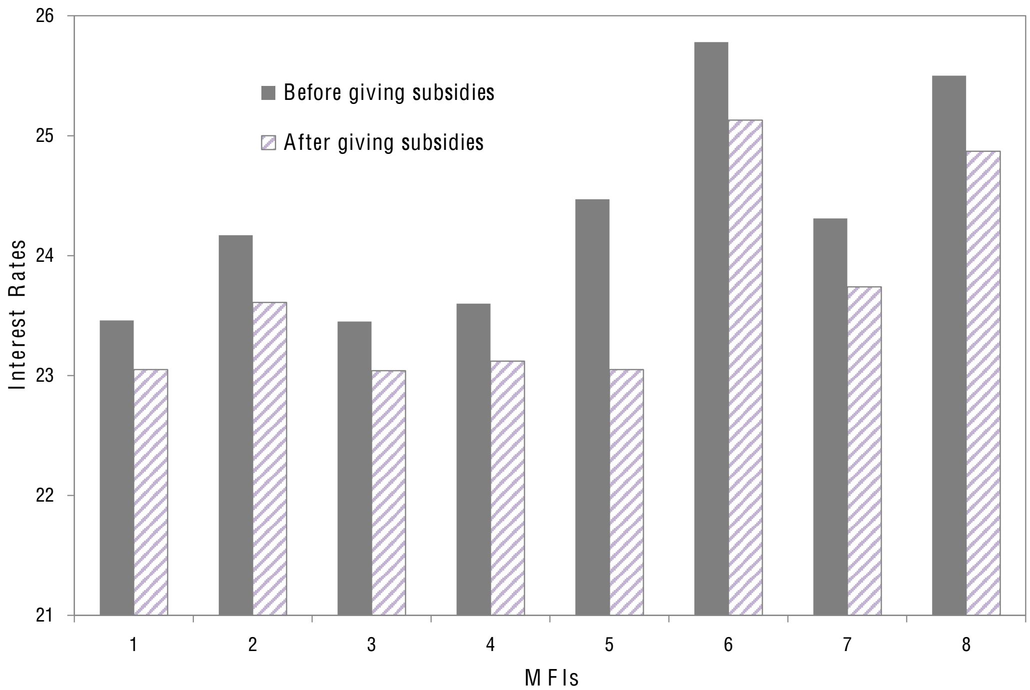

Fig. 13. **Equilibrium interest rates before and after adding subsidies to Bolivia's Eco Futuro (MFI 5).** As the plot suggests, subsidies help bring down the interest rate of not only Eco Futuro but also the other MFIs.

In contrast, the strawman structural equation model does not predict any changes in interest rates when an MFI is shut down (for the same reasons mentioned above). However, the strawman AGE model does predict an increase in the interest rates of the other MFIs, although it predicts very sharp rises in interest rates. When the government-owned MFIs in Bangladesh are shut down, the AGE model predicts that the other MFIs will increase their rates by a value between 7.57% and 9.23% (e.g., BRAC's interest rate will increase from 15% to 24.23%).

6.5. The Effect of 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 Fassil (MFI 7) in Bolivia expands its business to all villages (by adding six new branches). It may seem at first that because of 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.

These effects are not predicted by either of the strawman models. In particular, the strawman structural equation model predicts no change in interest rates. On the other hand, the strawman AGE model predicts that the interest rate of Fassil will go up very sharply by 25.37% (from 22.38% to 47.75%) while those of the other MFIs will go down by a value between 0.69% and 2.63%. This is opposite to our model's prediction that every MFI's interest rate will go up due to supply remaining the same amid increasing demand.

6.6. Other Types of Interventions

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).

7. DISCUSSION AND CONCLUDING REMARKS

In this paper, we studied causal strategic inference in the setting of network-structured microfinance markets. The ultimate objective for modeling microfinance markets, learning 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.

There are several aspects of our model that require further research. For example, we would like to better understand how the effect of the constant parameter C in our learning formulation, which modulates the relative importance of how close the *observed* allocations and interest rates should be to their closest equilibrium in the learned model, on the final results, and how we may be able to use cross-validation to select the "right" value. Similarly, it would be interesting to explore the use of cross validation to select the diversification parameter, or even individual diversification parameters for each village. Another research direction is the exploration of alternative ways to account for the apparent diversification we observe on the real-world data. It would also be interesting to explore non-linear revenue generating functions (e.g., by considering revenue generating functions exhibiting diminishing returns). A systematic theoretical or empirical study of the sensitivity of our learned models and (causal) predictions on different modeling choices and assumptions is also an interesting direction.

Although the Cobb-Douglas utility function for the village side optimization (16) looked promising, the simple strawman AGE model does not capture the interdependencies among the villages and MFIs. A closer look at the AGE utility function of the villages reveals that its derivate is proportional to $\frac{1}{x_{j,i}}$, whereas in our model the derivative is proportional to $\log \frac{1}{x_{j,i}}$, which is much less drastic. In the strawman AGE, this leads to r_i being independent of $x_{j,i}$'s in Equation (19). It is an interesting future direction to investigate how one can adapt standard utility functions like Cobb-Douglas to rectify this and capture the interdependencies among the villages and MFIs. The preference parameters $\tau_{j,i}$'s could also be made specific to MFIs only (i.e., τ_i), which would require more sophisticated learning schemes.

Along the same line, the objective function of the villages can be investigated further. We have modeled the villages as non-corporate agents that maximize diversified loans subject to repayment. There could be other alternatives to modeling the village side that would still maintain the non-corporate nature of the villages. In fact, the diversification parameter λ abstracts away potentially multiple factors behind diversification. As a result, the same diversification approach may not work in all microfinance markets. Modeling the details of diversification is an interesting future direction. Furthermore, in this work, we have modeled the market with villages and MFIs as the basic units. It is an exciting future direction to consider finer levels of granularity, such as households. Going one step further, we have not modeled donor agencies that typically subsidize MFIs. Future work can incorporate this important part of a microfinance system. There are several other directions for extending this work. First, the model 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 other systems as well, e.g., smart grid systems. We should mention here that a game-theoretic approach to modeling markets and performing learning and inference based on 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. This allows us to use the classical models of abstract economies that offer little or no analytic insight. In sum, an algorithmic approach like ours can help with modeling complex systems with nonlinear utility functions.

APPENDIX

A. TABLE OF ACRONYMS

The following table provides a summary of the different acronyms used throughout the article, for the benefit of the reader.

Acronym	*Stands for*
AI	Artificial Intelligence
ACH	Automated Clearinghouse
ATM	Automated Teller Machine
CS	Computer Science
CSP	Constrained Satisfaction Problem
E.U.	European Union
i.i.d.	independent identically distributed
ISP(s)	Internet Service Provider(s)
KKT	Karush-Kuhn-Tucker
MFI(s)	Microfinance Institution(s)
MSNE	Mixed-strategy Nash Equilibrium (or Equilibria, when clear from context)
NGO	Non-governmental Organization
PSNE	Pure-strategy Nash Equilibrium (or Equilibria, when clear from context)
U.S.A. or U.S.	United States of America

B. NOTATION LEGEND

The following table provides a legend of the different notation used throughout the article, for the benefit of the reader.

Symbol	*Brief Description*
$[c] \equiv \{1, 2, \ldots, c\}$	given positive integer c, set of natural numbers up to c
n	number of MFIs
m	number of villages
i	Index MFIs
$V_i \subset [m]$	Set of *villages* where MFI i operates
j	Index villages
$B_j \subset [n]$	Set of *MFIs* that operate on village j
$T_i > 0$	Finite total amount of loan available to MFI i to be disbursed
$l \in \mathbb{R}^+ \cup \{0\}$	a loan amount
$d_j > 0$	initial endowment of village j
$e_j \geq 1$	rate of revenue generation of village j
$g_j : \mathbb{R}^+ \cup \{0\} \to \mathbb{R}^+$	revenue generating function of village j (here assumed linear in l, with slope e_j and offset d_j)
$r_i > 0$	flat interest rate at which MFI i lends
$x_{j,i} \geq 0$	amount of loan that village j borrows from MFI i
λ	villages' diversification parameter for loan portfolios
$\mathbf{r} \equiv (r_i)_{i \in [n]}$	vector of all interest rates for all MFIs
$\mathbf{x}_j \equiv (x_{j,i})_{i \in B_j}$	total amount of loan that village j borrows over all MFIs in B_j
P_M^i	individual optimization program for MFI i
P_V^j	individual optimization program for village j
P_E	Eisenberg-Gale optimization-program formulation
$(\mathbf{r}^*, \mathbf{x}^*)$	equilibrium point corresponding to individually, locally optimal interest rates and allocations

C. 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 [Little 1995; Pearl 2000; Hoover 2001]. 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 importance and wide applicability. This paper presents a framework for studying causality within networked microfinance systems. This may also be applicable to other strategic scenarios. For simplicity, 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 [Little 1995]. 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 [Granger 1969]. 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 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 [Little 1995]. 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 describe this notion of interventions in strategic settings in Section C.2 and later show how we apply it in Section 6. We next draw a contrast between causal probabilistic inference and causal strategic inference.

C.1. 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 [Pearl and Verma 1991; 1992; Pearl 1994; Pearl and Verma 1995; Pearl 1998; 2000]. Applications of causal probabilistic inference can now be seen in very diverse disciplines, such as economics, public policy, sociology, CS, 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 [Fisher 1935] was published [Pearl 2000, 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 [Pearl 2000, 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 observational inferences. 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 $do(X = x)$ [Pearl 2000, 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 original form, 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 Pearl, the study of causality can be hierarchically organized in three natural types of queries with increasing levels of difficulty: predictions, interventions, and counterfactuals [Pearl 2000, 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 paper is to study causal inferences in game-theoretic settings at the second level of queries, interventions, from an AI perspective. Because we believe that game theory provides a reasonably reliable mathematical framework in which to encode strategic interactions among a set of distributed decision-makers (i.e.,"players") in AI, we refer to this type of inference *causal strategic inference*.

C.2. Causal Strategic Inference

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.

C.2.1. 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 [Kleinberg 2007], 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 [Sharma and Williamson 2009]. One motivation is the case where there are two types of users of a network application: *premium* users, who have the privilege of choosing their own route of traffic, and *ordinary* users, who must go by whatever the network administrator has chosen for them [Roughgarden 2002]. 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 the point of the previous paragraph, 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.

C.2.2. 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 this paper, 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 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 model again, because we do not have any data under those "hypothetical" conditions.) The difference in equilibrium outcomes before and after the removal of the banks gives us the desired answer.

D. 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 an AI-based framework for studying causality within networked microfinance systems, which may be more broadly applicable to handle 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: (1) mathematically *modeling* a complex system by exploiting advances in compact representation from the AI literature, (2) *learning* models from *real-world data* using machine learning techniques, and (3) designing *algorithms* to predict the effects of interventions based on advances in computational game theory. We will now review some earlier research relevant to these components, especially from economics literature.

D.1. Causal Strategic Inference in Econometrics

Modeling strategic scenarios and learning 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 the econometrics literature, a major application scenario for studying strategic decision-making is the setting where two or more firms simultaneously decide whether to enter a market or not [Aguirregabiria and Suzuki 2015]. 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 [McFadden 1974; 1981]. 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*.

D.2. Review of Related Work in Econometrics

In this section of the appendix, we provide a reasonably comprehensive, detailed review of specific work related to causal strategic inference in econometrics.

D.2.1. Bjorn and Vuong's Model of Labor Force Participation. The first discrete game model is attributed to Bjorn and Vuong [1984], who studied the case of simultaneous decision-making by a husband and a wife on whether to enter the labor force or not. 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 1974's *random utility model* as follows.

$$\widetilde{u}_i(x_i, x_{-i}) = u_i(x_i, x_{-i}) + \eta_i(x_i, x_{-i}) \tag{20}$$

The above payoff function (Equation 20) consists of an *observed, deterministic* part u_i and an *unobserved (to the researcher), random* part η_i. The first accounts for observed attributes of the players, such as age, education level, assets, and number of kids, among others. 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 η_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 particular model of Bjorn and Vuong is one of complete information. Here, the best response x_i^* of player i can be written as follows.

$$x_i^* = 1 \iff \widetilde{u}_i(1, x_{-i}) - \widetilde{u}_i(0, x_{-i}) > 0. \tag{21}$$

In other words, according to Equation 21, player i's best response is to choose the action that maximizes its payoff with respect to the other player's action (in the case of indifference, the player favors action 0). A PSNE is given by (x_1^*, x_2^*) such that both of players are best responding to each other simultaneously.[17] 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 [Heckman 1978; Schmidt 1981]. 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 [18] of the two players and express the probability of each PSNE in terms of these probabilities. They then use *maximum-likelihood* to estimate the probability of observing any particular PSNE, along with the 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 the standard, classical econometrics literature.

[17] Equation (21) 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 (20). 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 possible correlation between the two players.

[18] 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.

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.

D.2.2. Entry Models of Bresnahan and Reiss. Bresnahan and Reiss [1990] investigate entry in monopoly markets using two types of models: a simultaneous, game-theoretic model and a sequential decision making model. 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 $\widetilde{u}_i^M$ and $\widetilde{u}_i^D$ are the payoffs of firm i in a monopoly and a duopoly market, respectively. Similar to Bjorn and Vuong [1984], here the payoff function consists of two parts: an *observable* part and an *unobserved, random* part. The random part is observed by both the players, but not by the researcher. The entry decision (x_1^*, x_2^*) would be a PSNE if and only if the following best response condition holds for all firms $i = 1, 2$:

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

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 no longer strategic. 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 because the market size is already very big when the second firm enters the market.

In a related paper, Bresnahan and Reiss [1991] discuss the issues of the existence and the uniqueness of PSNE in *discrete game models*. They also discuss how one could deal with MSNE in discrete game models and how these models can be extended to the cooperative games setting. Although they motivated the needs for 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 these [certain probability] distributions," which needs to be done on a case by case basis if we would like to consider MSNE.

D.2.3. 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 [Levine 1987]. Berry [1992] takes 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. 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, or the like. Heterogeneity can also the result of unobserved factors. In Berry's model, the payoff of an airline i in a market k is defined as follows.

$$u_{ik}(\mathbf{x}_k) = v_k(f(\mathbf{x}_k)) + \phi_{ik} \tag{22}$$

In Equation 22, we denote by $\mathbf{x}_k$ the vector of the entry decisions of each airline i in a market k, so that $x_{ik} \in 0, 1$ denotes an airlines entry decision. The first term in the payoff function is market specific and captures the competitive effect that comes from the entry decisions of the airlines. The second term is specific to the airline-market pair and is treated as a single index of profitability. Berry imposes several assumptions in order to guarantee the existence of a PSNE and the uniqueness of the *number* of PSNE; this is one way to deal with the equilibrium selection problem as mentioned in the subsection above (Section D.2.2). First, the airlines in a market k can be sorted according to the profitability index ϕ_{ik}, and this ordering is independent of the entry decisions. Second, the function $f(\mathbf{x}_k)$ in the first term of Equation 22 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.

Each of the two terms in the above payoff function (Equation 22) 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 comes from the large number of airlines: It leads to an exponential number of integrations over the random parts. Berry proposes two ways to address this. One way is to impose additional restrictions on the model, such as removing the part of strategic interaction from the payoff function (i.e., the number of entrants does not affect an airline's profit). The other way is to apply simulation estimators [McFadden 1989]. 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 comes from the analytical approach to the problem. An alternative to deal with this would be an *algorithmic approach*, which we pursue in this paper (although not on the same problem).

Berry's model has been subsequently extended by others. Ciliberto and Tamer [2009] 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. 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 and the probabilities estimated from the data. Apart from modeling the airline industry, a different work by Berry et al. [1995] 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.

D.2.4. Other Models. The three early models that we reviewed above in Sections D.2.1, D.2.2, and D.2.3 of this Appendix exhibit some of the key aspects of the general econometrics approach to modeling strategic scenarios, such as the adoption of a random utility model [McFadden 1974; 1981], 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.

D.2.5. Seim's Model of Product Differentiation. The early game-theoretic entry models, such as the one by Berry [1992], focus mostly on firm-specific profits and the competitive effects of multiple entrants, but do not model product differentiation by the firms. Seim [2006] 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. Seim's 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 [1991, p. 60] say, "Games of private information pose much more complicated estimation issues." 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 [Rust 1996] and that the estimation can also be done in a straightforward two-step method [Bajari et al. 2010a]. 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 way.

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 (i.i.d.)* 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. 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 much.

D.2.6. 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. [2006] study the adoption of the 56K modem technology by the *Internet Service Providers (ISPs)* in a market during the late 1990s. 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).

D.2.7. 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 [Sweeting 2009]. In Sweeting's model, the radio stations in a market choose timings for advertisement strategically. Their payoff function captures 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 [Bajari et al. 2010a]. 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.

D.2.8. 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 *Automated-Teller-Machine (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 with many ATMs often charge more surcharge than smaller banks and credit unions [Hannan et al. 2003].

Ishii [2008] models banks and customers as strategic decision makers in a two-sided market to understand the effects of this surcharge. 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)*. 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 the study of various counterfactuals. For example, what would happen if the surcharges are declared illegal? 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 because of 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.

D.2.9. 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. One example is the adoption of *automated clearinghouse (ACH)* technology. ACH provides an electronic equivalent of paper checks and is commonly used in direct deposits and automated bill payments. Ackerberg and Gowrisankaran [2006] estimate the magnitude of network effects in ACH adoption. 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.

Ackerberg 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.

D.2.10. Recent Developments. The subject of discrete game models is an active area of research in econometrics. Recently, Bajari et al. [2010b] presented methods for identification and estimation of discrete game models of complete information, which are also applicable to general normal-form games. 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 [1996]. This is certainly a computationally expensive step for large games. In a separate work, Bajari et al. [2010c] address estimation of discrete game models of incomplete information. 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 Bajari et al. [2010a].

D.3. 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 problems. 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.

D.3.1. Analytic vs. Algorithmic Approaches. First, an econometrics approach to dealing with strategic settings is in large part analytic. For the most part, algorithms designed in econometrics are driven by analytic techniques (e.g., algorithms for estimating parameters as discussed in Appendix D.2). In contrast, our approach here is rooted in algorithms, computation, AI, and machine learning ideas.

Note that we do not claim that an algorithmic approach is better than an analytic approach. However, it is our view that, in certain situations, an algorithmic approach might provide a good alternative to an analytic approach. We believe 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 [Berry 1992]. 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 paper. We show that looking through an algorithmic lens, grounded in an engineering approach in AI and machine learning within CS, helps us solve problems that would otherwise be impossible to manage analytically.

D.3.2. Modeling. A common modeling approach in the econometrics literature that we reviewed is the adoption of the random utility model [McFadden 1974; 1981]. Besides giving a reasonable way of modeling unobservables, this approach sometimes also leads to simple closed-form expressions that are easy to use. For example, the choice probabilities in a discrete game model can be expressed in a closed form when the random parts of the payoff functions are i.i.d. type-1 extreme-value distributed.

In contrast, the modeling approach in this paper is inherently different. We model two-sided microfinance markets using the well-studied concept of abstract economies in classical economics [Debreu 1952; Arrow and Debreu 1954b]. 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 [Fisher 1892], which has been a subject of intense algorithmic study by computer scientists in recent years [Vazirani 2007].

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 such purpose best. 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.

D.3.3. Estimation. Econometrics and CS, and *machine learning* in particular, have diverging views on the issue of learning a model. 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, because the estimated parameters are very often the quantities of interest. To ensure identification, additional restrictions are sometimes imposed on the model. Identification with infinite samples is most common [Bajari et al. 2010b; Bajari et al. 2010c]. In contrast, predictive performance is of primary interest in machine learning.

There are many contested issues between these two disciplines, which are out of scope for this paper. It is not that one of the approaches is good and the other is bad. It is just that they are different. In this paper, we have taken a machine learning approach to estimation. The estimation is done using a bi-level optimization program. In our work, 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 CS and statistics. For example, a common technique for understanding the effects of new product features on consumers is known as bucket testing [Backstrom and Kleinberg 2011], 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. The reason is that in the setting of an online social network, the effect of a new feature can be more accurately tested if a user is exposed to it together with some of her friends.

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 [Backstrom and Kleinberg 2011]. 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 [Ugander et al. 2013]. 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 Baysian approach that provides more accurate estimates under network uncertainties [Toulis and Kao 2013]. In contrast to this line of work, our goal is not causal estimation. We rather want to estimate models that capture strategic interactions.

D.3.4. 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 [Bajari et al. 2010b; Bajari et al. 2010a]. First, the probability that an equilibrium, which is generated by the model, is observed in the data is estimated [Bjorn and Vuong 1984; Bajari et al. 2010b]. Sometimes, instead of considering all possible equilibria, only a few equilibria are considered in this probabilistic approach [Ackerberg and Gowrisankaran 2006]. 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 [Bresnahan and Reiss 1990; Berry 1992]. Third, the choice probabilities can sometimes be bounded between two limits, which guides the selection of an equilibrium [Ciliberto and Tamer 2009].

In our study of microfinance markets, our model can potentially generate multiple equilibria. We select one of these equilibria that is “closest” to the observed data. This equilibrium selection mechanism is embedded in the machine learning 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.

D.3.5. 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 learning 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 [Ishii 2008]. Ackerberg and Gowrisankaran show the comparative effects of subsides to customers and banks on ACH adoption [Ackerberg and Gowrisankaran 2006].

The main focus of this work is a wide range of interventions. 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 [Ballester et al. 2006]. As we discuss in more detail in Sections D.4 and D.5, there are several other examples of this and other similar kinds of interventions within the CS, AI, and machine learning literature.

D.4. Connection to Existing Models of Networked Economies

The model introduced is essentially one of networked economy. This is important because the running times of our algorithms depend *directly* on the number of edges connecting banks and villages. In particular, our most general computational schemes are iterative in nature, and each round takes time linear in the number of edges, which is indirectly related to the number of villages each bank serves. That number could be considerably smaller than the product of the number of banks times the number of villages (i.e., a "fully connected" bipartite graph): indeed, it is not hard to envision cases in which the number of villages that a bank serves may be constant, independent of the total number of banks or villages. Thus the network structure of our model can have immediate practical implications: Exploiting the network structure can yield significant savings in computation time in practice if the number of villages, or banks, is very large.

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 CS 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 [Kakade et al. 2004], modeling networked markets, such as labor markets and trades (see, for example, Chapter 10 of [Jackson 2008]), designing mechanisms with desirable properties for such markets [Babaioff et al. 2004], analyzing how properties such as competition [Blume et al. 2009] and price variation [Kakade et al. 2005] are influenced by the underlying network structure, to the most fundamental algorithmic question of computing an equilibrium point in such settings [Sarma et al. 2007; Kakade et al. 2004]. 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 [Fisher 1892], 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 consists of a set of buyers and *a set of sellers* [Kakade et al. 2005]. 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 the 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 [Kakade et al. 2005; Vazirani 2007]. There is, however, an interesting connection between our model and that of Fisher through an Eisenberg-Gale convex program formulation, which we show in Section 2.4.1.

Arrow and Debreu gave a very generalized mathematical model of games for competitive markets in their seminal work on competitive economies [Arrow and Debreu 1954b]. In fact, a Fisher economy is a special case of an Arrow-Debreu economy (see, for example, [Vazirani 2007]). Arrow and Debreu's proof of the existence of an equilibrium point uses Debreu's concept of abstract economy [Debreu 1952]. Interestingly, an *abstract economy* generalizes the concept of games in normal form in non-cooperative game theory [von Neumann and Morgenstern 1947], and leads to extensions of Nash's seminal result on the "universal" existence of equilibrium in mixed strategies for "finite" normal-form games [Nash 1950], 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. Our model is an instance of an abstract economy. However, for the same reasons discussed in the previous paragraph (i.e., variable endowment), the specific instantiation of an abstract economy leading to the classical Arrow-Debreu model for markets or more specifically, the recently developed graphical extension to the Arrow-Debreu model [Kakade et al. 2004], does not really capture our setting.

As a brief aside, Kakade et al. [2005] performed some preliminary experiments based on "ad hoc" instantiations of their two-sided network economies. They randomly generated economic network structures and combined them with real-word international trade data from the United Nations. They did not apply machine learning or other classical estimation methods to infer their models *per se*. They applied those specifically generated instances to study policy making in international trade. In particular, they studied the price variations that would result from merging the then 15 members of the European Union (E.U.) into a single economic nation.

Our work is different from various others in networked economies [Kranton and Minehart 2001; Babaioff et al. 2004; Sarma et al. 2007; Gale and Kariv 2007; Even-dar et al. 2007; Blume et al. 2009] 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 [Kranton and Minehart 2001]. 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 [Even-dar et al. 2007]. In contrast to that previous work, we do not study network formation here; that is, we treat the spatial structure of the branch-banking MFIs as exogenous (arbitrary, but fixed).

D.5. Connections to Other Work in AI and Machine Learning

There have been a few examples of specific applications of causal strategic inference, including applications to airline security [Kearns and Ortiz 2004; Kearns 2005] and computational biology [Pérez-Breva et al. 2006; Pérez-Breva 2007]. A lot of that work also exploits strategic complementarity for analytical and computational advantage. In some cases, the resulting algorithms led directly to a constructive proof of the existence of equilibria, just like some of our technical results presented here.

Kearns and Ortiz [2004] instantiate a game-theoretic model of interdependent security based on real-world data on airline reservations and explore the effect of "forcing" the "largest" airlines to fully invest in security. Using the instantiated game, the authors illustrate the "tipping behavior" induced by such interventions, leading all other airlines to also fully invest in security, including "smaller" airlines. That behavior happens to also corresponds to a PSNE of the instantiated game (see also Kearns 2005). In that work, there is no machine learning *per se*, just a somewhat "ad hoc" instantiation of the game parameters based on the available real-world dataset. They do use methods from “learning in games” (i.e., best-response dynamics), as a heuristic for equilibrium computation in the instantiated games.

The game-theoretic model of Pérez-Breva et al. [2006] may also be viewed as a type of two-sided network economy. That model also uses an entropic term for the purpose of “diversification” of resources, which in that case were the chemical concentrations of different proteins around the various potential binding sites along the DNA strand. The authors performed interventions akin to removing players (e.g., “knocking out” a gene), or controlling or “fixing” the actions of players (e.g., varying the total level of protein chemical concentrations). That work also applied machine learning ideas similar to ours to design optimization schemes to infer models from real-world micro-array data [Pérez-Breva 2007]. But, the specific nature of the payoff function differs considerably, leading to different analytic, computational, and machine learning challenges.

We believe it is fair to say that none of that work just described in the first few paragraphs of this subsection deals with the higher level of generality of the work described in this article, in terms of the broad range of interventions and inferences.

We introduced the concept of causal strategic inference within AI in our work on *influence games* [Irfan and Ortiz 2014]. Our particular emphasis there was to provide an alternative game-theoretic approach to the study of influence in social networks. We focused on *linear influence games (LIGs)*, which are multiplayer graphical games [Kearns et al. 2001; Kearns 2007] with parametric, linear-quadratic payoff functions, but over a finite set of pure strategies, with particular emphasis on binary values. [19] Honorio and Ortiz [2015] proposed a learning framework for that specific setting, with emphasis on LIGs.

We applied LIGs to two real-world scenarios: congressional voting and court rulings. In the first, we learned LIGs from data on voting records from the the US Senate, which comprises of 100 senators (i.e., each state has 2 senators). The other involved the 9 Justices of the U.S. Supreme Court [Irfan and Ortiz 2014]. We used data on their judicial rulings to infer an LIG for that setting based on the machine learning framework that Honorio and Ortiz [2015] proposed. (Honorio and Ortiz 2015 used the same dataset on congressional voting, but not the one on judicial rulings, to empirically evaluate their proposed learning framework.)

E. PROOFS MISSING FROM THE MAIN BODY OF THE ARTICLE

In this section of the appendix we provide all the proofs left out of the main body of the article.

E.1. Proof of Property 2.7

The proof is by contradiction. 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 $r^*_i \left(T_i - \sum_{j \in V_i} x^*_{j,i} \right) = 0$ (Equation 3) 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 $\sum_{i \in B_j} x_{j,i}$ (Equation 5), and with $r^*_i = 0$ and $e_j \geq 1$, the term $(1 + r^*_i - e_j)$ in the first constraint of P^j_V (Equation 6 in Definition 2.4) becomes ≤ 0. Since $d_j > 0$, that constraint is satisfied for $x^*_{j,i} = +\infty$. But $x^*_{j,i} = +\infty$ con-

[19] One can view LIGs as generalizations of “discrete choice with social interactions” models in econometrics [Brock and Durlauf 2001]. Yet, the focus of study and inference approach is again different. The work in econometrics is mostly analytic, based on a highly constrained version of the models; while ours deals with more general versions of the model and employs an algorithmic approach to the analysis of large, complex systems.

tradicts the constraint $\sum_{j \in V_i} x^*_{j,i} \le T_i$ in the MFI side (Equation 4 in Definition 2.3). Therefore, for every MFI i, $\sum_{j \in V_i} x^*_{j,i} = T_i$ must hold at any equilibrium point.[20]

For the second claim, in order to maximize its objective function $\sum_{i \in B_j} x^*_{j,i}$ (Equation 5), 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 (Equation 6) 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.

E.2. Proof of Property 2.9

The proof is by contradiction. Consider any MFI i. Let $k \in \arg\max_{j \in V_i} e_j$. Suppose that $r^*_i \le e_k - 1$. Then we get by rearranging the terms: $(1 + r^*_i - e_k) \le 0$. This allows $x^*_{j,i}$ to go to $+\infty$ (by Equations 5 and 6 for the special case of $\lambda = 0$), and that violates the constraint $\sum_{j \in V_i} x^*_{j,i} \le T_i$ (Equation 4), contradicting the assumption that this is an equilibrium point.[21]

E.3. Proof of Property 2.8

The proof is by contradiction. Suppose that for some village j, and for all $i \in B_j$ we have $x^*_{j,i} = 0$. Because $d_j > 0$, the constraint $\sum_{i \in B_j} x^*_{j,i}(1 + r^*_i - e_j) \le d_j$ (Equation 6) of P^j_V (Definition 2.4) is satisfied, but village j is not maximizing $\sum_{i \in B_j} x^*_{j,i}$ (Equation 5). This contradicts that $(\mathbf{r}^*, \mathbf{x}^*)$ is an equilibrium point.

E.4. Proof of Property 2.10

The proof is by contradiction. 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^i_M (Equation 3) in the following way. By Property 2.9, $r^*_i > 0$, and by our modeling assumption, $T_i > 0$ (Section 2.1). Therefore, we have $r^*_i \left(T_i - \sum_{j \in V_i} x^*_{j,i}\right) > 0$.

E.5. Proof of Theorem 2.13

We begin by stating the KKT conditions for P_E (Definition 2.12).

[20] It is important to note that in the above argument, the village side has been allowed 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^i_M (Definition 2.3) treats $x^*_{j,i}$ as exogenous and does not have a direct control over it inside P^i_M. Moreover, the village-side optimization problem P^j_V (Definition 2.4) for village j selects $(x^*_{j,i})_{i \in B_j}$ in order to maximize its objective function $\sum_{i \in B_j} x^*_{j,i}$ (Equation 5), without considering the MFI-side constraint $\sum_{j \in V_i} x^*_{j,i} \le T_i$ (Equation 4). The contradiction results from the necessary condition that *at any equilibrium point* $(\mathbf{x}^*, \mathbf{r}^*)$, all the constraints of all programs $(P^i_M)_{i \in [n]}$ and $(P^j_M)_{j \in [m]}$ must be satisfied.

[21] Once again, village j's demand $x^*_{j,i}$ is determined by P^j_V (Definition 2.4) without trying to satisfy the constraints of any P^i_M (Definition 2.3), given by Equations 3 and 4, as well as the non-negative constraint. However, at an equilibrium point $(\mathbf{r}^*, \mathbf{x}^*)$, both the village-side and the MFI-side problems must be optimized simultaneously.

Stationary condition:

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

Evaluating this at $z^*_{j,i}$ for any $i \in \{1, ..., n\}$ and any $j \in V_i$, we obtain the following.[22]

$$-\frac{1}{\sum_{k \in B_j} z^*_{j,k}} + \gamma^*_i - \mu^*_{j,i} = 0 \tag{23}$$

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:

$$\gamma^*_i \left(\sum_{j \in V_i} z^*_{j,i} - T_i \right) = 0, \qquad 1 \leq i \leq n \tag{24}$$

$$\mu^*_{j,i} \left(-z^*_{j,i} \right) = 0, \qquad 1 \leq i \leq n, \; j \in V_i \tag{25}$$

Note that if $\gamma^*_i > 0$, then Equation 24 gives us the following.

$$\sum_{j \in V_i} z^*_{j,i} - T_i = 0, \qquad 1 \leq i \leq n \tag{26}$$

Furthermore, if $z^*_{j,i} > 0$ then Equation 25 implies $\mu^*_{j,i} = 0$. In that case, we obtain the following from the stationary condition (Equation 23):

$$\gamma^*_i = \frac{1}{\sum_{k \in B_j} z^*_{j,k}} \,. \tag{27}$$

We obtain the following properties from the optimality conditions of the Eisenberg-Gale convex program P_E (Definition 2.12).

LEMMA E.1. *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 of the objective function 7 in Definition 2.12, 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. □

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 E.2. *For any j, $|I^*(j)| > 0$.*

[22] We remind the reader that the quantities at an optimal solution are denoted by *.

PROOF. Suppose that $z^*_{j,i} = 0$ for some j and for all $i \in B_j$. Rearranging the terms of (23), 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 (24), because $T_i > 0$ by our modeling assumption (Section 2.1). Therefore, for any j and some $i \in B_j$, we have $z^*_{j,i} > 0$, which completes the proof. $\square$

Another way of proving Lemma E.2 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 (2.12) goes to $+\infty$. This cannot happen, because P_E (Definition 2.12) is minimizing the objective function (Equation 7), 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 $i \in \{1, ..., n\}$ and all $j \in V_i$).

We can use Lemma E.2 to rewrite Equation 27 in terms of $I^*(j)$. For any j and any $i^*(j) \in I^*(j)$, the following holds.

$$\gamma^*_{i^*(j)} = \frac{1}{\sum_{k \in B_j} z^*_{j,k}} \tag{28}$$

To present an Eisenberg-Gale formulation of our market model given by the programs $(P^i_M)_{i \in [n]}$ (Definition 2.3) and $(P^j_V)_{j \in [m]}$ (Definition 2.4), we define the following terms.

$$x^*_{j,i} \equiv z^*_{j,i}, \qquad \text{for all } i \in \{1, ..., n\} \text{ and all } j \in V_i \tag{29}$$

$$r^*_{i^*(j)} \equiv \gamma^*_{i^*(j)} d + e - 1, \qquad \text{for all } j \in \{1, ..., m\} \text{ and all } i^*(j) \in I^*(j) \tag{30}$$

Note that in Equation 30 above, the equilibrium value r^*_i has not been explicitly defined for each $i \in 1, ..., n$. We first prove that $\cup_{j \in \{1,...,m\}} I^*(j) = \{1, ..., n\}$ to ensure that r^*_i is defined for all i. 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 $I^*(.)$, 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, we have $z^*_{j,i} = 0$, which violates Lemma E.1.

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 given in Equations 29 and 30 above to show that none of the villages has any left-over money.

$$\begin{aligned} d &= \frac{1 + r^*_{i^*(j)} - e}{\gamma^*_{i^*_(j)}} && \text{[Rearranging (30)]} \\ &= (1 + r^*_{i^*(j)} - e) \sum_{k \in B_j} z^*_{j,k} && \text{[Using (28)]} \end{aligned} \tag{31}$$

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

$$\gamma^*_{i^*(j)} = \min_{k \in B_j} \gamma^*_k.$$

By Equation 28, 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 $i^*(j) \in I^*(j)$, we have

$$\gamma^*_{i^*(j)} = \min_{k \in B_j} \gamma^*_k.$$

Thus, using the definition of r^* from (30), we have $r^*_{i^*(j)} = \min_{k \in B_j} r^*_k$. Furthermore, for any $l \in B_j - I^*(j)$, $z^*_{j,l} = 0$. We obtain from Equation 31,

$$\begin{aligned} d &= (1 + \min_{l \in B_j} r^*_l - e) \sum_{k \in B_j} z^*_{j,k} \\ &= \sum_{k \in B_j} z^*_{j,k}(1 + r^*_k - e). \end{aligned}$$

Using the definition of $x^*_{j,k}$ from (29), we obtain

$$d = \sum_{k \in B_j} x^*_{j,k}(1 + r^*_k - e).$$

Furthermore, by Lemma E.1, 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 (Equation 23) 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 (30). Also, (26) holds for $\gamma^*_i > 0$. Again, using the definition of $x^*_{j,i}$ from Equation 29, the other equilibrium condition for our model can be obtained from Equation 26:

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

Thus, the theorem follows.

E.6. Proof of Proposition 2.15

The proof is by contradiction. 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 with respect to $x_{j,i}$ is

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

which is positive (by Assumption 2.1 and equilibrium condition $x^*_{j,i} \leq T_i$).

E.7. Proof of Proposition 2.16

The proof of $e^i_{max} - 1 < r^*_i$ is similar to the proof of Property 2.9.

To prove the upper bound, note that the total amount of loan that villages in V_i can seek from MFI i is at most $\sum_{j \in V_i} \frac{d_j}{1+r_i-e_j}$; we obtain this bound from the first constraint (Equation 6) in the village-side optimization programs $(P^j_V)_{j\in[m]}$ (Definition 2.4), when each village in V_i seeks loan *only* from MFI i). We have that the following holds at an equilibrium point:

$$\begin{aligned} T_i &\leq \sum_{j \in V_i} \frac{d_j}{1+r^*_i - e_j} \\ &\leq d^i_{max} \sum_{j \in V_i} \frac{1}{1+r^*_i - e_j} \\ &\leq d^i_{max} \frac{|V_i|}{1+r^*_i - e^i_{max}} \, . \end{aligned} \tag{32}$$

Rewriting the last condition (Equation 32), we obtain $r^*_i \leq \frac{|V_i| d^i_{max}}{T_i} + e^i_{max} - 1$.

E.8. Proof of Lemma 3.1

Following is the Lagrangian function [Boyd and Vanderberghe 2004] of the village-side optimization program P^j_V (Definition 2.4) for village j:

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

At an optimal solution, we have $\frac{\partial L}{\partial x_{j,i}} = 0$ for any $i \in B_j$. Thus, we have

$$\begin{aligned} &-1 - \lambda \log \frac{1}{x^*_{j,i}} + \lambda x^*_{j,i} \frac{1}{x^*_{j,i}} + \alpha_j (1 + r_i - e_j) = 0 \\ &\Leftrightarrow x^*_{j,i} = \exp\left(\frac{1 - \lambda - \alpha_j(1 + r_i - e_j)}{\lambda} \right) . \end{aligned} \tag{33}$$

By Property 2.15, we have $\sum_{i \in B_j} x^*_{j,i}(1 + r_i - e_j) = d_j$. Substituting the expression for $x^*_{j,i}$ we obtain the second equation (Equation 33) claimed in the statement. Moreover, we have that α^*_j must be unique; otherwise, by the same expression for $x^*_{j,i}$, we would have multiplicity in the best response of village j, which is precluded by the convex optimization P^j_V.

E.9. Proof of Lemma 3.2

We prove the case of increasing r_i. The other case is similar. First, observe that we cannot simply invoke Equation 9 to prove the statement, because α^*_j also changes with the change in r_i. Furthermore, the direction of change of α^*_j (i.e., increase or decrease) is not immediately clear from Equation 10.

Suppose that the value of r_i has been increased in Line 6. Suppose, for the sake of obtaining 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 2.15, the first constraint (Equation 6) of the village-side program P^j_V (Definition 2.4) 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 strictly 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 3.1 in terms of α^*_j, we obtain the first equation below (Equation 34). The second equation (Equation 35) follows similarly.

$$\alpha^*_j = \frac{1-\lambda-\lambda\log x^*_{j,i}}{1+r_i-e_j} \tag{34}$$

$$\alpha^*_j = \frac{1-\lambda-\lambda\log x^*_{j,k}}{1+r_k-e_j} \tag{35}$$

By Equation 34, our assumption that $x^*_{j,i}$ has increased or remained the same in response to the increase of r_i implies that the value of α^*_j has strictly decreased from its previous value. Therefore, by Equation 35, the value of $x^*_{j,k}$ must strictly increase, which gives us a contradiction (to see this, 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 strictly decrease, for all $j \in V_i$.

E.10. Proof of Lemma 3.3

Consider a village $j \in V_i$. By Lemma 3.2, when an MFI i strictly increases its interest rate in Line 6, village j must strictly decrease $x^*_{j,i}$ in Line 7. Considering Equation 34, it may at first seem possible that the value of α^*_j can increase, decrease, or even remain the same, depending on how much $x^*_{j,i}$ has decreased. However, we will next show that α^*_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 10 as follows.

$$\begin{aligned}
&\sum_{i\in B_j} \exp\left(\frac{1-\lambda}{\lambda}\right)\exp\left(-\beta^*_j\rho_{j,i}\right)\rho_{j,i} = d_j \\
\Leftrightarrow &\sum_{i\in B_j} \frac{\rho_{j,i}}{\exp\left(\beta^*_j\rho_{j,i}\right)} = d_j \exp\left(\frac{-1+\lambda}{\lambda}\right)
\end{aligned} \tag{36}$$

Here, the right-hand side of Equation 36 is constant, since λ 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 β^*_j unchanged, then the left-hand side must decrease.[23] 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

[23]Note that increasing β^*_j will only further decrease the left hand side.

that the derivative of this term with respect to $\rho_{j,i}$ is non-positive:

$$\begin{aligned}
&\frac{1}{\exp(\beta_j^*\rho_{j,i})} - \frac{\beta_j^*\rho_{j,i}}{\exp(\beta_j^*\rho_{j,i})} \leq 0\\
&\Leftrightarrow \rho_{j,i}\beta_j^* \geq 1\\
&\Leftrightarrow (1 + r_i - e_j)\frac{\alpha_j^*}{\lambda} \geq 1\\
&\Leftrightarrow 1 - \lambda - \lambda \log x_{j,i}^* \geq \lambda, \;\; \text{by Equation 34}\\
&\Leftrightarrow \lambda \leq \frac{1}{2 + \log x_{j,i}^*} \;,
\end{aligned}$$

which holds by Assumption 2.1. Therefore, the value of α_j cannot increase when r_i increases.

Because α_j can only decrease when r_i increases, using Equation 35, 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 3.2, decreasing its interest rate cannot be MFI k's best response.

E.11. Proof of Theorem 3.4

Algorithm 1 begins with initial values of interest rates arbitrarily close to their lower bound established in Property 2.16. By Lemma 3.3, at the end of each iteration of the for loop over the MFIs in Algorithm 1, the interest rate r_i cannot go down (i.e., its value cannot be lower than that at the start of that iteration). By Lemma 3.1, every village has a unique best response to these interest rates. Now, the interest rates are upper bounded by Property 2.16. 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.

ACKNOWLEDGMENTS

We thank the anonymous TEAC reviewers for suggesting many improvements. We also thank David Parkes for his generous feedback.

REFERENCES

Jaap H Abbring and James J Heckman. 2007. Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation. *Handbook of econometrics* 6 (2007), 5145–5303.

Daniel A Ackerberg and Gautam Gowrisankaran. 2006. Quantifying equilibrium network externalities in the ACH banking industry. *The RAND journal of economics* 37, 3 (2006), 738–761.

Victor Aguirregabiria and Junichi Suzuki. 2015. Empirical games of market entry and spatial competition in retail industries. (2015).

Kenneth J. Arrow and Gerard Debreu. 1954a. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22, 3 (July 1954), 265–290.

Kenneth J. Arrow and Gerard Debreu. 1954b. Existence of an equilibrium for a competitive economy. *Econometrica* 22, 3 (1954), 265–290.

Angelique Augereau, Shane Greenstein, and Marc Rysman. 2006. Coordination versus differentiation in a standards war: 56K modems. *The Rand Journal of Economics* 37, 4 (2006), 887–909.

Moshe Babaioff, Noam Nisan, and Elan Pavlov. 2004. Mechanisms for a spatially distributed market. In *Proceedings of the 5th ACM conference on Electronic commerce (EC '04)*. ACM, New York, NY, USA, 9–20. DOI:http://dx.doi.org/10.1145/988772.988776

Lars Backstrom and Jon Kleinberg. 2011. Network bucket testing. In *Proceedings of the 20th international conference on World wide web*. ACM, 615–624.

Henry K Bagazonzya, Zaid Safdar, AKM Abdullah, Cecile Thioro Niang, and Aneeka Rahman. 2010. *Linking up and reaching out in Bangladesh: information and communications technology for microfinance*. World Bank Publications.

Patrick Bajari, Han Hong, John Krainer, and Denis Nekipelov. 2010c. Estimating static models of strategic interactions. *Journal of Business & Economic Statistics* 28, 4 (2010).

Patrick Bajari, Han Hong, and Denis Nekipelov. 2010a. *Game theory and econometrics: A survey of some recent research*. Working paper. University of Minnesota, Department of Economics.

Patrick Bajari, Han Hong, and Stephen P Ryan. 2010b. Identification and estimation of a discrete game of complete information. *Econometrica* 78, 5 (2010), 1529–1568.

Coralio Ballester, Antoni Calv-Armengol, and Yves Zenou. 2006. Who's Who in Networks. Wanted: The Key Player. *Econometrica* 74, 5 (09 2006), 1403–1417. http://ideas.repec.org/a/ecm/emetrp/v74y2006i5p1403-1417.html

Jonathan F Bard. 2013. *Practical bilevel optimization: algorithms and applications*. Vol. 30. Springer Science & Business Media.

Leonardo Becchetti and Stefano Castriota. 2011. Does microfinance work as a recovery tool after disasters? Evidence from the 2004 tsunami. *World Development* 39, 6 (2011), 898–912.

Steven Berry, James Levinsohn, and Ariel Pakes. 1995. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society* (1995), 841–890.

Steven T Berry. 1992. Estimation of a Model of Entry in the Airline Industry. *Econometrica: Journal of the Econometric Society* (1992), 889–917.

Paul A Bjorn and Quang H Vuong. 1984. *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.

Lawrence E. Blume, David Easley, Jon Kleinberg, and Éva Tardos. 2009. Trading networks with price-setting agents. *Games and Economic Behavior* 67, 1 (2009), 36 – 50. DOI:http://dx.doi.org/10.1016/j.geb.2008.12.002

Stephen Boyd and Lieven Vanderberghe. 2004. *Convex Optimization*. Cambridge University Press.

Timothy F Bresnahan and Peter C Reiss. 1990. Entry in monopoly market. *The Review of Economic Studies* 57, 4 (1990), 531–553.

Timothy F Bresnahan and Peter C Reiss. 1991. Empirical models of discrete games. *Journal of Econometrics* 48, 1 (1991), 57–81.

William A. Brock and Steven N. Durlauf. 2001. Discrete Choice with Social Interactions. *The Review of Economic Studies* 68, 2 (2001), 235–260. http://dx.doi.org/10.1111/1467-937X.00168.

Jeremy I. Bulow, John D. Geanakoplos, and Paul D. Klemperer. 1985. Multimarket Oligopoly: Strategic Substitutes and Complements. *The Journal of Political Economy* 93, 3 (1985), 488–511. http://www.jstor.org/stable/1832005

Federico Ciliberto and Elie Tamer. 2009. Market structure and multiple equilibria in airline markets. *Econometrica* 77, 6 (2009), 1791–1828.

B.Armendariz de Aghion and J. Morduch. 2005. *The economics of microfinance*. MIT Press.

Gerard Debreu. 1952. A Social Equilibrium Existence Theorem. *Proceedings of the National Academy of Sciences* 38, 10 (1952), 886–893.

Pradeep Dubey, Ori Haimanko, and Andriy Zapechelnyuk. 2006. Strategic complements and substitutes, and potential games. *Games and Economic Behavior* 54, 1 (2006), 77 – 94. DOI:http://dx.doi.org/DOI:10.1016/j.geb.2004.10.007

E. Eisenberg and D. Gale. 1959. Consensus of subjective probabilities: The Pari-Mutuel method. *Annals Math. Stat.* 30 (1959), 165–168.

Eyal Even-dar, Michael Kearns, and Siddharth Suri. 2007. A Network Formation Game for Bipartite Exchange Economies. In *ACM-SIAM Symp. on Discrete Algorithms (SODA)*.

Irving Fisher. 1892. *Mathematical Investigations in the Theory of Value and Prices*. Yale University.

Ronald A. Fisher. 1935. *The Design of Experiments*. Oliver and Boyd.

Drew Fudenberg and Jean Tirole. 1991. *Game Theory*. The MIT Press.

Douglas M. Gale and Shachar Kariv. 2007. Financial Networks. *American Economic Review* 97, 2 (2007), 99–103. DOI:http://dx.doi.org/10.1257/aer.97.2.99

A. Gelman, J.B. Carlin, H.S. Stern, and D.B. Rubin. 2003. *Bayesian Data Analysis*. Chapman & Hall/CRC. http://books.google.com/books?id=TNYhnkXQSjAC

M. Ghatak and T.W. Guinnane. 1999. The Economics of Lending with Joint Liability: Theory and Practice. *Journal of Development Economics* 60 (1999), 195–228.

Clive W J Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society* (1969), 424–438.

Timothy H Hannan, Elizabeth K Kiser, Robin A Prager, and James J McAndrews. 2003. To surcharge or not to surcharge: An empirical investigation of ATM pricing. *Review of Economics and Statistics* 85, 4 (2003), 990–1002.

James J Heckman. 1978. Dummy Endogenous Variables in a Simultaneous Equation System. *Econometrica* 46, 4 (1978), 931–959.

James J Heckman, Rosa L Matzkin, and Lars Nesheim. 2003. Simulation and estimation of hedonic models. (2003).

James J Heckman and Edward J Vytlacil. 2007. Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation. *Handbook of econometrics* 6 (2007), 4779–4874.

Jean Honorio and Luis Ortiz. 2015. Learning the Structure and Parameters of Large-Population Graphical Games from Behavioral Data. *Journal of Machine Learning Research* 16 (2015), 1157–1210. http://jmlr.org/papers/v16/honorio15a.html

K.D. Hoover. 2001. *Causality in Macroeconomics*. Cambridge University Press. http://books.google.com/books?id=4pGZLXO71jkC

Institute of Microfinance (InM). 2009. *Bangladesh Microfinance Statistics*. The University Press Limited (UPL), Dhaka 1000, Bangladesh. www.inm.org.bd/publication.htm

Mohammad T. Irfan. 2013. *Causal Strategic Inference in Social and Economic Networks*. Ph.D. Dissertation. Stony Brook University. http://www.bowdoin.edu/~mirfan/papers/Mohammad_Tanvir_Irfan_Dissertation.pdf.

Mohammad T. Irfan and Luis E. Ortiz. 2014. On influence, stable behavior, and the most influential individuals in networks: A game-theoretic approach. *Artificial Intelligence* 215 (2014), 79 – 119. DOI:http://dx.doi.org/10.1016/j.artint.2014.06.004

Joy Ishii. 2008. Compatibility, competition, and investment in network industries: ATM networks in the banking industry. *Unpublished working paper* (2008).

Matthew O. Jackson. 2008. *Social and Economic Networks*. Princeton University Press, Princeton, NJ, USA.

Sham M. Kakade, Michael Kearns, and Luis E. Ortiz. 2004. Graphical economics. In *In Proceedings of the 17th Annual Conference on Learning Theory (COLT*. Springer, 17–32.

Sham M. Kakade, Michael Kearns, Luis E. Ortiz, Robin Pemantle, and Siddharth Suri. 2005. Economic properties of social networks. In *Advances in Neural Information Processing Systems 17*. MIT Press, 633–640.

Michael Kearns. 2005. Economics, Computer Science, and Policy. *Issues in Science and Technology* (Winter 2005). http://www.cis.upenn.edu/~mkearns/papers/ist.pdf

Michael Kearns. 2007. Graphical Games. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vaziran i (Eds.). Cambridge University Press, Chapter 7, 159–180.

M. Kearns, M. Littman, and S. Singh. 2001. Graphical Models for Game Theory. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*. 253–260.

Michael Kearns and Luis E Ortiz. 2004. Algorithms for Interdependent Security Games. In *Advances in Neural Information Processing Systems 16*, S. Thrun, L. K. Saul, and B. Schölkopf (Eds.). MIT Press, 561–568. http://papers.nips.cc/paper/2495-algorithms-for-interdependent-security-games.pdf

Timothy J Kehoe, Thirukodikaval N Srinivasan, and John Whalley. 2005. *Frontiers in Applied General Equilibrium Modeling: In Honor of Herbert Scarf*. Cambridge University Press.

Shahidur R Khandker. 2005. Microfinance and poverty: Evidence using panel data from Bangladesh. *The World Bank Economic Review* 19, 2 (2005), 263–286.

Jon Kleinberg. 2007. Cascading Behavior in Networks: Algorithmic and Economic Issues. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, Chapter 24, 613–632.

Rachel E. Kranton and Deborah F. Minehart. 2001. A Theory of Buyer-Seller Networks. *American Economic Review* 91, 3 (2001), 485–508. DOI:http://dx.doi.org/10.1257/aer.91.3.485

Michael E Levine. 1987. Airline competition in deregulated markets: Theory, firm strategy, and public policy. *Yale J. on Reg.* 4 (1987), 393–503.

D. Little. 1995. *On the Reliability of Economic Models: Essays in the Philosophy of Economics*. Kluwer Academic Publishers. http://books.google.com/books?id=NqcY21jzDEwC

R. Duncan Luce. 1959. *Individual Choice Behavior: A Theoretical Analysis*. Wiley New York.

Rubana Mahjabeen. 2008. Microfinancing in Bangladesh: Impact on households, consumption and welfare. *Journal of Policy modeling* 30, 6 (2008), 1083–1092.

Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, and others. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.

Daniel McFadden. 1974. Conditional logit analysis of qualitative choice behavior. In *Frontiers of Econometrics*, P. Zarembka (Ed.). Academic Press, New York, NY, 105–142.

Daniel McFadden. 1981. Econometric models of probabilistic choice. In *Structural analysis of discrete data*, C. Manski and D. McFadden (Eds.). MIT Press, Cambridge, 198–272.

Daniel McFadden. 1989. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society* (1989), 995–1026.

Richard McKelvey and Thomas Palfrey. 1995. Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior* 10 (1995), 6–38.

Richard D. McKelvey and Andrew McLennan. 1996. Computation of equilibria in finite games. In *Handbook of Computational Economics*. Vol. I. 87–142.

J. Morduch. 1999. The role of subsidies in microfinance: evidence from the Grameen Bank. *Journal of Development Economics* 60 (1999), 229–248.

John Nash. 1951. Non-cooperative games. *Annals of Mathematics* 54 (September 1951), 286–295. Issue 2.

John F. Nash. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences* 36, 1 (1950), 48–49.

J. Pearl. 1994. From Bayesian networks to causal networks. *Proceedings of the Adaptive Computing and Information Processing Seminar* (1994), 25–27.

J. Pearl. 1998. Graphical models for probabilistic and causal reasoning. *Handbook of Defeasible Reasoning and Uncertainty Management Systems* 1 (1998), 367–389.

J. Pearl. 2000. *Causality: Models, Reasoning, and Inference*. University Press. http://books.google.com/books?id=wnGU_TsW3BQC

J. Pearl and T.S. Verma. 1991. A theory of inferred causation. (1991).

J. Pearl and T.S. Verma. 1992. A statistical semantics for causation. *Statistics and Computing* 2, 2 (1992), 91–95.

J. Pearl and T.S. Verma. 1995. A theory of inferred causation. *Studies in Logic and the Foundations of Mathematics* 134 (1995), 789–811.

Luis Pérez-Breva. 2007. *DNA Binding Economies*. Ph.D. Dissertation. Massachusetts Institute of Technology.

Luis Pérez-Breva, Luis E. Ortiz, Chen-Hsiang Yeang, and Tommi Jaakkola. 2006. Game Theoretic Algorithms for Protein-DNA binding. In *NIPS*, Bernhard Schölkopf, John C. Platt, and Thomas Hoffman (Eds.). MIT Press, 1081–1088.

David Porteous. 2006. Competition and Microcredit Interest Rates. *Focus Note* 33 (2006).

Tim Roughgarden. 2002. *Selfish routing*. Ph.D. Dissertation. Cornell University.

John Rust. 1996. Estimation of dynamic structural models, problems and prospects: discrete decision processes. In *Advances in econometrics: Sixth world congress*, Vol. 2. Cambridge Univ Pr, 119–170.

Mir M Salim. 2013. Revealed objective functions of microfinance institutions: evidence from Bangladesh. *Journal of Development Economics* 104 (2013), 34–55.

Atish Das Sarma, Deeparnab Chakrabarty, and Sreenivas Gollapudi. 2007. Public advertisement broker markets. In *Proceedings of the 3rd international conference on Internet and network economics (WINE'07)*. Springer-Verlag, Berlin, Heidelberg, 558–563. http://dl.acm.org/citation.cfm?id=1781894.1781967

Peter Schmidt. 1981. Constraints on the parameters in simultaneous tobit and probit models. *Structural analysis of discrete data with econometric applications* (1981), 422–434.

Katja Seim. 2006. An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics* 37, 3 (2006), 619–640.

Yogeshwer Sharma and David P Williamson. 2009. Stackelberg thresholds in network routing games or the value of altruism. *Games and Economic Behavior* 67, 1 (2009), 174–190.

Andrew Sweeting. 2009. The strategic timing incentives of commercial radio stations: An empirical analysis using multiple equilibria. *The RAND Journal of Economics* 40, 4 (2009), 710–742.

Jan Tinbergen. 1956. On the theory of income distribution. *Weltwirtschaftliches archiv* (1956), 155–175.

Panos Toulis and Edward Kao. 2013. Estimation of Causal Peer Influence Effects. In *Proceedings of international conference on Machine learning (ICML)*. ACM, New York, NY, USA.

Johan Ugander, Brian Karrer, Lars Backstrom, and Jon M. Kleinberg. 2013. Graph cluster randomization: network exposure to multiple universes. *CoRR* abs/1305.6979 (2013).

V. V. Vazirani. 2007. Combinatorial Algorithms for Market Equilibria. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, Chapter 5, 103–134.

John von Neumann and Oskar Morgenstern. 1947. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ. Second Edition.

D. Wright and D. Alamgir. 2004. Microcredit Interest Rates in Bangladesh 'Capping v Competition'. *Donors Local Consultative Group on Finance* (March 2004).

D.S. Karlan X. Gine, P. Jakiela and J. Morduch. 2006. *Microfinance Games*. Discussion paper 936. Yale University Economic Growth Center.

Received April 2016; revised ???? 20??; accepted ???? 20??