Information and Persuasion in Competitive Markets and Politics

A DISSERTATION

SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL

OF THE UNIVERSITY OF MINNESOTA

BY

Brian C. Albrecht

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Advisor: David Rahman

July, 2020

Acknowledgements

I owe many thanks to David Rahman for his guidance throughout the whole PhD. His passion while teaching the first-year course in microeconomic theory led me to work on economic theory. Each of the chapters in this dissertation grew out of conversations with him, where he more than generously offered up his time, help, and advice.

Thanks to the Minnesota Economic Theory Workshop for feedback and encouragement along the way. A few participants deserve special mention for their help throughout: Sergio Ocampo-Diaz, David Ruiz, Aldo Rustichini, Jan Werner, and especially Rafael Guthmann who I am lucky enough to co-author with on Chapter 3.

Outside of the wonderful Department of Economics at the University of Minnesota, I would like to thank the Hayek Program at the Mercatus Center for its support throughout my PhD and especially to Pete Boettke. Finally, I want to thank Josh Hendrickson for all his guidance on how to do economic research and for letting me be his co-author on other papers we wrote during my PhD.

Dedication

To Brynn. Te quiero.

Abstract

This dissertation studies the complex interplay between information and competition, using a variety of formal models from general equilibrium and game theory. I first study two models that extend a baseline, general equilibrium model. In the chapter "Investment without Coordination Failures", I prove that, even with incomplete markets, the prospect of competitive markets in the future provides sufficient information for decision-makers today to choose efficient investment levels.

In "On the Informational Efficiency of Decentralized Price Formation", written jointly with Rafael R. Guthmann, we prove that decentralized search economies which approximate a competitive economy in terms of allocation do not approximate the competitive economy in terms of the required information. We prove that a model of decentralized market-makers can approximate the informational efficiency of a competitive economy.

I then move on to study two different models, using tools from the recent field of information design. Again, I focus on the role of competition. In "Price Competition and the Use of Consumer Data", I study how changes in the information available to competing firms affects the distribution of consumer and producer surplus. I prove that, unlike under a monopoly, complete information–which generates first-degree price discrimination–is optimal for consumers.

I conclude by studying a model where people compete through information. I study a voting model where two political parties compete in an election by persuading voters of their candidate's quality. In the unique equilibrium, both parties design campaigns that generate uniform distributions of beliefs. The equilibrium means that voters are equally likely to have a range of beliefs about the candidate's quality after the campaign: some voters believe the candidate is perfect, some believe the candidate is terrible, and some believe everything in between. The level of disagreement is driven by competition; each party does not want to design a campaign that is easy to beat.

Contents

List of Figures

Chapter 1

Introduction

> "Competition may be the spice of life, but in economics it has been more nearly the main dish. Competition has been a major force in the organization of production and the determination of prices and incomes: *economic theory has accorded commensurate importance to the concept.*"
>
> George Stigler (1968, 3:181)

While the nature and meaning of competition (and especially perfect competition) have always been central topics within economic theory, the nature and meaning of information have taken up a similarly important role within economic theory since George Stigler penned the above epigraph. Using a series of different models, this dissertation studies the interplay between these two central features of economic theory: competition and information.

This dissertation asks the following types of questions about that interplay:

Can the belief in future competition create incentives for efficient investment today? (Chapter 2) Do certain approximately competitive mechanisms require approximately as little information as competitive markets? (Chapter 3) When firms have information on consumers and are competing through price, what information is optimal for those consumers? (Chapter 4) How can political parties compete by providing voters with information? (Chapter 5) These—and related questions—are the direct focus of this dissertation.

The remaining chapters are divided into two sections. The first section includes two chapters where I build on standard general equilibrium models of perfect competition. In Chapter 2, I study games with incomplete markets where agents must sink their investments before they can join a match that generates value. I focus on competitive matching markets where there is a public price to join any match. Despite the First Welfare Theorem, coordination failures can still arise because of the market incompleteness. Armen does not invest because Bengt does not invest, vice versa, and they are unable to write enforceable contracts to ensure joint investments. In these games, multiple equilibria can exist, with both efficient investment and not. However, the standard, Nash solution concept used in these games does not help in determining if all equilibria are equally robust or stable. I argue that we should replace the Nash solution concept in this context with a mild refinement: trembling-hand perfection. I prove that—in a general class of models with general heterogeneity of types, costs of investments, and matching surpluses—small trembles rule out coordination failures. My main theorem is a modified First Welfare Theorem: *even with endogenous and incomplete markets*, every perfect, competitive equilibrium is efficient. This result shows how competitive markets, even when they are incomplete, can provide the necessary information to economic actors for successful coordination.

It is well understood that the competitive allocation process is informationally efficient, meaning it requires the minimum amount of information possible to implement a Pareto efficient allocation. In Chapter 3, co-authored with Rafael R. Guthmann, we study the informational efficiency in models with *strategic and decentralized* price formation. First, we show in a two-good economy that the standard random search and bargaining model requires an infinitely larger message space compared to the competitive equilibrium for a large economy–making it infinitely less informationally efficient. We propose here a model of price formation through market-makers. As the random search equilibrium, this model of price formation attains the competitive allocation in the limit, but, in a quasi-linear environment with L goods, we prove that the market-maker mechanism only requires a message space with L-1 more dimensions than the competitive process. This appears to be the most informationally efficient form of decentralized price formation process that implements the competitive allocation at the limit.

Chapters 4 and 5 expand recent insights from the fields of information design (Taneva 2019) and Bayesian persuasion (Kamenica and Gentzkow 2011) to environments with competitive actors. Chapters 4 starts from the observation that modern firms have access to vast amounts of data on consumers, which allows them to strategically vary prices across different consumers, *i.e.* price discriminate. To study the effects of this data on consumer welfare, I develop a Bertrand duopoly model where each consumer's valuation for each firm's good is uncertain. Instead of imposing that firms have access to specific data, I allow for general information structures; firms may vary in the quality and form of their data. Fixing the available data, due to the discontinuities in Bertrand competition, the unique equilibrium is only supported through price dispersion. I directly construct the unique equilibrium by harnessing features of each firm's residual demand curve. In equilibrium, each firm randomizes her price and generates a unit-elastic residual demand for the other firm. I then vary the available data and compare the welfare consequences. In the baseline model, contrary to common concerns regarding price discrimination derived from the monopoly case, under competition, completely public consumer data (perfect price discrimination) is *optimal* for consumers.

Chapter 5 studies how actors can actually use information, as in the Bayesian persuasion literature, as their means for competing. I construct a voting model where two political parties compete by designing campaigns that release information about their party's candidate. Campaigns generate distributions of voter beliefs about the candidate's quality. Under competition, each party must worry about the other party and does not want to design a campaign that is easy to beat. In the unique equilibrium, both parties design campaigns that generate uniform distributions. The uniform distribution means that each voter is equally likely to have a range of beliefs about the candidate's quality after the campaign: some voters believe the candidate is perfect, some believe the candidate is terrible, and some believe everything in between. This is the outcome of the election even though every voter has the exact same beliefs before the election. The uniform distribution also means that the voters as a whole have maximum uncertainty about the candidates. The uniform distribution is a common feature in zero-sum games where each party wants to win but only needs a small margin, such as all-pay auctions or Colonel Blotto games. The paper also highlights a similarity between buying votes with money and using persuasion.

Information and Extensions of Perfect Competition

Chapter 2

Investment without Coordination Failures

Many real-world investments exhibit coordination; one person's investment requires an appropriately matching investment from another person. Entrepreneurs who develop hardware need others to develop software. College students who invest in learning programming skills need firms that can harness those skills. These investments only have value together. When deciding to invest, people must trust that market forces will reward their costly investments. This requires two things to happen. First, people must trust that there will be ex-post competition to avoid hold out problems so they earn a return for any value their investment generates. Second, people must trust that there will not be widespread coordination failures so that value can be generated. This chapter focuses on this second concern, coordination, which is central to economics.

I study coordination in a model where agents must first choose their investment non-cooperatively. After sinking their investment, agents enter a competitive matching market where their previous investments determine the set of matches they can choose to join and the price they have to pay to join. There is room for possible coordination in the matching market if there are investment complementarities. For example, Armen's investment in software is complementary to Bengt's investment in hardware. Armen wants to coordinate with the Bengt and only invest if he does, and vice versa. However, since they make their investments non-cooperatively, they cannot write a contract for a joint-deviation; markets are incomplete. If equilibria can be Pareto-ranked, we call any equilibrium that is not Pareto-optimal a coordination failure.[1] Without the possibility to write contracts for joint-deviations of investments, people may be stuck in a coordination failure.

According to economic theory, in markets with coordination, we can expect two types of outcomes: efficient outcomes and coordination failures. However, the existing theory does not give us a means for differentiating the likelihood of these two scenarios. The Nash solution concept used in the literature selects for both stable and unstable outcomes. Are efficient coordination and coordination failures both equally stable or robust? My chapter proves that they are not equally stable.

Formally, Theorem 3 proves that every equilibrium that survives a trembling-hand perfect refinement is efficient. This implies that any coordination failure does not survive, it is not robust, and is not likely to be observed in the world. This

1. A game like Battle of the Sexes does not have an equilibrium which is a coordination failure. While there is coordination and therefore multiple equilibria, they are not Pareto-ranked.

robustness of efficient allocations and non-robustness of coordination failures and provides one reason that real-world market actors may trust in the coordination abilities of markets.[2] When there is a chance that people experiment or make mistakes—as people can do—competitive markets only maximize surplus.

In a sense, the model is a formal argument of how Adam Smith's "higgling and bargaining of the market" can fix coordination failures. If people tremble or experiment with a new strategy that is not prescribed by the coordination failure equilibrium, the coordination failure unravels. If one person in a large economy develops software, another person will observe the software. She will then willing to develop hardware because there will be large returns to the development. This process reveals the instability of the coordination failure and the economy moves to an efficient equilibrium with both software and hardware.

To be more explicit about my environment, I consider the following two-stage game illustrated in Figure 2.1. Besides the fact that the game is just two-stages for simplicity, the environment is extremely general. I start with a large number of buyers and sellers. They are endowed with some arbitrary, finite types. Their types determine how costly it is to invest; there can be good and bad software and hardware developers. In Stage 1, in a non-cooperative setting, all buyers and sellers choose their investment from a set of possible investments. These investments determine the value generated by any match of one buyer and one seller. After choosing their investments, buyers and sellers can buy a contractual right to enter

2. When discussing the formal results in words, I use the term "robustness" compared to more appropriate term "stability" (Kohlberg and Mertens 1986), because stability has a different, common meaning in the matching literature that I want to avoid discussing. All equilibria are stable in the matching sense, even those that are ruled out by trembling hand refinement and therefore not stable in the Kohlberg and Mertens-sense.

into a match with a member of the other side. An investment equilibrium requires that each side chooses its investment level and contract optimally and that prices clear the market so that all contracts are fulfilled.

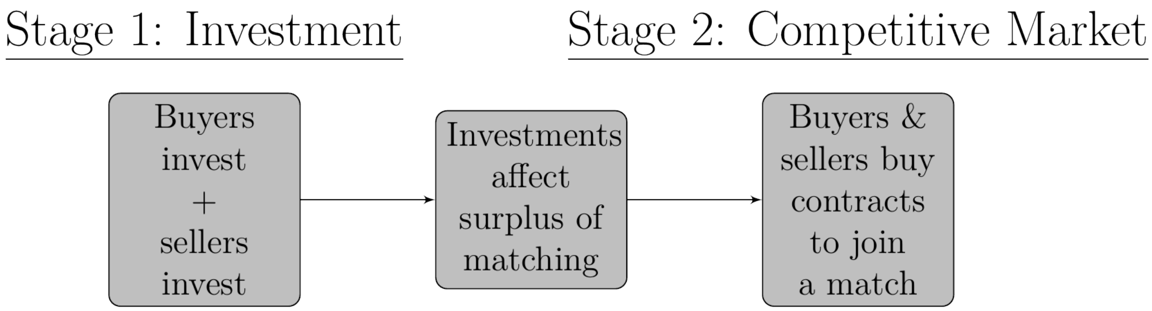

Figure 2.1: Timing of Investment Game

For efficient coordination in this two-stage game, two types of coordination must take place: *market* and *game* coordination.[3] First, the market must coordinate buyers and sellers of different contracts. This is just the coordination of supply and demand through prices.[4] Second, non-cooperative investments must coalesce around the appropriate investments. The competitive nature of the second-stage market guarantees the coordination of the supply and demand of contracts. The second type of coordination does not always take place, which leads to coordination failures. Incorporating both types of coordination better matches real-world situations, where "It is not sufficient for an individual to have complete knowledge of all objective conditions (technology, resources, and so on)" (O'Driscoll 1977, p. 23-4). Theorem 3 provides one situation where we should expect to see

3. Klein and Orsborn (2009) make a similar distinction between "concatenate coordination" and "mutual coordination" and Albrecht (2016) provides a model that ties together the two different forms of coordination through the effort of entrepreneurs.

4. In a standard Walrasian model, prices do *all* the work of coordinating supply and demand in a price-taking model. See Makowski and Ostroy (2013).

full coordination (market and game): whenever there is competition.

It is important to note that my result does not imply that coordination problems cannot exist in the real world. It does warn against searching for them in competitive environments. We should instead think of them as arising in environments with *imperfect* competition, as is the common in the macroeconomics literature since Cooper and John (1988), for example. The theorem also means that if we observe an outcome that looks like a coordination failure, we have a reason to look for the imperfections of competition in the market. The model serves as a foil to compare possible coordination failures against (Albrecht and Kogelmann 2018). As for policy, solving imperfections of competition should solve the coordination problem.

In this chapter, I strive to emphasize the mechanism and not to prove the most general theorem possible. Therefore, I make simplifying assumptions that make the argument easier along the way. For example, I assume quasi-linear (transferable) utility, which avoids the need to distinguish all of the different welfare benchmarks used in the literature. To justify price-taking, I assume throughout a large economy with a continuum of players.[5] Competition (and no externalities) ensures that any inefficiency that could arise comes from coordination failures, as compared to other failures like hold out problems.

The outline of the rest of the chapter is as follows. Section 2.1 briefly goes over related literature on investment, matching, and refinements of competitive models. Section 2.3 goes through a simple example that highlights all of the main

5. As Gretsky, Ostroy, and Zame (1999, p. 63) put it "if we seek (robust) perfect competition we must look to continuum economies."

results and mechanisms of the model. Readers who skim the chapter are encouraged to focus on the example. Section 2.4 lays out the full model and goes through the standard, Nash-style equilibrium. Section 2.5 then constructs the trembling-hand refinement and proves that coordination failures are not robust under the trembling-hand. Section 2.6 concludes.

2.1 Investment and Coordination Literature

To formally study the connection between market coordination and game coordination, I build on a series of papers that have a non-cooperative game before a competitive market.[6] Makowski and Ostroy (1995) showed that if there is full appropriation and investments are non-complementarity, then a First Welfare Theorem holds and any market equilibrium is efficient.[7]

As Makowski and Ostroy show, competition gives full appropriation. However, when there are complementarities, coordination problems can still arise in competitive markets. That means competition alone is not sufficient for efficiency. Following up on Makowski and Ostroy (1995), three important papers of competitive matching Cole, Mailath, and Postlewaite (2001a, 2001b) and Felli and Roberts (2016) show how coordination failures can manifest themselves: (1) under-investment equilibria, (2) over-investment equilibria, and (3) mismatch equilibria.

6. Brandenburger and Stuart (2007) call such games, with a non-cooperative game before a cooperative game, "biform games." Such games are grossly understudied.

7. Full appropriation means that each individual's private benefit from any investment coincides with his/her social contribution. Non-complementarity means that different player's investments cannot be complementary.

From all of these papers, one takeaway is always the same: coordination failures *exist* in competitive markets. Further follow up papers, such as Makowski (2004) and most recently Nöldeke and Samuelson (2015) have further generalized results and clarified the connection between competition and efficiency. Makowski (2004) considers a similar environment to mine but focuses on the hold out problem, which I assume away in my problem. I draw most heavily on Nöldeke and Samuelson (2015), who like me, look at the efficiency of coordination in competitive matching markets.[8]

2.2 Refinements for Competitive Economies

None of these papers examine whether these coordination failure equilibria are robust or not. By introducing a refinement, which I draw from an entirely separate literature on adverse selection in competitive markets, I can determine which equilibria are robust. As Gale (1992) points out, in these models, there are many equilibria. However, some of those equilibria are sustained by unreasonable off-equilibrium beliefs, like the belief that other people will not best-respond if a deviation occurs. To discipline off-equilibrium beliefs, Gale uses a form of a trembling-hand refinement (Selten 1975).

There is nothing in general about the refinement that selects for efficient outcomes. Whether the refinement leads to more or less efficient equilibria depends

8. The competitive matching literature that I follow, where no individual chooses prices and equilibrium prices can be thought as coming from a Walrasian auctioneer, is distinct from the competitive *search and matching* literature, following Shimer (1996) and Moen (1997), where one side of the market posts prices.

on the exact context. For example, in Gale (1992), the refined equilibria are inefficient, while in Gale (1996) they are efficient. The usual examples of perfect equilibria actually show that the perfect equilibria are the *inefficient* ones. See Selten's original example (Selten 1975, p. 33) or a textbook example in Maschler, Solan, and Zamir (2013, p. 263). More recently, refinements have been studied in the case of default (e.g. Dubey and Geanakoplos 2002; Dubey, Geanakoplos, and Shubik 2005), and adverse selection in the healthcare market (Scheuer and Smetters 2018).

The two closest papers to mine that incorporate refinements are Zame (2007) and Liu (2018). Zame (2007) considers an extremely general model of firm formation with moral hazard and adverse selection. Instead of allowing inefficiency from adverse selection, as is possible in Zame's model, my model shuts down the adverse selection to focus on the role of coordination, which is not examined in his paper. Liu (2018) studies stability in a two-sided matching model with asymmetric information, such as adverse selection. He also introduces a refinement to discipline off-path beliefs about players' types. In contrast, I focus on off-path beliefs about prices, which depend on the players' actions. Because I am focusing on coordination problems, actions are the key, not types. With these differences in mind, we can now move to the example.

2.3 Example

Consider a simple example of a two-sided matching market with measure one of the agents on both sides. Agents are endowed with a type $t \in T = \{b, s\}$, with measure one of each. For consistent language, I talk about buyers and sellers. The only reason the name matters is to determine who pays a transfer to whom; buyers pay sellers. There are two stages to the game. First, before matching, buyers and sellers must invest in an attribute, $a_b \in A_b = \{0,1\}$ and $a_s \in A_s = \{0,1\}$. The cost to a buyer is $\frac{1}{4}a_b$ and the cost to a seller is $\frac{1}{4}a_s$. These investments generate a surplus for any match: $v(a_b, a_s) = a_b a_s$. Second, after buyers and sellers sink their investment, buyers and sellers enter a competitive market with prices. People purchase the contractual right to join a match (a_b, a_s) but can only do so for the investment levels they bring to the market. Let $p : A_n \times A_s \to \mathbb{R}$, where the price $p(a_b, a_s)$ is a transfer from buyer a_b to seller a_s when they are matched. The final utility of buyers is $v(a_b, a_s) - p(a_b, a_s) - \frac{1}{4}a_b$. For sellers, it is $p(a_b, a_s) - \frac{1}{4}a_s$.

A buyer's problem has two parts. At the second stage, given their investment choice, they choose the optimal seller to match with. The matching market generates indirect utility of indirect utility for a_b:

$$v^*(a_b, p) = \max_{a_s} \{v(a_b, a_s) - p(a_b, a_s)\},$$

and similarly for a seller with a_s:

$$v^*(a_s, p) = \max_{a_b} \{p(a_b, a_s)\}.$$

At the first stage, the choose to choose investment to maximize their expected match surplus in the second stage. While prices are observed in the competitive market, when agents invest, they do not yet observe prices. Therefore, each type t makes investments before seeing prices make their decisions based on some price *conjectures*, $\tilde{p}^t(a_b, a_s)$. Taking as given they choose the optimal contract in the competitive market, a buyer's investment problem is

$$\max_b \left\{ \underbrace{v^*\left(a_b, \tilde{p}^t\right)}_{\text{Utility from Conjectured Optimal Contract}} - \underbrace{\frac{1}{4}a_b}_{\text{Investment Cost}} \right\}.$$

Similarly, the seller's investment problem is $\max_s \left\{ v^*\left(a_s, \tilde{p}^t\right) - \frac{1}{4}a_s \right\}$. The payoffs are given in Figure 2.2. The decision problem is straightforward, except that the prices and conjectures are endogenous, equilibrium objects.

		Buyer Payoffs	Seller Payoffs
Matching Contract	(0,0)	$-\tilde{p}^b(0,0)$	$\tilde{p}^s(0,0)$
	(0,1)	$-\tilde{p}^b(0,1)$	$\tilde{p}^s(0,1) - \frac{1}{4}$
	(1,0)	$-\tilde{p}^b(1,0) - \frac{1}{4}$	$\tilde{p}^s(1,0)$
	(1,1)	$1 - \tilde{p}^b(1,1) - \frac{1}{4}$	$\tilde{p}^s(1,1) - \frac{1}{4}$

Figure 2.2: Example Payoffs

An investment equilibrium[9] is a set of prices, conjectures, and allocations where:

1) each buyer chooses a_b to maximize utility, given her price conjectures: $\tilde{p}^b(a_b, a_s)$,

2) each seller chooses a_s to maximize utility, given her price conjectures: $\tilde{p}^s(a_b, a_s)$, and

3) everyone holds rational conjectures:

9. I take this terminology from Makowski (2004). Besides some technical details, an investment equilibrium is equivalent to what Cole, Mailath, and Postlewaite (2001b) and Nöldeke and Samuelson (2015) call an ex-post equilibrium, where contracting only happens ex-post investment, or what Makowski and Ostroy (1995) call an occupational equilibrium, where people choose an occupation before entering a market.

a) if positive mass of agents choose contract (a_b, a_s), conjectures agree with the posted price: $\tilde{p}^{t_b}(a_b, a_s) = \tilde{p}^{t_s}(a_b, a_s) = p(a_b, a_s)$,

b) otherwise conjectures are not pinned down, and

4) prices clear the matching market.

Because utility is quasi-linear (transferable), a profile of investments and matchings is efficient if and only if it maximizes $v(a_b, a_s) - \frac{1}{4}a_b - \frac{1}{4}a_s$. For this example, the efficient allocation is to maximize investment: $a_b = a_s = 1$.

The efficient allocation with full investment is an equilibrium. To see how, suppose that $p(1,1) = \frac{1}{2}$. All types conjecture that all other prices are zero, however, those prices are not always posted by the auctioneer, since the economy does not include such matches. The matching generates a positive surplus for the buyers and sellers, which is better than all other alternatives which lead to a conjectured utility of zero. This implies that at the investment stage, given conjectures, it is optimal for buyers and sellers to invest, choosing $(1,1)$. Price clears matching markets; all buyers and sellers (each with equal measure) want to match. Finally, the players' conjectures are not contradicted by the data.

There is also an equilibrium where no one invests: $a_b = 0$, $a_s = 0$, and $p(0,0) = 0$. This equilibrium is a coordination failure. This outcome is called a coordination failure because if at least one buyer and one seller could coordinate a joint deviation where they both invest, they both could achieve a higher level of utility.

It is important to note that the coordination failure equilibrium is only sustained by *contradictory* conjectures of buyers and sellers. Returning to the payoffs in Figure 2.2. Fix all conjectures besides $\tilde{p}^t(1,1)$ besides to zero. To prevent any buyer from deviating to $(1,1)$, the conjecture must be sufficiently high: $\tilde{p}^b(1,1) \geq \frac{3}{4}$. In contrast, to prevent any seller from deviating to $(1,1)$, the conjecture must be sufficiently low: $\tilde{p}^s(1,1) \leq \frac{1}{4}$.

Notice that the price conjectures for markets that do not exist in equilibrium—like $\tilde{p}^t(1,1)$ when there is a coordination failure–are a free parameter. They do not need to agree across agents; the theorist is free to pick prices to sustain the allocation. As I will show later, their conjectures *cannot* agree in a coordination failure.

With the free parameter of beliefs, many equilibria can be sustained in general. In this simple example, in particular, notice also that this coordination failure *minimizes* total value. In that sense, there is little predictive power from a Nash-style equilibrium, especially if we are looking to study the level of efficiency. The theory as-is gives us no way to pin down the welfare consequences; the best and worst allocations are both equilibria.

To study the robustness of equilibrium, I consider a mild refinement: trembling hand perfection. For simplicity, first, assume no trembles in the competitive matching market. Second, assume that at the investment stage there is only a simple type of tremble: a uniform trembling hand, where each attribute must be chosen with positive probability $\epsilon^t > 0$ by each buyer and seller. Since there is a continuum of buyers and sellers, I will assume that in the aggregate each attribute must be chosen by a positive mass of players. A perfect investment equilibrium is a price and allocation, such that there exists a sequence of ϵ that goes to zero where the sequence of equilibria converges to the perfect investment equilibrium.

With trembling hand, each a_b and a_s are played, so that the actual prices are pinned down and players cannot have contradictory beliefs: $\tilde{p}^b(a_b, a_s) = \tilde{p}^s(a_b, a_s) = p(a_b, a_s)$. If $p(1,1)$ is low ($< \frac{3}{4}$), no buyers want to choose $a_b = 0$. Instead, they choose $a_b = 1$ as much as possible: $1 - \epsilon^{t_b}$. This drives $p(1,1)$ up. At the same time, if $p(1,1)$ is high ($> \frac{1}{4}$), no sellers want to choose $a_s = 0$. Instead, they choose $a_s = 1$ as much as possible: $1 - \epsilon^{t_s}$. This drives $p(1,1)$ down. Some $p(1,1) \in \left[\frac{1}{4}, \frac{3}{4}\right]$ brings these two forces into balance clears all markets.

For any perturbation, equilibrium requires $p(1,1) \in \left[\frac{1}{4}, \frac{3}{4}\right]$. At that price, everyone wants to invest. As perturbations go to zero, $(1,1)$ is the unique equilibrium strategy profile, even though the prices are not unique in this example.

This example highlights four important features of such markets. First, even under competition, coordination failures can exist. Second, coordination failures are only sustained by off-path conjectures that are contradictory across buyers and sellers. Third, the possibility of mistakes/trembles is one justification to rule out such contradictions. Finally, the possibility of mistakes rules out coordination failures and proves coordination failures are not robust. The next section extends these three features to a more general model with arbitrary (1) finite types of agents, (2) finite investment options, (3) cost of investment, and (4) surplus functions.

2.4 Model

A continuum of agents are endowed with a type $t \in T$, which can be partitioned into "buyers" ($t_b \in T_b$) and "sellers" ($t_s \in T_s$), such that $T = T_b \cup T_s$ and $T_b \cap T_s = \emptyset$. I assume the set of types is finite. An economy is defined by a positive measure on the set of types,

$$E \in M_+(T).$$

There are two stages to the model. In the first stage, each individual must acquire/invest in one *attribute*, $a \in A$. For simplicity, the set of attributes is finite. The attributes are partitioned into those that have a finite cost for buyers, $a_b \in A_b$, and those that have a finite cost for sellers, $a_s \in A_s$, such that $A = A_b \cup A_s$ and $A_b \cap A_s = \emptyset$. There exists a cost function of acquiring an attribute

$$c : T \times A \to \mathbb{R} \cup \infty,$$

so that $c(t, a)$ is the cost of acquiring a for type t. By definition, there is an infinite cost for a buyer type to acquire a seller attribute and vice versa. After attribute investments are made, there is a distribution of attributes $\mu \in M_+(A)$. For any attribute $a \in A$, $\mu(a)$ is the mass of individuals with attribute a. A distribution of attributes is feasible if $\sum_{t_b} E(t_b) = \sum_{a_b} \mu(a_b)$ and $\sum_{t_s} E(t_s) = \sum_{a_s} \mu(a_s)$. Sometimes it will be helpful to work with the distribution of only buyers, $\mu_b \in M_+(A_b)$, or only sellers, $\mu_s \in M_+(A_s)$. We will often be interested in the support for the μ functions, which is just supp $\mu = \{a \in A | \mu(A) > 0\}$.

The second stage involves a people forming matches. To allow individuals to remain unmatched, define $A_b^{\varnothing} \equiv A_b \cup \varnothing$ and $A_s^{\varnothing} \equiv A_s \cup \varnothing$. The value generated by an specific match is given by a bounded value function: $v : A_b^{\varnothing} \times A_s^{\varnothing} \to \mathbb{R}$. In general, I will impose no further assumptions. A matching is a distribution

$$x \in M_+(A_b^{\varnothing} \times A_s^{\varnothing}).$$

A matching x is *feasible* for μ if $x(\varnothing, \varnothing) = 0$, and

$$\sum_{a_s' \in A_s^{\varnothing}} x(a_b, a_s') = \mu(a_b) \quad \forall a_b$$

$$\sum_{b_s' \in A_b^{\varnothing}} x(a_b', a_s) = \mu(a_s) \quad \forall a_s.$$

The second-stage matching/assignment is done through prices. To focus on coordination under competition, I assume each player acts as a price-taker.[10] A price system is $p : A_b^{\varnothing} \times A_s^{\varnothing} \to \mathbb{R}$. People buy the contractual right to participate in a match with some seller with a specific attribute. Neither side can choose the type that they match with. However. this is without loss of generality in terms of payoffs since the value function does not depend on types.[11] A market is open if that pair is part of an equilibrium, that is $x(a_b, a_s) > 0$.

Definition 1. Fixing the distribution of investment, μ, a pair (x, p) is an (ex-post) *competitive equilibrium* for μ if x is feasible for μ, $p(a_b, \varnothing) = p(\varnothing, a_s) \equiv 0$,

10. By assuming price-taking, I follow most of the related matching literature, such as Cole, Mailath, and Postlewaite (2001b) and Nöldeke and Samuelson (2015). See Gretsky, Ostroy, and Zame (1999) and Makowski (2004) for a rigorous analysis of when the price-taking assumption is justified in an assignment model.

11. Alternatively, we can think of a match, (a_b, a_s), as simply a standard good sold by a seller with attribute a_s to a buyer with attribute a_b.

1) for each $a_b \in$ supp μ_b and each $(a_b, a_s^*) \in$ supp x, the match maximizes a_b's utility:

$$a_s^* \in \operatorname{argmax}_{a_s \in \text{ supp } \mu_s} \{v(a_b, a_s) - p(a_b, a_s)\}, \textit{and}$$

2) for each $a_s \in$ supp μ_s and each $(a_b^*, a_s) \in$ supp x, the match maximizes a_s's utility:

$$a_b^* \in \operatorname{argmax}_{a_s \in \text{ supp } \mu_b} \{p(a_b, a_s)\}.$$

The equilibrium requires that when players are deciding whether to form a match given prices, they are optimizing. Nothing says that this maximization must be unique, so players with the same attributes can choose to match with players of different attribute levels. Unlike in the example, conjectures are not a part of a competitive equilibrium; all relevant markets are priced. Even though closed markets are not priced, those markets are irrelevant after investment decisions have been made. Social matching gains function for μ is given by

$$g(\mu) \equiv \max_x \sum_{a_b \in A_b^{\emptyset}} \sum_{a_s \in A_s^{\emptyset}} v(a_b, a_s) x(a_b, a_s) \quad s.t. x \text{ is feasible given } \mu.$$

An allocation that attains $g(\mu)$ is *constrained efficient*. Because of price-taking, we immediately have a "Constrained First Welfare Theorem": If a pair (x, p) is competitive for μ, then it is constrained efficient. It is constrained because maximization only holds within the support of attributes.[12] This immediately rules out any

12. Because of complementarities the equilibrium price is not unique. There is a pie $v(a_b, a_s) = 1$ to divide by $p(a_b, a_s)$. The division which occurs is indeterminate, even though the optimal "quantity traded" is when all buyers and sellers match. This is exactly the setup and outcome in Figure 4 of Smith (1982, p. 171). As Smith finds, even though the number of trades is the efficient and equilibrium amount, the price moves between each round of play.

mismatch equilibria found by Felli and Roberts (2016). Besides that, it is a very weak notion of efficiency. In our example, the equilibrium where no one invests is constrained efficient, even though surplus is minimized. For our current purposes, the result is important because it establishes how the matching market is working effectively, given investments.

Even though all competitive equilibria are constrained efficient, investment coordination failures can still arise such that joint deviations would make everyone better off. In the example, $a_b = 0, a_s = 0$ can be part of an ex-post competitive equilibrium. Even with competition, players are stuck in a coordination failure since a joint deviation to $a_b = 1, a_s = 1$ would make both sides better off. The next subsection asks, given investments are chosen in a non-cooperative setting, do people choose the efficient a_b and a_s?

2.4.1 Investment Equilibrium

Fix the population of types, E. An allocation of attributes is a measure $\nu \in M_+(T \times A)$, where ν_T and ν_A a are the respective marginal distributions. An allocation ν is *feasible* for E if $\nu_T = E$.

Each agent of type t has price conjectures: $\tilde{p}^t : A_b^\varnothing \times A_s^\varnothing \to \mathbb{R}$. For simplicity of notation, I do not allow conjectures to vary for players of the same type. A buyer of type t_b with attribute a_b who conjectures $\tilde{p}^{t_b}$ conjectures an indirect utility from matching of

$$v^* \left(a_b, \tilde{p}^{t_b}\right) \equiv \max_{a_s \in A_s^\varnothing} \left\{ v(a_b, a_s) - \tilde{p}^{t_b}(a_b, a_s) \right\},$$

and a seller of type t_s with attribute a_s who conjectures $\tilde{p}^{t_s}$:

$$v^*\left(a_s, \tilde{p}^{t_s}\right) \equiv \max_{a_b \in A_b^{\varnothing}} \left\{\tilde{p}^{t_s}(a_b, a_s)\right\}.$$

We now have all of the pieces to define the relevant notion of equilibrium.

Definition 2. A tuple $\left(\nu, \left\{\tilde{p}^t\right\}_{t \in T}, p, x\right)$ is (ex-ante) *investment equilibrium* for E if ν is feasible, (x, p) is a competitive equilibrium for ν_A,

1) for all $(t_b, a_b) \in \text{supp}\ \nu$

$$v^*\left(a_b, \tilde{p}^{t_b}\right) - c(t_b, a_b) \geq v^*\left(a_b', \tilde{p}^{t_b}\right) - c(t_b, a_b') \quad \forall\, a_b' \in A_b,$$

2) for all $(t_s, a_s) \in \text{supp}\ \nu$

$$v^*\left(a_s, \tilde{p}^{t_s}\right) - c(t_s, a_s) \geq v^*\left(a_s', \tilde{p}^{t_s}\right) - c(t_s, a_s') \quad \forall\, a_s' \in A_s, \text{ and}$$

3) for all $t \in \text{supp}\ \nu_T$,

$$\tilde{p}^t(a_b, a_s) = p(a_b, a_s) \quad \forall\, (a_b, a_s) \in \text{ supp } \mu_b \times \text{ supp } \mu_s.$$

The first condition is that all type-attribute pairs for the buyers in the support, the attribute maximizes the buyers' utility. Again, nothing forces these to be unique; the same type can choose different attributes. The second condition is the same for the seller and the third condition is that conjectures are rational.

Because the buyer is maximizing over all $a_b \in A_b$, he is allowed to consider deviations outside of the equilibrium support of a_b, and the same holds for sellers. However, because the price conjectures for those deviations are a free-parameter, when constructing an equilibrium we can effectively rule out such deviations by picking appropriately high prices conjectures for buyers and low price conjectures for sellers.

This equilibrium is like a Nash equilibrium, where each player is best-responding, given what everyone else does. However, the equilibrium does not involve the standard Nash equilibrium epistemic justification; people are not best-responding to actions. Instead, in line with the competitive nature of the market, they are best-responding to prices. Each person is choosing her best attribute, given the indirect utility implied by prices and the cost of acquiring that attribute. But beyond the normal conditions for a competitive equilibrium, when players are deciding how much to invest, they must form conjectures about what prices will be in the future. The equilibrium disciplines those conjectures, as Hayek (1937, p. 41) pointed out, "the concept of equilibrium merely means that the foresight of the different members of the society is in a special sense correct." However, the exact meaning of correctness is not clear since some prices never materialize so people can hold contradictory, but in a sense correct, things. I will further discipline conjectures below when I consider refinements to address this issue.

There is a total cost of attributes in the economy ν is $\sum_A \sum_T c(t,a)\nu(t,a)$, and a total surplus from ν,

$$G(\nu) = g(\nu_A) - \sum_A \sum_T c(t,a)\nu(t,a).$$

Definition 3. The allocation ν is unconstrained *efficient* for E if it is feasible and $G(\nu) \geq G(\nu')$ for all other feasible allocation ν'.

The previous literature has documented the existence of efficient equilibria in similar matching models, *e.g.* Nöldeke and Samuelson (2015, Corollary 1, p. 858) and Dizdar (2018, Proposition 2, p. 98). With the above definitions in order, we can immediately show the same in this environment.

Proposition 1. *For any economy, there exists an unconstrained efficient investment equilibrium.*

Proof. The existence proof is by construction. To construct the equilibrium, assume conjectures are consistent: $\tilde{p}^t(a_b, a_s) = p(a_b, a_s)$ for all $t \in \text{supp } \nu_T$. Then we can write down the welfare as

$$\sum_{t_b} \left[\max_{a_b, a_s} v(a_b, a_s) - p(a_b, a_s) - c\left(t_b, a_b\right) \right] E(t_b) + \sum_{t_s} \left[\max_{a_b, a_s} p(a_b, a_s) - c\left(t_s, a_s\right) \right] E(t_s).$$

For any match that occurs, the price to the seller equals the price to the buyer. For anyone that is unmatched, the price is normalized to zero. Since everyone is maximizing given their conjectures and, by assumption, the price conjectures are the same, this means that the price drop out of the overall welfare.

$$\max_{a_b, a_s} \left\{ \sum_{t_b} \left[v(a_b, a_s) - p(a_b, a_s) - c\left(t_b, a_b\right) \right] E(t_b) + \sum_{t_s} \left[p(a_b, a_s) - c\left(t_s, a_s\right) \right] E(t_s) \right\}.$$

Therefore, the sum of everyone's individual maximization is identical to an overall maximization of welfare. ∎

The proof relies on all players having consistent conjectures. If conjectures are consistent, it is as if markets open in the first stage, since everyone is optimizing given the same prices. With open markets in the first period, efficiency comes directly from the First Welfare Theorem, which holds even with increasing returns in models such as Acemoglu (1996).[13]

But we also know that not all investment equilibria are efficient. Returning to the example, we already showed that $a_b = 0$ for all buyers, $a_s = 0$ for all sellers, and $p(a_b, a_s) = 0$ can be sustained as an investment equilibrium with certain contradictory conjectures. Moreover, this is the worst possible outcome; it minimizes the surplus. The next subsection shows that this type of surplus minimizing equilibria exist for many economies that are relevant in the matching literature.

2.4.2 Weak Predictions with Unconstrained Beliefs

One problem with the equilibrium concept used in the literature, and why it leads to so many different equilibria as shown in the last section, is that off-path beliefs are a free parameter for the theorist. As Robert Lucas taught us, "beware of theorists bearing free parameters." In related papers of adverse selection mentioned above, economists have recognized this issue in other competitive contexts. For example, Zame (2007) notes that "imposing no discipline would admit equilibria which are *viable only because different agents hold contradictory beliefs*." The same is

13. In a footnote, Acemoglu (1996, p. 785, fn. 7) discusses the connection between markets that open in the first period (as in his paper) and endogenous markets that open in the second period (as in this chapter). He suggests in the context of a specific example, a refinement rules out the inefficient equilibria even with endogenous markets. Theorem 3 formalizes and generalizes this point. Thanks to Georg Nöldeke for pointing this out.

true in this model. When the equilibrium concept allows agents to hold contradictory beliefs, many equilibria can be sustained.

To show just how weak the solution concept is, in this section, instead of focusing on the most general forms of the surplus and cost functions that we have used so far, let us consider a smaller set that is still relevant for models of investment and matching.

I will consider two definitions:

Definition 4. A cost function has costly investment if there exists an attribute, $0 \in A$, such that for all types t and for all $a \neq 0$, $c(t,a) \geq c(t,0) = 0$

Definition 5. Investment is mutually-necessary if surplus is zero unless both sides invest: $v(0,\emptyset) = v(\emptyset,0) = v(a_b,0) = v(0,a_s) = 0$ for all a_b, and a_s and $v(a_b,a_s) \geq 0$ for all a_b and a_s.

These are strong restrictions on the cost and surplus functions, but they include economies that are relevant for any researcher who is looking at the interaction of investment with matching.

Definition 6. An allocation is individually rational if every type receives weakly positive utility.

The following proposition shows that for all economies like this, the surplus-minimizing outcome (out of those that are individually rational) is an equilibrium.

Proposition 2. *For any economy with costly and necessary investment, there exists an investment equilibrium where the total surplus is zero.*

The proof is immediate and highlights the Nash-style equilibrium.

Proof. Suppose all types but t_b are not investing. The two cases are matching with a seller with $a_s = 0$ or not matching. Because conjectures are a free parameter for those two types of deviations, I can choose them to ensure no deviation. For the case of matching with an $a_s = 0$ type, pick conjectures so that:

$$\begin{aligned} v(0,0) - p(0,0) - c\left(t_b, 0\right) &\geq v(a_b, 0) - \tilde{p}^{t_b}(a_b, 0) - c\left(t_b, 0\right) \quad \forall a_b \\ 0 - 0 - 0 &\geq v(a_b, 0) - \tilde{p}^{t_b}(a_b, 0) - c\left(t_b, a_b\right) \quad \forall a_b \\ \tilde{p}^{t_b}(a_b, 0) &\geq v(a_b, 0) - c\left(t_b, a_b\right) \quad \forall a_b. \end{aligned}$$

The same process can be done for not matching. Therefore, we have found conjectures so that it is optimal for t_b to not invest, but t_b was arbitrary so the same process can be used for all types.

For sellers, choose price conjectures of zero and complete the same process. Now we have found conjectures so that no types (buyers or sellers) want to invest. Therefore, everyone receives a utility of zero, which is the minimum possible utility profile that is individually rational. ■

Once written out the proposition is so obvious that it seems not worth mentioning. I include it simply to show that for a reasonable class of models of investment and matching, the equilibrium concept allows the worst-case outcome—out of those that are individually rational so that everyone receives non-negative utility—and the best-case outcome.

The proposition holds regardless of the shape of the cost and surplus functions. Even if the cost of investment is arbitrarily small and the surplus generation is arbitrarily big, there exists an investment equilibrium with zero total surplus. In this case, we still cannot rule out that either the best or the worst possible allocation can occur.[14] For doing welfare analysis though, it may be desirable to say more than "either the best or worst outcome can occur."

The proposition also means that any economy that rules out the surplus minimizing outcome does so because of decisions made that ignore the matching process; people invest regardless of the matching market. In that case, we can rule out the worst outcomes, but it does not have anything to do with the matching market.

To discipline the set of possible outcomes, I follow Gale (1992), who argued that "*some refinement of the equilibrium concept is required to give the theory predictive power*. One such refinement is based on the notion of the 'trembling' hand." The next section shows the under such a refinement, all equilibria are efficient.

2.5 Disciplined Beliefs and Perfect Equilibrium

To discipline believes, I will consider a perturbed strategy vector for all buyers. For simplicity of notation when there is a continuum of agents, I assume that all buyers are subjected to the same tremble, $\epsilon_{A_b} = (\epsilon(a_b))_{a_b \in A_b}$, satisfying $\epsilon(a_b) > 0$ for all $a_b \in A_b$ and

$$\sum_{A_b} \epsilon(a_b) \leq 1.$$

14. If we introduced random actions, we can say that anything in-between could happen too.

Similarly, all sellers are subjected to the same tremble, $\epsilon_{A_s} = (\epsilon(a_s))_{a_s \in A_s}$, satisfying $\epsilon(a_s) > 0$ for all $a_s \in A_s$ and

$$\sum_{A_s} \epsilon(a_s) \leq 1.$$

A perturbed game is indexed by the vector of perturbed strategies $\epsilon = (\epsilon_{A_b}, \epsilon_{A_s})$. An allocation $\nu(\epsilon)$ is ϵ-feasible for E if $\nu_T = E$ and for all $a \in A$

$$\nu_A(\epsilon(a)) \geq \epsilon(a).$$

Instead of jumping directly to the analysis of the limit of perturbed games, it is helpful to say something about the perturbed games themselves. In particular, we can consider their respective efficiency. To do so, let us say that an allocation $\nu(\epsilon)$ is ϵ-efficient for E if it is feasible and $G(\nu(\epsilon)) \geq G(\nu'(\epsilon))$ for all other ϵ-feasible allocation ν'. Formally,

Definition 7. A tuple $\left(\nu(\epsilon), \left\{\tilde{p}^t\right\}_{t\in T}, p, x\right)$ is an ϵ-investment equilibrium for E if ν is ϵ-feasible, p is a competitive price for ν_A, conjectures are rational, and for all (t, a) such that $\nu_A(\epsilon) > \epsilon$,

$$v_a^*(\tilde{p}^t) - c(t,a) \geq v_{a'}^*(\tilde{p}^t) - c(t,a') \quad \forall a' \in A.$$

Note that by construction, with a trembling hand, $\operatorname{supp} \nu_A(\epsilon) = A$. Because there is full support and all markets are open, coordination failures cannot arise. This is shown through the following lemma.

Lemma 1. *If $(\nu(\epsilon), p)$ is an ϵ-investment equilibrium, then it is ϵ-efficient.*

Proof. Let $Q(\epsilon)$ be the utility generated by the trembling actions

$$Q(\epsilon) = \sum_{A_b} \left\{ \sum_{t_b} \left[v^* \left(a_b, \tilde{p}^{t_b} \right) - c(t_b, a_b) \right] \nu_T(t_b) \right\} \epsilon(a_b) + \sum_{A_s} \left\{ \sum_{t_s} \left[v^* \left(a_s, \tilde{p}^{t_s} \right) - c(t_s, a_s) \right] \nu_T(t_s) \right\} \epsilon(a_s)$$

Then the total utility can be written as:

$$\sum_{t_b} \left[\max_{a_b, a_s} v(a_b, a_s) - \tilde{p}^{t_b}(a_b, a_s) - c\left(t_b, a_b\right) \right] \nu_T(t_b) \left(1 - \epsilon(a_b)\right) + \sum_{t_s} \left[\max_{a_b, a_s} \tilde{p}^{t_s}(a_b, a_s) - c\left(t_s, a_s\right) \right] \nu_T(t_s) \left(1 - \epsilon(a_s)\right) + \underbrace{Q(\epsilon).}_{\text{Constrained Choice}} \quad .$$

But since all actions are played by trembles, $\tilde{p}^{t_b}(a_b, a_s) = \tilde{p}^{t_s}(a_b, a_s)$ for all (a_b, a_s), not just those in a subset. Therefore agents optimize the entire left expression. The rest follows the proof of Proposition 1. ■

The possibility of mistakes rules out coordination failures. Lemma 1 is closely connected to Proposition 3.1 in Nöldeke and Samuelson (2014, p. 860), which proves that with complete prices, the matching is efficient. The reason is standard; for any positive tremble, all actions are played and so all markets are open. With complete markets, all possible trades are price and any competitive equilibria are efficient. except for the lemma does not require that all individuals are matched. Now we can consider the limit of trembles, which need not have complete prices.

A tuple $\left(\nu, \left\{\tilde{p}^t\right\}_{t\in T}, p, x\right)$ is a *perfect investment equilibria* if there exists a sequence of ϵ, such that $\lim_{k\to\infty} M(\epsilon^k) = 0$ such that $\left(\nu(\epsilon), \left\{\tilde{p}^t\right\}_{t\in T}, p, x\right) \to \left(\nu, \left\{\tilde{p}^t\right\}_{t\in T}, p, x\right)$.

Theorem 3. *If (ν, p) is a perfect investment equilibrium, then it is efficient.*

Proof. The theorem is immediate from Lemma 1 since $\underbrace{Q(\epsilon)}_{\text{Constrained Choice}} \to 0$. ■

Any perfect investment equilibrium is efficient, even though markets are 1) endogenous, because they depend on investment choices, and 2) incomplete, because investments must be sunk and not all markets are priced. It is a modified First Welfare Theorem. It is weaker than the standard version because it only applied to perfect equilibria, but it is stronger because it does not require complete markets. Still, the logic is directly from the First Welfare Theorem, since people make decisions as if markets are complete.

The theory's predictive power comes from imposing more restrictions on beliefs than just rational conjectures. The mathematical mechanism is that the mistakes caused by trembles generate complete markets. The trembling with a large number of agents rules out contradictory beliefs, as in Zame (2007), and ensures "price consistency", as in Makowski and Ostroy (1995). However, instead of assuming price consistency, the tremble gives a justification for sure price consistency in terms of the robustness and stability of the equilibria considered.

There are other justifications for non-contradictory beliefs. For example, Dubey and Geanakoplos (2002) consider a fictitious seller who contributes an infinitesimal to each health insurance pool. Dubey, Geanakoplos, and Shubik (2005) assume that the government intervenes to sell infinitesimal quantities of each asset and fully delivers on its promises. In both cases, since all markets are open, all markets have public prices, and everyone's price conjectures agree in equilibrium.

2.6 Conclusion and Implications

In this chapter, I argue against focusing too much on coordination failures when there is competition. Those coordination failures, emphasized by the previous literature, rely on using beliefs as a free parameter and constructing overly pessimistic conjectures. With the free parameter, there are many equilibria. If we want predictive power, we must use a refinement, such a trembling hand perfection.

When we consider perfect equilibrium in a competitive matching model with investment, every perfect equilibrium is efficient. If we are interested in the efficiency properties of only those competitive equilibria which are robust, then Theorem 3 strengthens the standard First Welfare Theorem because it proves the efficiency of competitive markets, even with incomplete markets of the type studied.

The theorem does not imply that coordination problems cannot exist, just that those in competitive environments are not stable and therefore are unlikely to last. We should instead think of them as arising in environments with imperfect competition. If we observe an outcome that looks like a coordination failure, we have a reason to look for the imperfections of *competition* in the market. When looking for how to use policy, solving imperfections of competition should solve the coordination problem.

Chapter 3

On the Informational Efficiency of Decentralized Price Formation (with Rafael R. Guthmann)

> "To gain an advantage from better knowledge of facilities of communication or transport is sometimes regarded as almost dishonest, although it is quite as important that society make use of the best opportunities in this respect as in using the latest scientific discoveries."
>
> F. A. Hayek (1945)

At least since Hayek (1937, 1945), economists have assessed institutions by their relative ability to incorporate and communicate information. According to this framing, markets are a desirable institution because they achieve an efficient allocation of resources through decentralized decision making where each decision-maker only needs to be aware of market-prices besides their individual endowments and technology: That is, markets achieve allocative efficiency without requiring the explicit communication of the information that is dispersed through the economy.

Studies such as Mount and Reiter (1974), Hurwicz (1977a, 1977b, 1977c), Jordan (1982), and Chander (1983) developed a formal concept of informational efficiency. Mount and Reiter (1974) shows that the competitive equilibrium is informationally efficient in the sense that competitive prices communicate the minimum amount of information necessary to implement a Pareto efficient allocation in an environment where information is dispersed.[1] Moreover, Jordan (1982) proves that competitive prices are the unique decentralized mechanism that achieves this informational efficiency and satisfies the individual rationality constraint that the utility in participating in the mechanism is at least as high as their initial endowment.[2]

While these informational efficiency results are important, for other reasons, the competitive mechanism is not without its critics. Particularly, the concept of competitive equilibrium as an allocation mechanism lacks strategic foundations

1. Such as information regarding individual endowments, preferences, or technology.

2. Further illustrating the informational power of competitive equilibrium, even when types are private information, Hammond (1979) has shown in an economy with a continuum of agents that an incentive-compatible mechanism that implements a Pareto optimal allocation must give rise to the same net trades as the competitive equilibrium mechanism.

(*e.g.* Gale 2000): decision-makers take prices as given and the price is determined by the underlying market data through an exogenously imposed market-clearing condition: the strategic process that determines market prices is not modeled. That is, only the quantities are determined by the decentralized choices of decision-makers while the prices are determined by the "Walrasian auctioneer" or "the economist" who solves the model by equating the quantity supplied with the quantity demanded. Without any way to measure the information required to reach the competitive allocation through a fully decentralized strategic procedure, it is unclear how seriously economists should take these informational efficiency results.

While there are different strategic foundations for competitive equilibrium, we show that the main one in the literature—random search and bargaining—is unattractive from an informational efficiency perspective. In particular, we prove that for large economies the search allocation mechanism requires infinitely more information than the competitive mechanism. The reason is that in a search economy, each buyer needs to know about the price he or she would trade with each seller, which requires a different message. If there are k buyers and k sellers who trade one good, the message space must include k^2 prices. This seems like an unattractive feature of a search model for large economies.

Therefore, we introduce a new strategic foundation for competitive equilibrium which requires less information but qualitatively replicates several empirical observations. Following the intermediation models of Spulber (1996) and Rust and Hall (2003) we introduce a model of market-makers who intermediate trade between buyers and sellers in a model of dynamic price formation. We show that a model of exchange based on intermediators can approximate the informational efficiency properties of the competitive equilibrium while endogenizing the equilibrium prices as strategic choices of decision-makers. The informational efficiency is obtained through an "intellectual division of labor" between the consumers in the economy who "outsource" the price formation process to the market-makers who are the strategic players, who in turn make profits by extracting part of the surplus of the gains from trade. These market-makers can be thought of as arbitrageurs who have the ability to buy low and sell high, exploiting opportunities that other actors are unaware of.[3] As frictions decrease the equilibrium approximates the no-arbitrage limit which is the competitive equilibrium, which is well understood in the case of product markets (*e.g.* Makowski and Ostroy 1998) and financial markets (*e.g.* Werner 1987).

Among informationally decentralized allocation mechanisms where terms of trade are set strategically, this study suggests that the one that is the most informationally efficient is the allocation mechanism implied by a situation where the market for each good is monopolized by a single market-maker who sets prices, that way consumers in the economy only need to be aware of one price for each good, mimicking the competitive equilibrium. Allowing for entry puts pressure on the market-maker to lower prices to deter entry in what we call a monopolistic deterrence equilibrium. Taking discount rates to zero implies that the allocation

3. Rubinstein and Wolinsky (1987), Duffie, Gârleanu, and Pedersen (2005), Nosal, Wong, and Wright (2019) are other examples of models that feature middlemen or market-makers that intermediate trade. In addition, there is a literature following Kirzner (1973) which would call these market-makers "entrepreneurs" that are "alert" to profit opportunities that exist in the market at a point in time.

corresponding to the monopolistic deterrence equilibrium converges to the competitive allocation.[4]

The main objectives of this chapter are as follows:

1. Construct a model of an economy that features strategic and decentralized trading that attains as its "frictionless limit" the allocation corresponding to the competitive equilibrium while allowing for asymmetric information.

2. Show that the equilibrium of this economy can replicate several of the features reported in the empirical literature that cannot exist in competitive equilibrium.

3. In addition, this model of an economy satisfies the condition that the amount of information that needs to be communicated does not diverge from the competitive allocation mechanism.

4. Find what is the minimum amount of information required by a strategic model of price determination that approximates the competitive allocation.

Strategic Foundations of Competitive Equilibrium

There is a large literature on strategic foundations for competitive equilibrium which justifies the assumption that markets can be described by the model of competitive equilibrium by showing that under a variety of different conditions that the strategic equilibrium of decentralized economies (modeled as the core of a

4. This study implies that the presence of platforms such as bitcoin exchanges, Amazon, and Uber, with large market-shares might be more informationally efficient than industries with many agents on both sides of the market.

cooperative game or random matching and bargaining games) does generate the same allocation as the competitive equilibrium.[5] Also, the literature on random matching and bargaining games is also able to explain several deviations of competitive conditions verified empirically (such as price dispersion is generated as an equilibrium outcome in models such as Mortensen and Wright (2002) while accommodating the competitive equilibrium as a special case when frictions of trading are zero. However, this literature has assumed that decision-makers have complete information regarding market conditions.

More recent papers have shown, however, in the case of a market for a single good that asymmetric information regarding private valuations (Satterthwaite and Shneyerov 2007, 2008) and imperfect information regarding aggregate market conditions (Lauermann, Merzyn, and Virág 2018) can be shown to be consistent with decentralized random matching games that attain competitive equilibrium allocations. However, these studies assume that decision-makers have full awareness of the data of the economy: there exists only uncertainty regarding the "quality" of the information (that is, the values of individual parameters of the economy which are private valuations in the case of Satterthwaite and Shneyerov (2007, 2008) and whether the supply of a good is greater than demand, so the competitive equilibrium price is 0 or 1 (Lauermann, Merzyn, and Virág 2018)) but there are no constraints on the "quantity" of information regarding the economy that can be utilized by each decision-maker.

5. Starting in cooperative game theory with Edgeworth (1881), Debreu and Scarf (1963), and Aumann (1964) and moving into non-cooperative random matching game-theoretic arguments (Gale 1986a, 1986b, 1987, 2000; Osbourne and Rubinstein 1990; McLennan and Sonneschein 1991; Mortensen and Wright 2002; Lauermann 2013).

The literature on informational efficiency argues that the competitive equilibrium mechanism allows an economy to achieve an efficient allocation while minimizing the amount of information that needs to be communicated by the individual decision-makers. That is, in a competitive equilibrium individual decision-makers only need to be aware of the minimum amount of parameters regarding the state of the market to achieve an efficient allocation. The decision-makers can be simply completely unaware of the rest of the economy besides market prices (unawareness being interpreted as not "generally taking into account", "being present in mind" (Modica and Rustichini 1999, p. 274), "thinking about" (Dekel, Lipman, and Rustichini 1998) or "paying attention to" (Schipper 2014)) - and still achieve the same allocation which is obtained by a central planner in possession of all the data in the economy who is maximizing a social welfare function. If the allocation corresponding to a competitive equilibrium is reached through a distinct mechanism (such as in a Gale-type random matching and bargaining model) then this informational efficiency property does not hold and we show that in the case of the typical random search model the implied allocation mechanism is infinitely less informationally efficient than the competitive mechanism.

In contrast, an economy with market-makers can require little information. In particular, this mechanism has a message space that has one additional dimension over the competitive mechanism for each good in the economy besides the numeraire good. This extra-dimension represents the profit margin between the bid and ask price that the strategic market-maker obtains from setting prices. Another issue with the concept of competitive equilibrium is empirical: most markets exhibit features that are very different from the idealized features of competitive equilibrium. For example, there is substantial price dispersion for individual goods instead of a single price (Sorensen 2000; Kaplan et al. 2019). Individual plants face plant-specific demand instead of taking market prices as given, and this plant-specific demand depends on the history of plant activity and grows over time (Foster, Haltiwanger, and Syverson 2016). Both of these features can be replicated by the market-maker model presented here. Informational efficiency and the empirical reality are both reasons to further study models with market-makers.

3.1 Physical Environment and Competitive Mechanism

Consider a class of environments E where there are $L = 2$ goods: an indivisible good and a divisible numeraire good. There is a continuum of consumers with measure normalized to 1 and indexed by i. For an environment $e \in E$ a fraction $s \in [0, 1]$ are sellers (set denoted $S = [0, s]$) who are willing to supply a unit of the indivisible good for a higher price (in the numeraire good) than their cost c^i. A fraction $b = 1 - s$ are buyers (set denoted $B = (s, 1]$) who are willing to buy a unit of the indivisible good for a lower price (in the numeraire good) than their valuation v^i. Valuations for the good are distributed according to cumulative distribution functions G and F for sellers and buyers, respectively which are strictly increasing and differentiable. The CDFs G and F have supports

$[\underline{c}, \overline{c}] \subset \mathcal{R}_{++}, [\underline{v}, \overline{v}] \subset \mathcal{R}_{++}$, respectively, where $\overline{v} > \underline{c}$ (so that there is possibility for mutually beneficial trade and the market does not shut down).

Let E^i be the set of possible types for consumer i which indicates whether i is a buyer or a seller and i's valuation. A realization of $e \in E$ is an environment $e = (e^i)_{i \in [0,1]}$, where $e^i \in E^i$ with $e^i = c^i$ if $i \in S$ and $e^i = v^i$ if $i \in B$ which are distributed according to the CDF's G and F.

Let Y be the set of net-trades which satisfy feasibility conditions: $y_1(i) \in \{-1, 0, 1\}$, $y_2(i) \in \mathcal{R}$, $\int_0^1 y_1(i) = 0$, [6] and $\int_0^1 y_2(i) = 0$. The last two conditions mean that net trades of both the indivisible and numeraire good must add up to 0. Formally, we can define the set of feasible allocations as

$$Y = \left\{ y : [0,1] \to \{-1,0,1\} \times \mathcal{R} : \int_0^1 (y_1(i), y_2(i)) di = (0,0) \right\},$$

To formalize the notion of the informational size, we need to define messages. The set M is an abstract message space. The non-empty valued correspondence $\mu : E \rightrightarrows M$ specifies the messages for each environment. Finally, the outcome function $g : M \to Y$ maps messages to trades/allocations.

We are interested in allocation mechanisms that are informationally decentralized, which are mechanisms that feature a message process (μ, M) that is privacy-preserving.

Definition 1. A message process (μ, M) is *privacy-preserving* if there exists a correspondence $\mu^i : E^i \rightrightarrows M$ for each i such that for each $e \in E$ $\mu(e) = \cap_{i \in [0,1]} \mu^i(e^i)$.

6. That is, the agent buys one unit of the indivisible good,sells a unit or does not change his or her endowment.

Putting this together, we can define an allocation mechanism, which is the object of interest.

Definition 2. An *allocation mechanism* is a triple (μ, M, g).

3.1.1 Competitive Mechanism

Let $M_c = \{(p, y) \in \mathcal{R}_{++} \times Y : py_1(i) + y_2(i) = 0, \forall i\}$. For an environment $e \in E$, each $i \in S \cup B$ define the correspondence $\mu_c^i : E \rightrightarrows M_c$ by

$$\mu_c^i(e^i) = \left\{ (p, y) \in M_c : y(i) = \begin{cases} (1, -p) & \text{if } i \in B \text{ and } v^i \geq p \\ (0, 0) & \text{if } i \in B \text{ and } v^i < p \text{ or } i \in S \text{ and } c^i > p \\ (-1, p) & \text{if } i \in S \text{ and } c^i < p \end{cases} \right\},$$

that is, a buyer purchases the good for p if his valuation is higher and a seller sells the good for p is her cost is lower. Define

$$\mu_c(e) = \cap_{i \in [0,1]} \mu_c^i(e^i).$$

Then (μ_c, M_c) is the competitive message process and (μ_c, M_c, g_c) is the competitive allocation mechanism, where $g_c : M_c \rightarrow Y$ is the outcome function given by $g_c(p, y) = y$.

Remark 1. Note that to satisfy feasibility a competitive equilibrium price p^* satisfies $sG(p^*) = b[1 - F(p^*)]$, since G and F are continuous and strictly increasing the competitive equilibrium price p^* is unique. The competitive mechanism is also privacy preserving by construction.

3.2 The Informational *In*efficiency of the Search Mechanism

3.2.1 Environment

In a search environment, there is an entry rate of potential $s > 0$ sellers and $b = 1 - s > 0$ buyers. Given stocks of B buyers and S sellers currently in the market, buyers and sellers meet at the rate $M(B, S)$. Let the buyer/seller ratio $\theta = B/S$ be the market tightness parameter, and $m(\theta) = M(B, S)/S$, be the rate a seller meets buyers and $m(\theta)/\theta$ be the rate a buyer meets sellers. There are discount rates and search costs for buyers and sellers, respectively given by $(r, c_b, c_s) \in \mathcal{R}_+$.

When buyers and sellers meet if there is a positive surplus in trading they trade at terms determined by the generalized Nash solution over the joint surplus. Sellers' bargaining power is given by the parameter $\beta \in (0, 1)$ (buyers' bargaining power is then $1 - \beta$).

For this search environment to be directly comparable to the competitive mechanism the flows of buyers and sellers exiting the market should be equal to the flows entering which are s and b, so that the corresponding net trades in the competitive equilibrium have an analogous implementation in this environment. The search equilibrium where the rate of entry of new buyers and sellers in the market is the same as the exit rate is called a steady-state search equilibrium. Given a pair of marginal types of buyers and sellers (R_b, R_s), where buyer with valuation $x > R_b$ enters and seller with cost $y < R_s$ enters, who are indifferent between participating in the market or not in steady-state search equilibrium this pair satisfies the condition $sG(R_s) = b[1 - F(R_b)]$, and the distribution of participating types is constant.

As described in Mortensen and Wright (2002), these parameters determine the steady-state search equilibrium which is characterized by $(V_b, V_s, R_b, R_s, \Phi, \Gamma)$, the value functions (V_b, V_s), cutoff valuations, and costs to participate in the market (R_b, R_s), and the distributions of participating types (Φ, Γ) of buyers and sellers, respectively. Transaction prices between buyer with valuation x and seller with cost y satisfy

$$p(x,y) = y + V_s(y) + \beta[x - y - V_b(x) - V_s(y)]. \tag{3.1}$$

Mortensen and Wright (2002) show that if search costs c_b, c_s are strictly positive and r is lower than some threshold $\hat{r} > 0$ then all meetings result in trade. This implies that steady-state equilibrium distribution of operating types (Φ, Γ) is given by the densities of (F, G) on the types who participate ($v \geq R_b, c \leq R_s$). For $r \in (0, \hat{r})$, then all meetings result in trade and there is price dispersion. In this case the equilibrium price for the transaction with seller with cost y and buyer with valuation x satisfies

$$p(x,y) = \beta\left[\frac{r\theta x + (1-\beta)m(\theta)R_b}{r\theta + (1-\beta)m(\theta)}\right] + (1-\beta)\left[\frac{ry + \beta m(\theta)R_s}{\beta m(\theta) + r}\right]. \tag{3.2}$$

In the case where the common discount rate is zero, $r = 0$, then the Law of One Price holds and there is an equilibrium price

$$\hat{p} = \beta R_b + (1-\beta)R_s. \tag{3.3}$$

If search costs (c_b, c_s) converge to zero then $\hat{p}$ converges to the competitive price and R_s, R_b both converge to the same value R which is the competitive equilibrium price p^* and the search equilibrium allocation in terms of quantity also converges to $sG(p^*)$ which is the quantity sold in competitive equilibrium.

3.2.2 Steady-State

Throughout we focus on the steady-state of the search mechanism since that is where the search mechanism involves the least information, providing a best-case scenario for the search mechanism. Note that in steady-state constant distribution of types currently in the market implies that the distribution of types leaving the market is the same as the distribution of types entering the market, which is given by (F, G) with the cutoffs (R_b, R_s), the allocation in the steady-state can be described by a pair (p_s, y) where $p_s : [0, 1] \to \mathcal{R}_+$ is a function, where $p_s(i)$ for describes the equilibrium transaction price for agent i if i participates in the market that is, if $i \in B$,$x^i \in [R_b, 1]$ and if $i \in S, y^i \in [0, R_s]$. If i does not participate then if i is a seller for $c^i \in (R_s, \overline{y}]$, $p_s(i) = R_s$ and if i is a buyer, $x^i \in [\underline{x}, R_b)$, then $p_s(i) = R_b$.

Since buyers and sellers meet randomly and the transaction price depends on the pair of valuations of buyers and sellers $p(x, y)$, prices are not deterministic in the search equilibrium but the distribution of realized transaction prices is deterministic as there is a continuum of traders and can be described by a CDF $P : [\underline{p}, \overline{p}] \to [0, 1]$. Any function p_s consistent with the search equilibrium implies in an equilibrium distribution of prices P. Let $\overline{y}(x)$ and $\underline{x}(y)$ be the highest seller's cost and lower buyer's valuation such that there is positive joint surplus in trading given buyer's and seller's valuations (x, y), respectively, then p satisfies $p_s(i) \in \{p(x^i, y), y \in [\underline{y}, \overline{y}(x^i)]\}$ if i is a buyer and $p_s(i) \in \{p(\underline{x}, y^i), x \in [\underline{x}(y^i), \overline{x}]\}$ if i is a seller.

The privacy-preserving message process (μ_s, M_s) is constructed as follows:

The message space of the search mechanism is

$$M_s = \{(p_s, y) \in \mathcal{F} \times Y : p_s(i)y_1(i) + y_2(i) = 0, \forall i\},$$

where $\mathcal{F}$ is the space of functions on $[0, 1]$ to $\mathcal{R}_{++}$.

Let μ_s^i be a correspondence from E^i to M_s. Let $\mu_s^i : E^i \rightrightarrows M_s$ given by

$$\mu_s^i(x^i) = \left\{ (p_s, y) \in M_s : y(i) = \begin{cases} (1, -p_s(i)) & \text{if } i \in B \text{ and } x^i \geq p_s(i) \\ (0, 0) & \text{if } i \in B \text{ and } x^i < p_s(i) \\ (-1, p_s(i)) & \text{if } i \in S \text{ and } x^i > p_s(i) \end{cases} \text{or } i \in S \text{ and } x^i > p_s(i) \right\}.$$

Define the correspondence $\mu_s : E \rightrightarrows M_s$ by

$$\mu_s(e) = \cap_i \mu^i(e^i) \cap (p_s(e) \times Y), \tag{3.4}$$

where $p_s(e)$ is the pricing function determined by the search equilibrium in the environment e (with buyers and sellers types distributed according to F and G) so that means that μ_s is restricted to the subset of M_s consistent with the search environment described in subsection 3.2.1. Note that μ_s is privacy-preserving by construction.

The search mechanism is a triple (μ_s, M_s, g_s) where $g_s(p, y) = y$ is a projection from M_s to Y.

Note that $p_s(i) > R_s$ if $c^i \leq R_s$ and $p_s(i) < R_b$ if $v^i \geq R_b$ since prices must compensate for search costs, while agents who do not trade are the types with costs/valuations in (R_s, R_b).

Remark 2. If the discount rate $r = 0$ and search costs (c_b, c_s) converge to zero then the search equilibrium prices all converge to p^* which means that the search mechanism becomes the competitive mechanism. Then, clearly, it is informationally efficient at this frictionless limit.

3.2.3 Informational Efficiency

In the environment has a continuum of agents and smooth distributions of types any allocation $y \in Y$ is an infinite-dimensional object. To articulate the argument of the size of information messages in the terms of Hurwicz (1977b) and Jordan (1982) on the dimensional size of message space which are finite-dimensional manifolds, we study here sequences of environments with finitely many types of buyers and sellers and as the number of types grows to infinity, the distributions of types approximate the continuous distributions of buyers F and sellers G.

Let $\{e_k\}_{k\geq 2}$ be a sequence of environments where buyers and seller types are distributed according to $\{F_k, G_k\}$, sequences of step-functions. A pair (F_k, G_k) that represent the cumulative distributions of types of buyers and sellers, respectively, in an environment e_k with k types of buyers and k types of sellers, each type of measure $1/k$ (that is, the sets of buyers and sellers are partitioned into subsets of the same measure whose elements are all identical). The pair of sequences

$\{F_k, G_k\}_k$ converges to F and G, respectively. We call an environment with k types of buyers and sellers a k-environment. In addition F_k, G_k are such that p^* is consistent with competitive equilibrium in economy e_k.

Then the allocation mechanisms can be written in terms of types. Let $b \in \{1, \ldots, k\}$ index buyer types and $v \in \{1, \ldots, k\}$ seller types, let $x(b)$ be the valuation of a buyer of type b and $y(v)$ be the cost of a seller of type v.

Let $y^B(b) = (y^B(b)_1, y^B(b)_2)$ be the net trades for a buyer of type b in the competitive equilibrium and $y^S(v) = (y^S(v)_1, y^S(v)_2)$ the net trades for a seller of type v. A profile of net trades specifies a net trade for each of the $2k$-types of traders: $y = \left((y^B(b))_{b=1}^k, (y^S(v))_{v=1}^k\right)$. The set of net trades for the competitive mechanism is then

$$Y_k^c = \left\{ y \in ((-1, 0, 1) \times \mathcal{R}_+)^{2k} : \sum_{b=1}^{k} y^B(b) + \sum_{v=1}^{k} y^S(v) = (0, 0) \right\}.$$

as there are k types of buyers and k types of sellers. The competitive message space in the k-environment specifies a price of p and allocation y where $y \in Y_k^c$.

For the search mechanism, the set of net trades is a higher dimensional object. To see that consider the search equilibrium. A buyer of type $b \in \{1, \ldots, k\}$ can match with a seller of type $v \in \{1, \ldots, k\}$. Consider a partition of the set of buyers of each type b into k subsets $\{m(b, v)\}_{v=1}^k$, corresponding to each seller type that a buyer could be matched with. Each subset $m(b, v)$ can transact at a price $p(b, v)$. If a pair of types (b, v) do not transact in the search equilibrium then we can divide these into two cases: (1) Buyer of type b does not participate in the market because his valuation is too low (so $b < R_b^k$, where R_b^k is the marginal buyer type for the

k-types economy). (2) The seller of type s does not trade with the buyer because either s does not participate or the discount rate r is too high for all participating types to trade with each other. In case (2) the set $m(b,v)$ is empty. In case (1), the buyer of type b does not participate in the market and we can just assume that $m(b,v)$ has the same measure for each seller type v.

Let $\lambda(b,v)$ be the probability that a transaction is between a buyer of valuation $x(b)$ and a seller of valuation $y(v)$. Among buyer type b who does participate in the market $\lambda(b,v)$ is given by the measure of the set $m(b,v)$ divided by the sum of the measure of sets $\{m(b,v') : v' \in \{1,2,\ldots,k\}\}$. If buyer type b does not participate in the market $\lambda(b,v)$ is zero for all seller types.

The search message space in the k-environment specifies prices for each possible pairing of buyers-seller types, which means that there are k^2 prices for each pairing between the k-types of buyers and k-types of sellers. Let $y^B(b,v) = (y^B(b,v)_1, y^B(b,v)_2)$ be the net trades of the set of buyers of type b with sellers of type v and $y^S(b,v) = (y_1^S(b,v), y_2^S(b,v))$ be the net trades of the set of sellers of type v with buyers of type b. A profile of net-trades is $y = \left(y^B(b,v), y^S(v,b)\right)_{b,v\in\{1,\ldots,k\}}$.

Therefore, the set of net trades for the search mechanism is described by:

$$Y_k^s = \left\{ y \in ((-1,0,1) \times \mathcal{R}_+)^{2k^2} : \sum_b \sum_v \lambda(b,v)[y^B(b,v) + y^S(b,v)] = (0,0), y^B(b,v) = y^S(b,v) = 0 \text{ if } b < R_b^k \right\}. \tag{3.5}$$

Let $(\mu_c^k, M_c^k, g_c^k), (\mu_s^k, M_s^k, g_s^k)$ be the k-environment versions of the competitive and search allocation mechanisms, where

$$M_c^k = \{(p,y) \in \mathcal{R}_{++} \times Y_k^c : py_1^j(i) + y_2^j(i) = 0, \forall j \in \{B,S\}, \forall i \in \{1,\ldots,k\}\}$$
$$M_s^k = \{(p,y) \in \mathcal{R}_{++}^{k^2} \times Y_k^s : p(b,v)y_1^j(b,v) + y_2^j(b,v) = 0, \forall j \in \{B,S\}, \forall b,v \in \{1,\ldots,k\}\}.$$

and μ_c^k, μ_s^k are the finite analogues of μ_c, μ_s: correspondences that map E^k into M_c^k, M_s^k, respectively, and g_c^k, g_s^k are projections from M_c^k, M_s^k, respectively, to Y^k.

Informational Size of the Message Spaces

In the competitive mechanism of the k-economy, the message space includes only one price (as the price of the numeraire good is normalized to 1) and $2k$ types of buyers and sellers. However, market-clearing implies that if $2k-1$ types trade then the net-trades for the last type are implied, therefore the message space of the competitive mechanism M_c^k is a $2k$-dimensional manifold.

In the search mechanism the k-economy there are k^2 prices and effectively $2k^2$ different types of agents (each player's endowed type and who they are matched with) in the search equilibrium, that means that for each buyer (or seller), they form expectations regarding prices for transactions with each of the k-types of sellers (buyers) in the other side of the market, therefore each buyer (seller) individually has to form expectations regarding k distinct prices (each of which depends on the opportunity costs that depend on the distribution of buyers' types). This implies that the dimensional size of the message space is k^2 (prices) plus $\sum_{i=1}^{k}(k+k)$ different types of agents minus one dimension 1 due to the market-clearing condition. Therefore, M_s^k is a $3k^2-1$ dimensional manifold, which has approximately $1.5k$ times more dimensions than the message space of the competitive mechanism M_c^k. Clearly, this difference converges to infinity as $k \to \infty$, stated in the Proposition 4.

In other words, the search mechanism requires that each participant of the market be aware of all types of participants operating in the market to form expectations regarding payoffs from participating in the market and to bargain with the other participants. This is precisely the inverse of the intuition regarding the informational efficiency of the market as articulated by the literature on the informational efficiency of competitive markets: that each participant of the market can use prices as an efficient way to substitute for the information they would otherwise require to allocate resources without access to market prices.

Proposition 4. *M_c^k and M_s^k are $2k$ and $3k^2 - 1$ dimensional manifolds. Therefore, as $k \to \infty$ the ratio of the dimensional size of M_s^k to M_c^k converges to infinity.*

Proof. See Appendix Subsection A.1 ∎

Remark 3. On private information: If valuations are private information then upon matching buyers and sellers cannot play pricing strategies that are functions of the pair of matched types. Following Satterthwaite and Shneyerov (2007), consider the trade mechanism where buyers offer prices and sellers post reservation prices. If the offered price is higher than the reservation price the transaction occurs, otherwise there is no transaction.

Let the offered price by the buyer of type b be $\overline{p}(b)$ and the reservation price of a seller of type v be given by $\underline{p}(v)$. Then, there are $2k$ different "prices" in the search-mechanism under asymmetric information instead of k^2 prices. But there are still $2k^2$ types of consumers as each distinct pairing of buyers and sellers types can imply in a different net trade.

In this case, the dimensional size of the message space is smaller than in the search mechanism with perfect information but the result of Proposition 4 still holds as the ratio of dimensional sizes of the message spaces of the search mechanism to the competitive mechanism also diverge to infinity as the number of types of consumers increases.

3.3 The Informational Efficiency of Market-makers

3.3.1 Environment

Suppose that in addition to buyers and sellers (the consumers), there is a finite set J of market-makers in this economy (with abuse of notation, J has cardinality J). Market-makers are firms who act as profit-maximizing intermediaries who "make the market" by posting bid and ask prices for the indivisible good, intermediating trade between the suppliers (sellers) and the final consumers (buyers) in the economy.[7] Consumers valuations are private information so we assume market-makers are constrained to uniform pricing policies where there is no price discrimination. [8] Unlike the search mechanism, buyers and sellers do not directly

7. In here market-makers perform the same role as in Spulber (1996), but in this model the number of market-makers is finite and consumers are matched with different probability to each market-maker.

8. A market-maker could practice price discrimination but the trader's valuations and awareness are private information. There is two-sided asymmetric information in the sense that both buyers and sellers have private information regarding their valuations and costs as in Satterthwaite and Shneyerov (2007, 2008) and since the good is indivisible then the only direct revelation mechanism that is truthfully implementable consists of a pair of prices that the market-maker is

match with each other. Instead, both buyers and sellers trade through the market-makers. Buyers purchase from the lowest priced market-maker they have access too as long as it is lower than their valuation, sellers sell at the highest-priced market-maker as long as the posted price is higher than their cost.

Consider the case of the absence of any form of frictions of trade: All traders have costless access to all contracts posted by all market-makers. Consider a market-maker $j \in J$ who posts a pair of prices which the market-maker offers to buy and sell the good, respectively, (p_b, p_s) which are, respectively, higher and lower than prices posted by all other market-makers, then its profits are

$$\pi(p_s, p_b) = (p_s - p_b) \times b[1 - F(p_s)], \tag{3.6}$$

subject to the market clearing constraint that the quantity brought from sellers is equal to the quantity demanded by buyers:

$$b(1 - F(p_s)) = sG(p_b). \tag{3.7}$$

If the market-maker posts bid prices lower than some other market-maker no seller will sell to him or her and so its profits are zero. If the maker posts bid prices higher than all others but not the lowest ask prices the market-maker has monopolized the supply and profits also satisfy 3.6 subject to the resource constraint 3.7.

Proposition 5. *If at least two market-makers are operating then there is only one Nash equilibrium: for at least two market-makers to post a pair of bid-ask prices* $(p_b, p_s) =$

willing to buy or sell the good for.

(p^*, p^*), *that is market-makers post the competitive equilibrium price.*

Proof. See Appendix Subsection A.2 ■

Proposition 5 states that this environment of strategic price determination by market-makers implements the competitive equilibrium in a frictionless setting. The interesting application of such a model, which is developed over the next-subsections, is to consider the case of imperfectly functioning markets where there are frictions of trading. We describe frictions of trading in this setting by the hypothesis that traders might not have full access to all market-makers because they might be unaware of the full set of market-makers operating in the market.

Frictions of Trading: Consideration Sets Constrained by Unawareness

The market-maker model incorporates frictions of trading which allow it to yield results such as market-power, price dispersion, and other features of markets that do not exist in the competitive market mechanism. We model the frictions of trading by assuming that consumers (buyers and sellers) have constrained choice sets regarding the market-makers that they can trade with. We use the term (borrowing from Perla (2019)) imperfect awareness to describe these constrained choice sets regarding the market-makers operating, we assume that awareness is randomly and independently distributed. In addition, we introduce a diffusion process of awareness among the consumers over time (as in the model of sticky-information by Mankiw and Reis 2002) which can result in departures from the conditions of competitive equilibrium that the standard search model does not have but that has found support in empirical studies.

For $j \in J$ there is subset $A^j \subset [0,1]$ of buyers and sellers who have access to market-maker j and let $A^i = \{j \in J : i \in A^i\}$ be the set of market-makers that i has access to. The set of environments E includes the market-makers and the awareness among buyers and sellers regarding them (use the term "aware of j" to mean $j \in A^j$[9]), this information is given by $\{A^j\}_{j \in J}$. That is, $e^i = (x^i, A^i)$, where x^i is agent i's valuation or cost and A^i is the set of markets that i is aware of, an environment $e \in E$ is specified by $\boldsymbol{A} = \{A^j\}_{j \in J}$ and the distributions of buyers valuations F and sellers costs G.

Let the awareness parameter m^j given by $m^j = \lambda(A^j) \in (0,1]$, the (Lebesgue) measure of A^j which is the fraction of all traders aware of j. We assume that A^j is a simple random sample of the traders. This means that it satisfies the properties of independence of type and awareness which means that:

Assumption A1 The fraction of traders who are buyers in A^j is b, sellers is $s = 1 - b$.

Assumption A2 The distribution of types of buyers and sellers conditional to being in A^j are F and G, respectively.

Assumption A3 The probability of being aware of a competing market-maker is independent, that is

$$\lambda(A^j \cap A^h) = \lambda(A^j) \times \lambda(A^h), \tag{3.8}$$

for the awareness sets A^j, A^h of marker-makers j and h.

9. Other studies such as Armstrong and Vickers (2019b) and McAfee (1994) use the term "consideration set" or "availability rate", respectively, to indicate the subset of agents that buyers or sellers have access to and to indicate the degree of access among consumers in the market.

Note that 3.8 implies that the fraction of buyers and sellers who are aware of the seller j conditional on being aware of a competitor is m^j. As valuations, consumer's choice sets are private information so market-makers cannot price discriminate based on the consumer's choice sets.

Static Equilibrium

The solution concept used here is Nash equilibrium in mixed strategies: a mixed strategy profile for the market-makers is a profile of distributions over a subset $F \subset [\underline{c}, \overline{c}] \times [\underline{v}, \overline{v}]$ of pairs of prices that are consistent with market-clearing (that is, the quantity brought by the market-makers is equal to the quantity sold) such that any price pair on the support of the distributions is profit maximizing. As stated in Proposition 6 given a profile of awareness parameters $\boldsymbol{m}$ such that consumers are fully aware of at most one market-maker the unique equilibrium in the this environment is a profile of pricing strategies described by $(\{P^j\}_{j=1}^J, \boldsymbol{p})$, where P_j is a cumulative distribution function on $[0,1]$ and $\boldsymbol{p}$ is a function that maps $[0,1]$ into a pair of prices for buying and selling in $[\underline{c}, \overline{c}] \times [\underline{v}, \overline{v}]$. This function is such that posting any pair of prices $\boldsymbol{p}(\alpha) = (p^s(\alpha), p^b(\alpha))$ for α on the support of P^j is profit-maximizing and the resulting allocation is feasible.

Proposition 6. *If $\boldsymbol{m}$ is such that $m^j < 1$ for at least $J-1$ market-makers then there is a unique equilibrium that consists of a profile of mixed pricing strategies $\{P^j\}_{j \in J}$ and the sharing rule such that for a pair market-makers h and j if $m^h < m^j$ then traders aware of both will trade with h if posted prices are the same. The profile of equilibrium strategies features connected supports $[\underline{\alpha}^j, \overline{\alpha}^j]$ for each $j \in J$, they share a common lower bound of the support $\underline{\alpha}$, and the distributions are continuous on $[\underline{\alpha}, 1)$. For each $j \in J$, for $\alpha \in [\underline{\alpha}^j, \overline{\alpha}^j]$,*

$P^j(\alpha)$ satisfies

$$P^j(\alpha) = \frac{a^{\bar{j}}}{a^j} P^{\bar{j}}(\alpha),$$

where $\bar{j}$ is the market-maker with the largest availability a. The distribution $P^{\bar{j}}$ is given by

$$\prod_{j \neq \bar{j}} (1 - a^j P^{\bar{j}}(\alpha)) Q(\alpha) = \prod_{j \neq \bar{j}} (1 - m^j) Q(1).$$

where $Q(\alpha) = [p_s(\alpha) - p_b(\alpha)] s G(p_b(\alpha))$ is the quantity that is transacted per unit of awareness if there was no competition.

Consider a sequence of awareness profiles $\boldsymbol{m}_n$ If, for at least two market-makers h, g, m_n^h and m_n^g both converge to 1 then the equilibrium pricing strategies $\{P^j, \boldsymbol{p}\}_{j \in J}$ converge in probability to the competitive equilibrium price p^*, which is the unique equilibrium if $m^h = m^j = 1$ for at least two market-makers.

Proof. See Appendix Subsection A.3. ■

The equilibrium mixed strategy profile described in Proposition 6 is similar to the equilibrium described in McAfee (1994): the distributions of prices posted by the market-makers are non-degenerate and are continuous on the interior of the support, the larger market-makers (in terms of the m^j) transact at higher margins than smaller market-makers in the sense that the distribution of margins between ask and bid prices of the larger market-makers first-order stochastically dominate those of the smaller market-makers). The reason for this result is that in an economy with finitely many market-makers it is less likely that buyers and sellers are aware of a competitor to large market-maker than a competitor to a smaller market-maker, so the larger market-maker loses fewer customers if the spread between sell and bid prices is increased.

Dynamic Equilibrium: Competitive Price Formation, Awareness Diffusion

Suppose that time is discrete: $t = 0, 1, 2, \ldots$ and let $\beta = 1/(1+r)$ be the discount factor. The good is perishable so at each period sellers can produce one unit of the good and buyers have unit demand for one unit per period.

Awareness Diffusion: Given a set J of market-makers there is a awareness profile $\{m_t^j\}_{j=1}^J \in (0,1]^J$. Suppose awareness regarding a market-maker diffuses through the market according to

$$m_{t+1}^j = (1-\delta)m_t^j + M(m_t^j, 1-m_t^j), \tag{3.9}$$

where M is a matching function that represents the diffusion of awareness through traders who hitherto had access to the market-maker and $\delta \in [0,1)$ is the awareness depreciation parameter (that is, the rate in which traders "forget" about the market-maker).

Each market-makers chooses in period 0 to post prices according to a sequence of distributions for each period. Since the choice of the pricing strategies does not have any effect on the state of the market the optimal strategy for each market-maker is to choose the profit-maximizing pricing behavior at each period given the action profile of the other market-makers at that period. This means that at a given point in time t, prices practiced in the market are given by $\{P_t^j\}_j$ described in the proof of Proposition 6.

We are interested in the convergence of equilibrium prices and allocation to the competitive equilibrium. Since the outcome of the equilibrium is stochastic as the market-makers randomize their bid and ask prices we use the notion of convergence in probability: Convergence in probability of equilibrium prices and allocation to the competitive equilibrium means that the probability at a period t that prices and the allocation are at some positive distance from the prices and allocation of the competitive equilibrium converges to zero as $t \to \infty$. The distance between two allocations x and x' that assign a consumption bundle $x'(i)$, and $x(i)$, respectively, to consumer $i \in [0,1]$ is described by a function $D(x,x')$ that satisfies

$$D(x,x') = \int_0^1 \sum_l \left[|x_l(i) - x_l'(i)| \right] di.$$

Proposition 7 follows from Proposition 6 as the expected equilibrium margin between buy and ask prices posted by the market-makers converge to zero if $\lim m_t^j = 1$ for $m_t^j > 0$ and $J \geq 2$. Therefore, the awareness by the consumers regarding the market-makers operating converges to 1 as $t \to \infty$ implies that a measure converging to 1 of consumers has access to buy and ask prices that are converging in probability to the competitive price. Therefore, the equilibrium allocation converges in probability to the competitive allocation.

Proposition 7. *If the law of motion for awareness diffusion 3.9 implies that* $\lim m_t^j = 1$ *for* $m_t^j > 0$ *then if* $J \geq 2$ *as* $t \to \infty$ *the equilibrium prices and the equilibrium allocation converge in probability to the competitive equilibrium.*

Steady-State Equilibrium

The analysis so far implied that a market where trade is mediated by market-makers approximates the outcome of the competitive equilibrium as awareness regarding the market-makers diffuses. However, by allowing awareness to depreciate with $\delta > 0$ (as consumers forget about a market-maker) then given some assumptions on the matching function M[10] then there is a unique steady-state awareness level $\hat{m}$ such that

$$\hat{m} = M(\hat{m}, 1 - \hat{m})/\delta.$$

There is a corresponding steady-state equilibrium if and only if all market-makers have the same awareness parameter $m^j = \hat{m}$ which is a symmetric mixed strategy described by a pair $\{P, \boldsymbol{p}\}$ that all market-makers follow. If we let $0\delta \to 0$ then awareness diffusion function M implies that $\hat{m} \to 1$, that is if awareness depreciates very slowly the steady-state level of awareness approximates full awareness, which implies that the steady-state equilibrium approximates the competitive equilibrium.

Entry and Exit

Suppose that at some period t there are some market-makers that are operating in the market and some that are not operating in the market (which means that their awareness parameter m^j is zero). There is an entry cost $E > 0$ which is the cost

10. More precisely, that M is continuously differentiable, concave in both arguments, satisfies the condition that $\lim_{a \to 0} \partial M(a, b)/\partial a = \infty, \forall b > 0$.

of setting up a starting consumer base represented by an entry-level awareness parameter $m_e \in (0,1)$.

The law of motion for the diffusion of awareness 3.9 implies that for M increasing and concave in both arguments, m_e that is not very large, and if the depreciation parameter is not very large then market-makers grow after entry (in the sense that m_t^j is increasing over time). That implies that incumbent market-makers are larger than entrants and therefore transact at higher expected margins. If we follow Spulber (1996) interpretation that market-makers are firms who intermediate between suppliers and consumers this equilibrium property replicates the findings of Foster, Haltiwanger, and Syverson (2008, 2016) that incumbents charge higher prices than entrants.

Contestable Market Equilibrium

Suppose that there are only two market-makers. The set of market-makers is $J = \{1,2\}$, suppose that $m_0^1 = 1, m_0^2 = 0$ at a date set to 0. That is, 1 is a monopolist market-maker and all traders have access to his posted contracts and 2 is out of the market, but 2 can decide to enter in the current period. That is, 2's action set is that she can choose $m_t^2 \in \{0, m_e\}$ besides the price if she has not entered the market in the period before. A monopoly deterrence equilibrium is a situation where the incumbent 1 chooses to post buying and selling prices for each period such that the profits from offering better prices to the consumers are too low to compensate for the cost of entering the market.

Definition 3. A *monopoly deterrence equilibrium* is an equilibrium where 1 chooses a pricing schedule and given this pricing schedule, 2 finds it optimal to not enter.

The pricing schedule is the profit-maximizing in the sense that a higher selling-buying margin that yields higher profits for 1 implies that 2 finds it optimal to enter and undercut 1's posted offers in every period and is profit-maximizing in the sense it yields a higher discounted expected value of the profit stream for 1 than the expected value of the profits in the equilibrium under a duopoly if 2 also enters the market.

The proposition below states that if entry costs are high enough and awareness diffusion is fast enough then the unique equilibrium is for the monopolist to deter entry. That is because entry cost is higher than the expected profits that can be obtained in the duopoly competition process where market-maker 2 competes with the former monopolist. However, the monopolist 1 must commit to a sequence of prices that still yield a low enough profit to deter the entrant. The unique equilibrium is the sequence of prices that makes 2 indifferent between entering and not but that maximizes the present value of 1's profit stream. As discount rates decrease the present value of the gains from entering the market increase. This implies that the buy and ask prices posted by the monopolist become closer to the price of competitive equilibrium. As the discount rate r converges to zero, the present value of any positive profit stream converges to infinity, which implies that the monopoly deterrence equilibrium converges to the competitive equilibrium as the discount rate converges to zero.

Proposition 8. *If awareness diffusion is fast enough so $\sum_{t=0}^{\infty}(1 - m_t^2) \leq C$ for some constant C conditional on seller 2 entry, and the discount rate r is low enough, then, for an entry cost E equal or higher than $C \times \pi^M$, the unique equilibrium is the monopoly* *deterrence: The monopolist commits to posts prices* $\mathbf{p}(\pi)$ *that yield a per-period profit of*

$$\pi = E / \left(\sum_{t=0}^{\infty} \beta^t m_t^2 \right)$$

to deter entry. As r *converges to zero the deterrence monopoly equilibrium profit flow* π *converges to zero, which means the posted buying and selling prices converge to the competitive equilibrium prices* p^*.

Proof. See Appendix Subsection A.4. ■

The existence of entry costs for 2 can be interpreted to represent the costs of communicating additional information to the market-participants so that if the costs of communicating additional information are higher than the (private) benefits which are the profits 2 obtains from entering the market, then it does not occur in equilibrium. [11]

3.3.2 Allocation Mechanism

The allocation mechanism in this case is the allocation mechanism that implements the allocation corresponding to a realization of the Markov perfect equilibrium of the market-maker economy at some period (which is a profile of realized pairs of prices for each market-maker). Note that if there is imperfect awareness regarding almost all the market-makers (i.e. $\boldsymbol{m} = (m^j)_{j=1}^J$ such that $m^j = 1$ for

11. We have not performed a welfare analysis to check if the monopoly deterrence equilibrium is more efficient than the duopoly after 1's entry. As the social benefits of 2's entry would be the reduction of the deadweight loss thanks to prices closer to perfect competition. The social benefits are different from the private benefits of entry.

at most one j) then for any market-maker the posted price for buying is strictly smaller than for selling and therefore profits are strictly positive.

Following Hurwicz (1977a), we interpret that the profits of the market-makers and the resulting deadweight losses are both components of the "cost" of operating the allocation mechanism: In that case, the allocation implemented by the mechanism features strictly negative net-trades for the numeraire good among the agents in the economy and as bid and ask prices diverge it means that not every buyer or seller who has access to a market-maker and would trade under competitive prices does so.

The set of net-trades incorporates the possibility of market-makers making profits by buying at lower prices than they sell:

$$Y_m = \left\{ y : [0,1] \to \{-1,0,1\} \times \mathcal{R} : \int_0^1 y_1(i)di = 0, \int_0^1 y_2(i))di \leq 0 \right\}.$$

Given a realized profile of prices $p_m = (p_1, \ldots, p_J)$ (which describes the equilibrium at some date in the dynamic version of the model) the message space is given by

$$M_m = \{(p_m, y) \in \mathcal{R}^J_{++} \times Y_m : \text{for each } i, \exists j \in A^i \text{ s.t. } i \in A^j \text{ and } p^j y_1(i) + y_2(i) = 0\},$$

and μ_m is a correspondence on E to M_m that satisfies

$$\mu_m = \cap_i \mu^i_m(e^i),$$

where $\mu^i_m : E^i \rightrightarrows M_m$ satisfies

$$\mu_m^i(e^i) = \left\{ (p_m, y) \in M_m : y_1(i) = \begin{cases} (0,0) & \text{if } i \in B \text{ and } v^i < \min\{p_s^j : j \in A^i\} \text{ or } A^i = \emptyset \\ (1, -\min\{p^j : j \in A^i\}) & \text{if } i \in B \text{ and } v^i \geq \min\{p_s^j : j \in A^i\} \\ (0,0) & \text{if } i \in S \text{ and } c^i > \min\{p_b^j : j \in A^i\} \text{ or } A^i = \emptyset \\ (-1, \min\{p^j : j \in A^i\}) & \text{if } i \in S \text{ and } c^i \leq \min\{p_b^j : j \in A^i\} \end{cases} \right\}.$$

Informational Efficiency

As in Subsection 3.2.3 consider a sequence of finite types economies $\{e_k\}$ that approximates the environment with the cumulative distributions of buyers and sellers valuations F and G. In this case, the dimensional size of the message space incorporates the different market-makers that make the market: If there are $N \leq J$ market-makers with non-zero awareness parameters then there are $2N$ different prices posted to the consumers, plus the subset of consumers who are not aware of any market-makers. As in the case of the search allocation mechanism, the set of consumer "types" increases to include differentiate consumers by their access to different prices (as awareness is heterogeneous).

The number of consumer types is determined by the discrete distributions of valuations (G_k, F_k) and N. The type of a consumer can be specified by a triple (r, κ, h) where $r \in \{b, s\}$ denotes if the consumer is a buyer or seller, $\kappa \in \{1, 2, \ldots, k\}$ denotes the valuation type of buyer or seller, $h \in J$ denotes the market-maker that the consumer transacted with (including $h = \emptyset$ if the consumer does not transact with any market-maker).[12] Therefore, there are $2kN$ or $2k(N+1)$ types of

12. Note that it is not necessary for the computation of the dimensional size of the message-space to include the information of which other market-makers the consumer was aware of besides the one he or she transacted with.

consumers if the subset of consumers who are not aware of any market-makers is respectively empty or non-empty. Therefore, market-clearing of the indivisible good among consumers who interact with each market-maker implies that the message space corresponding to environments with k different valuations for buyers and sellers is a Z-dimensional manifold where Z is equal to $2N + (2k - 1)N$ or $2N + (2k - 1)(N + 1)$, respectively if the subset of consumers who are not aware of any market-makers is respectively empty or non-empty. This implies in the following corollary:

Corollary 8.1. *As k increases to infinity the ratio of the dimensional size of the message spaces of the market-maker mechanism to the competitive mechanism $Z/2k$ converges to N or $N + 1$.*

That is, the ratio of the size of the message spaces between the competitive mechanism and the market-maker mechanism when the number of types of consumers is large is approximately the number of market-makers operating in the market. This result is intuitive since the competitive mechanism implicitly assumes a single monopolist market-maker called the Walrasian auctioneer whose bid and ask prices have zero spread.

Consider the case of the monopolist market-maker who deters entry, note that the subset of consumers who are not aware of any market-makers is empty. So the number of consumer types is $2k$ but a pair of prices is realized instead of one price in the case of the Walrasian auctioneer. It represents the most informationally efficient mechanism in this class of market-maker environments with informational size $2k + 1$, or only one dimension more than the competitive mechanism (outside of the limit case of $\delta = 0, \beta \to 1$ when it converges to perfect competition and there is only one price posted to all consumers). This additional dimension reflects the profit margin between purchase and sale to provide incentives for the market-makers to "produce" the price mechanism.

However, for the model to generate more sophisticated market-behavior in equilibrium, such as price dispersion with average spreads between the ask and bid prices that depends on the degree of tenure of the market-maker has been operating, then it requires the assumption that there are at least two market-makers operating. In that case, the minimum dimensional size of the message space of the market-maker equilibrium in a given period is $Z = 4 + 2(2k - 1)$, as there is a pair of bid and ask prices and there are two profiles of net trades for the $2k$ types of consumers.

3.4 Extending the Model to an Environment with L Goods

It is a simple exercise to extend the market-maker model of Section 3.3 to a more general environment with multiple goods, which is more comparable to the original informational efficiency papers of Mount and Reiter (1974), Hurwicz (1977b), and Jordan (1982). Consider an environment with $L > 1$ goods where good L is the numeraire. There is a continuum of consumers $i \in [0, 1]$, with consumption sets $X = \mathcal{R}_+^{L-1} \times \mathcal{R}$ and CRRA preferences, where

$$u_i(x) = \sum_{l=1}^{L-1} \theta_{il} \left(\frac{x_l^{1-\sigma_l}}{1 - \sigma_l} \right) + x_L, \theta_{il} > 0, \sigma_l \geq 0, \forall l.$$

Suppose θ_{il} is private information and distributed on interval $[\underline{\theta_l}, \overline{\theta_l}]$ according to some cumulative distribution F.

There is a set $J_l, |J_l| \geq 1$ of market-makers operating in the markets for each good $l \in 1, \ldots, L-1$ and they compete by posting bid and ask prices for the good in exchange for the numeraire.[13]

Quasi-linear preferences imply that the equilibrium price of the market for each good is independent of the other markets, and profits are continuous functions of the posted prices as long as they are more attractive to sellers/buyers than their competitor's prices. Therefore, the results of the analysis of Section 3.3 applies here. In particular, consider the ratio of the dimensional sizes of the message space of the market-maker allocation mechanism to the competitive mechanism. By a similar argument to the two good economies of subsection 3.3.2 this ratio is well defined in a finite types economy with k types of buyers and k types of sellers. As the number of consumer types converges to infinity and the distribution of types in a finite types economy approximates the distribution of types of the environment E the ratio of dimensional size converges to a number equal to $\sum_{l=1}^{L-1} J_l/(L-1)$. That is, the average number of market-makers operating in the market for each good $l \in \{1, \ldots, L-1\}$ yields the ratio of the size of the message space to the competitive economy (where there is one market-maker for each good: the Walrasian auctioneer).

13. Note that if we allow for non-linear pricing (as described, for instance, Bolton and Dewatripont (2005)) in this case the dimensional size of the message space is higher, with continuously varying prices it is also infinity. That provides a justification why we do not see many empirical examples of sophisticated non-linear pricing if allocation mechanisms are informationally constrained.

The dynamic deterrence equilibrium described in Subsection 3.3.1 can be implemented in this L-goods economy with one monopolist in the market for each good $l \in \{1, \ldots, L-1\}$. In equilibrium, there will be $2(L-1)$ prices: a buying and selling price for each good l in terms of the numeraire good L. This implies that the dimensional size of the message space of the monopolist market-maker economy has $L-1$ more dimensions than the competitive mechanism, one for each good besides the numeraire good. This means that the difference in informational requirements between this mechanism and the competitive market mechanism converges to zero as the finite types economy approximates an economy with a continuum of consumer types.

3.5 Concluding Remarks

In this chapter, we consider the informational efficiency of decentralized price formation. In particular, we are interested in economies with strategic agents where the allocation mechanism converges to the competitive mechanism. We study two such mechanisms: search and market-makers.

While the random search model has been extensively studied in the literature, we show that it is unattractive from an informational perspective. In particular, we showed that for large economies the search allocation mechanism requires infinitely more information than the competitive mechanism. A true search mechanism, where everyone must be able to search across all the people in the economy to find trading partners, is extremely inefficient in terms of information as it requires that each agent must have a complete model of the economy. That is one possible reason we do not often observe single buyers trading with single sellers in real-world economies.

In contrast, we propose a different decentralized mechanism with market-makers. Such a mechanism has a few attractive features. First, the market-maker mechanism better matches certain features of the data, such as exhibiting price dispersion and prices that depend on the tenure of firms in the market. The other attractive feature, which is the focus of this chapter, is that the market-maker mechanism almost requires as little information as the competitive allocation. Moreover, the mechanism requires relatively little information, even when it is used to explain deviations from the competitive allocation. This informational efficiency is one possible reason we observe intermediaries that facilitate trade between individual original sellers and individual final buyers and that their operation is restricted to individual markets.[14] It is a rather puzzling result that an economy where trading for each good is intermediated by few traders can be though as more informationally efficient than markets where trading is highly decentralized but it is an intuitive result: if consumers only need to be aware of a few intermediators for each good they purchase the informational requirements are much smaller than if consumers need to form a model of the whole market before engaging in random matching and bargaining for their consumption bundle.

But, perhaps the main contribution of this study is to point out that even if a theoretical model is able to replicate certain features of real economies the degree of informational efficiency of the allocation mechanism implicit in a model should

14. For example, real estate agencies specialize in the intermediation of trade in real estate.

also be an element to take into consideration when judging the degree of plausibility of that model. In the specific case of models that explain price formation should also explain how the message space of the allocation mechanism that is implicit in the model approximates the main feature of competitive mechanism: that agents can take terms of trade as a given without the need to "think" about how they are determined.

Information Design with Multiple Competitors

Chapter 4

Price Competition and the Use of Consumer Data

> "The complete theory of competition cannot be known because it is an open-ended theory; it is always possible that a new range of problems will be posed in this framework, and then, no matter how well-developed the theory was with respect to the earlier range of problems, it may require extensive elaboration in respects which previously it had glossed over or ignored. "
>
> George Stigler (1957, p. 14)

Technological change has given both consumers and firms access to information that was previously unimaginable. Consumers can access more data about firms and their products through online reviews. firms can access more data about consumers through deeper market research and tracking data on previous purchases, web searches, or social media posts. That consumer data is vital to a firm's ability to price discriminate.

At the same time, technological change has allowed firms to differentiate products. Identical goods from the supply side (same marginal cost) are not identical from the demand side (different marginal valuation). Gone are the days of one-size-fits-all goods (Neiman and Vavra 2019) To me, most beer tastes the exact same and is substitutable. To others, that claim is preposterous. But all beer companies need to pick prices with both types of consumers in mind. Crucially though, the degree of substitutability, and thus competition, depends on the market make-up and what firms know about consumers. This chapter asks, first, how do firms compete in a market where people differ in the degree of substitutability? Second, how does firm access to data on consumers affect how they compete? Finally, when does access to consumer data come back to help or hurt consumers?

To be more specific about the model, I study a stylized model where a continuum of consumers can buy a single good from one of two firms. The two firms compete *à la* Bertrand. However, contrary to a textbook Bertrand model with identical goods—as in Mas-Colell, Winston, and Green (1995, 388) or Tirole (1988, 310)—the consumer may have a different valuation for the goods from the two firms. A firm faces a trade-off between raising the price when the consumer prefers its good and lowering price to compete whenever the consumer prefers the other firm's good. Therefore, I extend a standard model of price discrimination to allow for general, varied access to consumer data by firms.

The paper first considers a simple model to capture the trade-off facing firms where the consumer may be one of three types: one type loyal to each firm for whom the firms' goods are not substitutes and each firm would be a monopolist if they knew the loyal consumers and a third type for whom the firms' goods are perfect substitutes. That setup captures the fundamental trade-off. Each firm wants to raise prices on the consumers loyal to them since they are monopolists. However, for the consumers that are not loyal, the force of Bertrand competition pushes the price down.

In this environment, learning serves two purposes. First, as in the standard monopoly case, learning about customers and being able to use price discrimination allows firms to charge a high price to their loyal consumers and extract more surplus. This hurts consumers. At the same time, a firm can learn which customers they can lure away from the other firm. They learn when it is worth competition on price. In that case, consumer data encourages more intense competition and drives down the price. The model sets up a horse race between the two forces.

My first contribution is methodological. Constructing equilibria is not obvious in this environment. In general, to balance this trade-off, the unique equilibrium involves each firm randomizing over prices. Due to the discontinuity inherent in Bertrand competition, the randomization allows each firm to avoid losing the consumer by only a small amount on price. This leads to price dispersion in equilibrium.

To construct the unique equilibrium, I use simple observations about the residual demand curve facing each firm and what each firm can guarantee itself by only pricing for loyal consumers.[1] The unique equilibrium on the game involves both

1. The construction of the demand curve is similar to Albrecht (2018)'s analysis of Roesler and

firms choosing their price to generate a unit-elastic demand curve for the other firm, making the firms indifferent over a range.

I then ask in a stylized model, given the equilibrium behavior under competition, when is access to consumer data helpful or harmful to consumers and producers? Without competition, the welfare implications of consumer data are clear; information that allows price discrimination weakly raises total profit. This is simply a manifestation of Blackwell (1951, 1953) that information is always valuable for single receiver.[2] Unless output increases, price discrimination lowers consumer surplus.

To study the role of information, I fix underlying valuations, i.e. demand curve in the aggregate market. I then vary the consumer data available to each firm, allowing price discrimination. However, I allow for the case where firms do not have the same information. Therefore, higher-order beliefs matter, e.g. Target doesn't know my exact address, *but* Target knows that Amazon knows... These beliefs affect equilibrium prices.

Finally, given the competitive pricing of the two firms, what is the consumer-optimal form of consumer data? Even in the simple example, the solution to this trade-off is not obvious. However, I show that, contrary to the monopoly case, giving firms *complete information* is consumer optimal, even in an efficient environment where trade always occurs. The increase in consumer surplus does not come from any risk-aversion or from increased consumption, as in models where price

Szentes (2017).

2. See Bergemann, Brooks, and Morris (2015) for the general information case of a monopolist.

discrimination increases consumer surplus by increasing trade. Even environments where production does not increase, consumer surplus is maximized under complete information, because the firms to compete more fiercely and charge a lower price when the consumer learns the goods are substitutes for him. The gain from fierce Bertrand competition when the firm's goods are perfect substitutes outweighs the loss when the consumer learns she only values one good and thus that firm is able to raise its price.

4.1 Price Discrimination Literature

This paper contributes to the literature on the relationship between price discrimination and competition is a standard issue in industrial organization–in classic form in Pigou (1920) and Robinson (1933), and in modern form in Borenstein (1985) and Holmes (1989) (See Stole (2007) for a summary).[3] One particular branch of the price discrimination literature shows the optimality of constant-elasticity demand curves, which also show up in my model. This shows up in more standard monopoly models, such as Aguirre (2008), Aguirre and Cowan (2015), Roesler and Szentes (2017), and Condorelli and Szentes (2019). Unit-elastic demand also shows up in non-Bayesian monopoly models, such as Neeman 2003; Bergemann and Schlag 2008; Renou and Schlag 2010. These models involve some

3. This literature has always included a large empirical component. For a recent example, Chandra and Lederman (2018) show that in the Canadian airline market price discrimination raises prices for some passengers and lowers prices for others. This is consistent with my main result, although in my model the average price is lower.

form of a maximin utility function or minimax regret, which relates to the outcome of my model where one firm can only achieve his maximin revenue and the rest is competed away by the other firm.

The constant-elasticity implies that there can be multiple prices where marginal revenue equals marginal cost. In my model, this leads to price dispersion. Contrary to the classic price discrimination papers of Stigler (1961) and Burdett and Judd (1983), in my model, price dispersion comes from *each firm* setting a random price, such as in Varian (1980). However, the market power that generates the distribution is due to the goods being imperfect substitutes, and thus having some brand loyalty as in Rosenthal (1980), compared to search frictions.

However, since each firm is randomizing over the set of prices, what Menzio and Trachter (2015) aptly call a "game of cat-and-mouse", the equilibrium prices are not only the lowest, consumer-optimal price found in Roesler and Szentes (2017), but a range of prices that the firm is indifferent between.

In addition, this chapter contributes to the literature on information design. In particular, I consider the range of possible outcomes that could arise for some information structure, as in Bergemann, Brooks, and Morris (2015). In particular, I am interested in the information structures that are consumer-optimal, as in Roesler and Szentes (2017). However, contrary to both papers, I introduce competition explicitly into the model. Competition is between the receivers of information, making it more like the literature on robust mechanism design (Bergemann and Morris 2013), and not between the senders, which is the topic of much of the literature on information design and competition (Boleslavsky and Cotton 2015, 2016; Gentzkow and Kamenica 2016b; Li and Norman 2018; Au and Kawai 2020).

As in Bergemann, Brooks, and Morris (2015), the consumer-optimal outcome involves extremal markets, where the firms are indifferent among all prices in the support of valuations.

4.2 Three Type Model

Throughout there are two firms with differentiated goods that cost 0 to produce. For now, we consider a model with three types of consumers with a total measure of one. First, there are consumers who are loyal to firm 1 and only buy from firm 1. Second, there are consumers loyal to firm 2. Third, there are indifferent consumers who buy from whoever has the cheapest price.[4] Each consumer has unit demand at $p = 1$ from the firm(s) that she is willing to buy from. A consumer's type is a pair of valuations (willingness to pay) $v = (v_1, v_2) \in \{(1,0),(0,1),(1,1)\} = V$. The firms start with a common prior on the market: (m_{10}, m_{01}, m_{11}). In general, we will be interested "interior" markets when no element is zero and all types exist in the economy with positive measure.

A market for firm i is the collection of consumers that a firm cannot differentiate and must set the same price for. Mathematically, a market is defined as a distribution of the types of consumers. Graphically, with three-types, a market is a point on the unit simplex. For all illustrations, I will consider a starting, aggregate market: $\left(\frac{1}{4}, \frac{1}{6}, \frac{7}{12}\right)$. The aggregate market with no price discrimination is

4. While I model consumers as having different preferences over sellers' goods, the baseline model could be interpreted as one where consumers differ in whether or not they are "aware" of 1 or 2 firms, as in Guthmann (2019a, 2019b) and Guthmann and Albrecht (2020) and Chapter 3, or "consider", as in Armstrong and Vickers (2020), or are "captive" as in Armstrong and Vickers (2019a) and Elliott and Galeotti (2019).

represented by the orange circle in Figure 4.1.

If both firms have complete data on consumers and know each consumer's type, the firms can segment consumers into three different markets and set a different price for the three consumers. A different price for each type is perfect price discrimination. For the example, complete information means that for each firm there is one market $(1, 0, 0)$ with measure $\frac{1}{4}$, another market $(0, 1, 0)$ with measure $\frac{1}{6}$, and a final market $(0, 0, 1)$ with measure $\frac{7}{12}$. These three markets are represented by the blue diamonds in Figure 4.1.

No price discrimination and perfect price discrimination are extreme cases. There are other types of markets. Consider a situation where each firm receives a signal that with probability q correctly reveals a consumer's type and with probability $(1 - q)$ the signal comes from a random consumer. This nests the two previous cases with $q = 1$ being a perfect signal and $q = 0$ being an uninformative signal. For any intermediate q, upon seeing a signal that says "this consumer is an indifferent type" the firm properly updates and puts more weight that the consumer being indifferent. With this set of three possible signals, there are three markets: one corresponding to each signal. On the simplex, each feasible market corresponds to a convex combination of the aggregate market and the perfectly discriminated markets. The three markets are plotted by the green triangles in Figure 4.1.

Formally, the firms' access to consumer data, also called their information structure, is a set of signals for each firm S_i, and a probability distribution which maps the profile of the consumer's values to the profile of signals: $\pi : V \to \Delta(S)$. The utility functions and the information structure (S, π) are the parameters for a

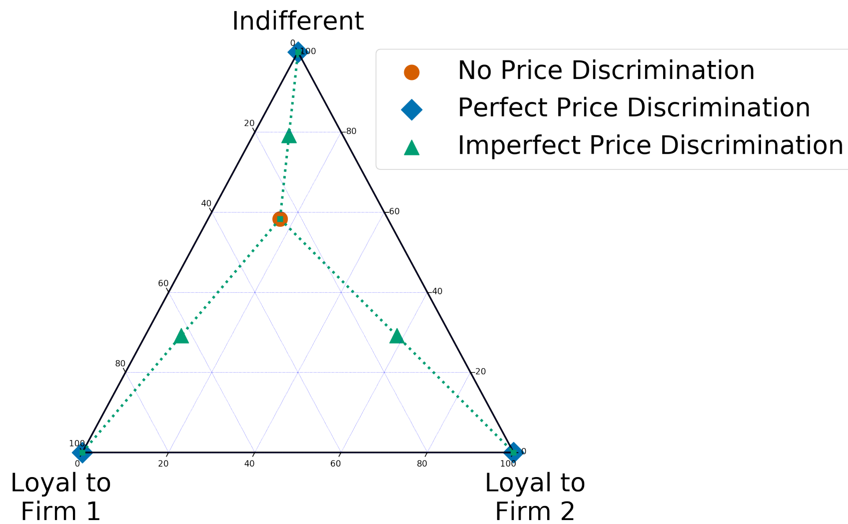

Figure 4.1: Feasible Markets

game of incomplete information. We will define the rest of the game after fixing (S, π).

Each firm i observes a signal $s_i \in S_i$. A pure strategy for firm i is a price for each signal $\{p_i\}_{s_i} \in \mathbb{R}_+^{|S_i|}$. The discontinuity of payoffs requires that we work with mixed prices for each firm, sometimes called price dispersion. A mixed strategy, $F_i(p|s_i)$, is the probability that $p_i \leq p$ given receiving a signal s_i. Let $f_i(p|s_i)$ be the density associated with F_i, when it is defined.

For a given (S, π), a strategy profile is a Bayes Nash equilibrium (BNE) if $f_i(p_i|s_i)$ is not defined (*i.e.* p_i is played with positive probability) or $f_i(p_i|s_i) > 0$ implies

$$p_i \in \arg\max_{p_i'} \quad \underbrace{p_i' \mathbb{E}[v_i = 1, v_j = 0|s_i]}_{\text{Loyal Customers}} + \underbrace{p_i' \mathbb{E}\left[\left(1 - F_j(p_i)\right), v_i = 1, v_j = 1|s_i\right]}_{\text{Indifferent Customers}},$$

given $F_j(p)$, for all s_i, s_j, i, j. Because we are interested in the effect of information on equilibria, we will be doing comparative statics with respect to the information structure. In addition, we will want to ask if there exists any information structure that generates an equilibrium outcome. For that, the following definition is helpful: a strategy profile is a Bayes correlated equilibrium (BCE) if it is a BNE for some information structure (Bergemann and Morris 2016). When looking for the highest possible consumer surplus, the problem requires searching over the set of Bayes correlated equilibria.

4.2.1 Solving for Equilibrium, Public Consumer Data

Before looking at the set of possible information structures, it is helpful to characterize the equilibrium for particular information structures. First, we will consider public information where both firms have the same data on consumers. The easiest case is complete information, where firms receive a signal that is perfectly reveals the consumers type: $s_1 = s_2 = v$. If the consumer will not buy from the competitor, $v_i = 0$, firm j sets the monopoly price of 1. If the consumer values both goods, $v_1 = v_2 = 1$, because of Bertrand competitive, the price of driven to 0. Combined, the expected price is $m_{10} + m_{01}$. If we plot the distribution of prices the expected price is the area above price distribution.

The equilibrium for the aggregate market, with no additional consumer data and therefore no price discrimination, is more complicated to solve for but has a known solution (Narasimhan 1988).[5] Lemma 2 characterizes the equilibrium.

5. Thanks to Mark Armstrong for this reference and making me aware of the related work of Armstrong and Vickers (2019a).

Lemma 2. *Let $m_{10} \geq m_{01}$. The unique BNE profit is m_{10} for firm 1 and $\frac{m_{10}}{m_{10}+m_{11}}(1 - m_{10})$ for firm 2. The unique strategies are given by*

$$F_1^*(p) = \begin{cases} 0 & p < \underline{p} \\ 1 - \dfrac{\underline{p}(m_{11} + m_{01}) - pm_{01}}{pm_{11}} & p \in \left[\underline{p}, 1\right) \\ 1 & p \geq 1 \end{cases}$$

$$F_2^*(p) = \begin{cases} 0 & p < \underline{p} \\ 1 - \dfrac{m_{10}(1-p)}{pm_{11}} & p \in \left[\underline{p}, 1\right], \end{cases}$$

Proof. Despite the result appearing before in Narasimhan (1988), I present a new proof here for completeness and because the proof technique is used in the extensions. First, let us plot the proposed equilibrium distribution of prices in Figure 4.2. Conditional on $v_1 = 1$, firm 1 assigns probability $\frac{m_{10}}{m_{10}+m_{11}}$ to being the monopolist. Regardless of what firm 2 does, firm 1 will never set a price below $\underline{p} = \frac{m_{10}}{m_{10}+m_{11}}$. Therefore, neither will firm 2.

The rest of the construction relies on simple observations of each firm's best-response when facing a residual demand curve. Translating demand curves into game-theoretic terms, best-responding by each seller means choosing a price and quantity subject to the constraint that the pair is on her demand curve.[6] To find the best-response, plot firm 1's indifference curves, which are just the iso-revenue curves, since the cost is zero.[7] This is also the worst-case (maximin) profit for firm

6. The best-response could also be found by constructing the marginal revenue (MR) curve. However, because MR is discontinuous, the optimal expected quantity is the maximum quantity such that MR > MC = 0, instead of the standard MR=MC condition used for continuous functions.

7. I thank Kevin M. Murphy for pointing out the benefits of plotting indifference curves in

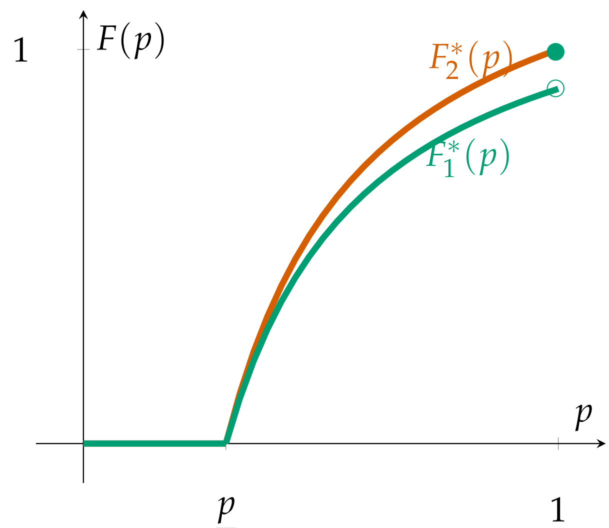

Figure 4.2: Equilibrium Price Distribution

1 since he could never do worse than if $p_2 = 0$.

To show this strategy is an equilibrium, in orange in Figure 4.3b, I plot firm 1's residual demand, given $F_2^*(p)$, $D(Q_1)$. Notice that to the right of m_{10}, where the residual demand comes from the indifferent buyer, the curve is simply one minus the $F_2^*(p)$ of firm 2's randomized price. The only complication is when a specific price is set with positive probability, which will be addressed shortly. Firm 1's problem is to maximize profits, subject to its residual demand curve. Because the marginal cost is zero, the isoprofit curves are simply given by $\pi = p_1 Q_1$.

Because of the shape of the residual demand curve, firm 1 is indifferent between any price between $\frac{m_{10}}{m_{10}+m_{11}}$ and 1. Therefore, firm 1 would be optimally mixing with any price in that range. For any mixing, the highest profit is simply given by m_{10}.

Marshallian, quantity-price coordinates.

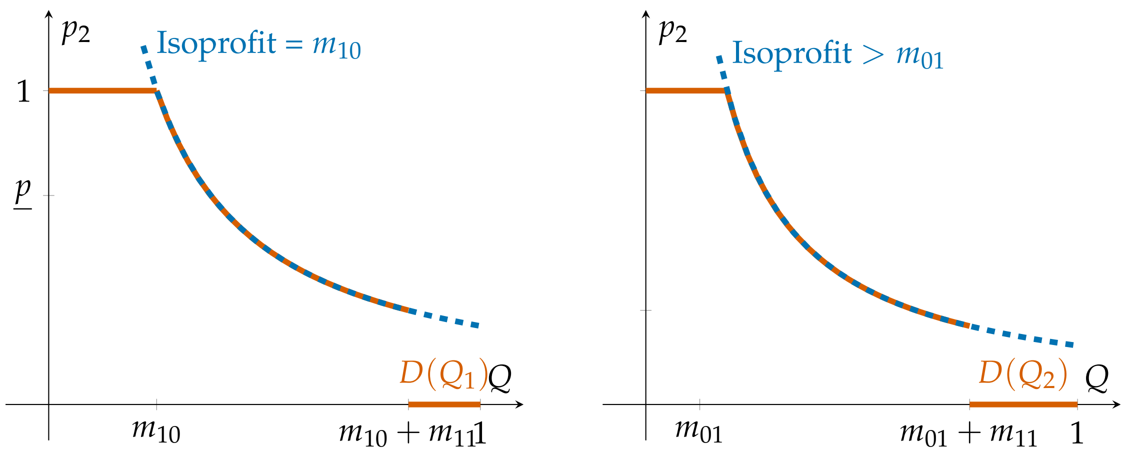

(a) Firm 1

(b) Firm 2

Figure 4.3: Equilibrium Residual Demand

We do the same thing for firm 2 in Figure 4.3b. Again, firm 1 has picked a strategy. However, this time, since firm 1 set $p = 1$, the top of firm 2's residual demand extends out beyond just it's loyal consumers with measure m_{01}, and firm 2 can still receive some consumers when $p_1 = 1$.

Therefore, we have shown that each strategy is the best-response to the other strategy and constructed an equilibrium. Uniqueness is immediate as well.

■

Due to the discontinuity of payoffs in Bertrand competition, if $p \neq 1$, every equilibrium involves a distribution of prices. For either firm to be indifferent, $p \times Q = \text{constant}$. The distribution of prices is, therefore, proportional to $-\frac{1}{p}$. The only exception is that there could be a possible mass point at $p = 1$. This distribution will show up in any equilibrium.

4.2.2 Construction of Mixed Strategy Equilibrium

This subsection constructs the equilibrium price distribution of Lemma 2, as compared to just verifying that the given strategy is an equilibrium. This is worthwhile because the construction can be applied to a wide range of markets. Translating demand curves into game-theoretic terms, best-responding by each seller means choosing a price and quantity subject to the constraint that the pair is on her demand curve. As a first guess of the equilibrium strategies, let us first find firm 1's maximin strategy; suppose that $p_2 = 0$. The best-response price is one and the expected quantity that firm 1 sells is m_{10}.[8]

However, setting a price of zero is not optimal for firm 2. Now I ask, what is the highest (pure strategy) price that firm 2 can charge without changing firm 1's behavior? That case is clearly a better-reply, even if it is not a best-reply. Increasing the price of firm 2 increases firm 1's demand curve over the range $(m_{10}, m_{10} + m_{11})$ (selling to the indifferent types), by p_2. By plotting firm 1's iso-profit curve, we can see that firm 2's price can be increased to $\frac{m_{10}}{m_{10}+m_{11}}$ without inducing firm 1 to change. While this is not going to be an equilibrium, it will allow us to bound prices. First, it shows that it is never a best-response for firm 2 to set a price lower than $\underline{p}$. She can raise her price to $\underline{p}$ while one is still best-responding with her price of one. It also creates a lower bound for firm 1's price. Even if firm 1 sold every time the buyer was not loyal to the firm 2, $Q_1 = m_{10} + m_{11}$, she would still make less profit for any price below $\underline{p}$ than if she just set a price of one. Therefore, both sellers will never drop their price below $\underline{p}$ in any equilibrium.

8. The best-response could also be found by constructing the marginal revenue (MR) curve. However, because MR is discontinuous, the optimal expected quantity is the maximum quantity

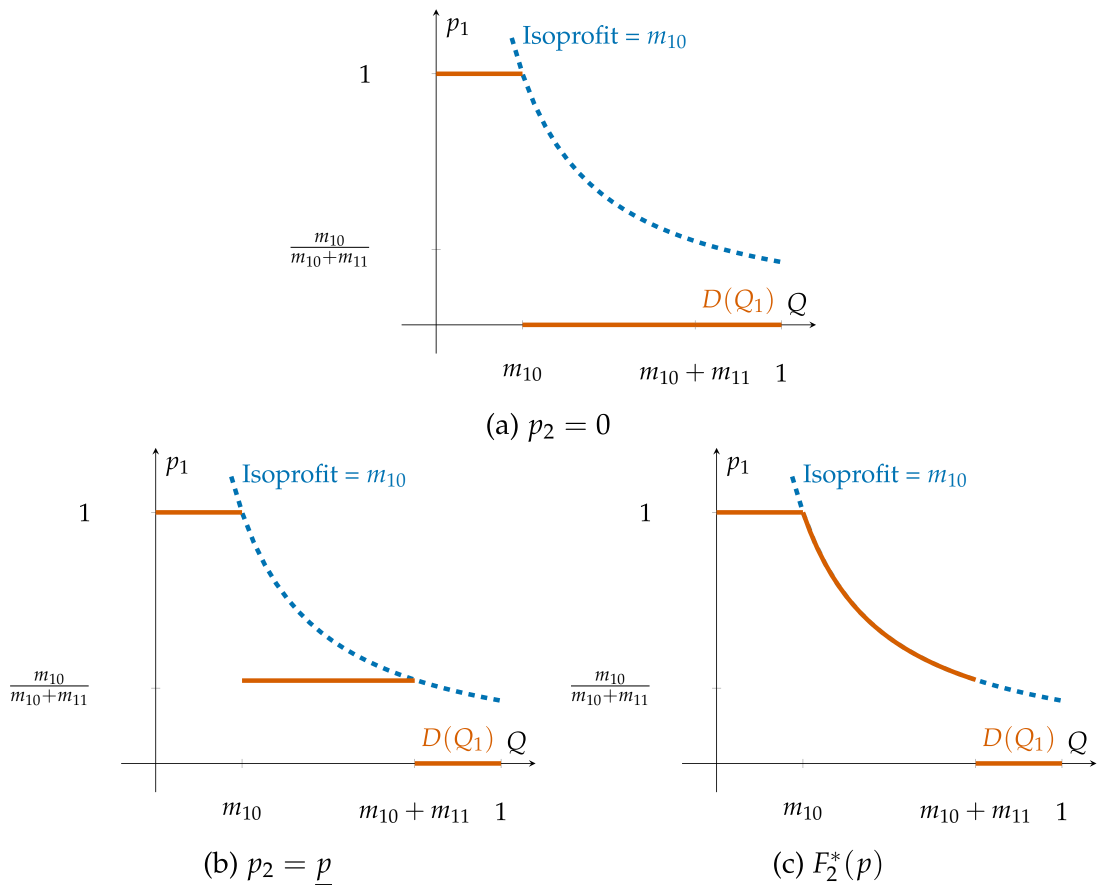

Figure 4.4: Firm 1's Residual Demand and Indifference Curves

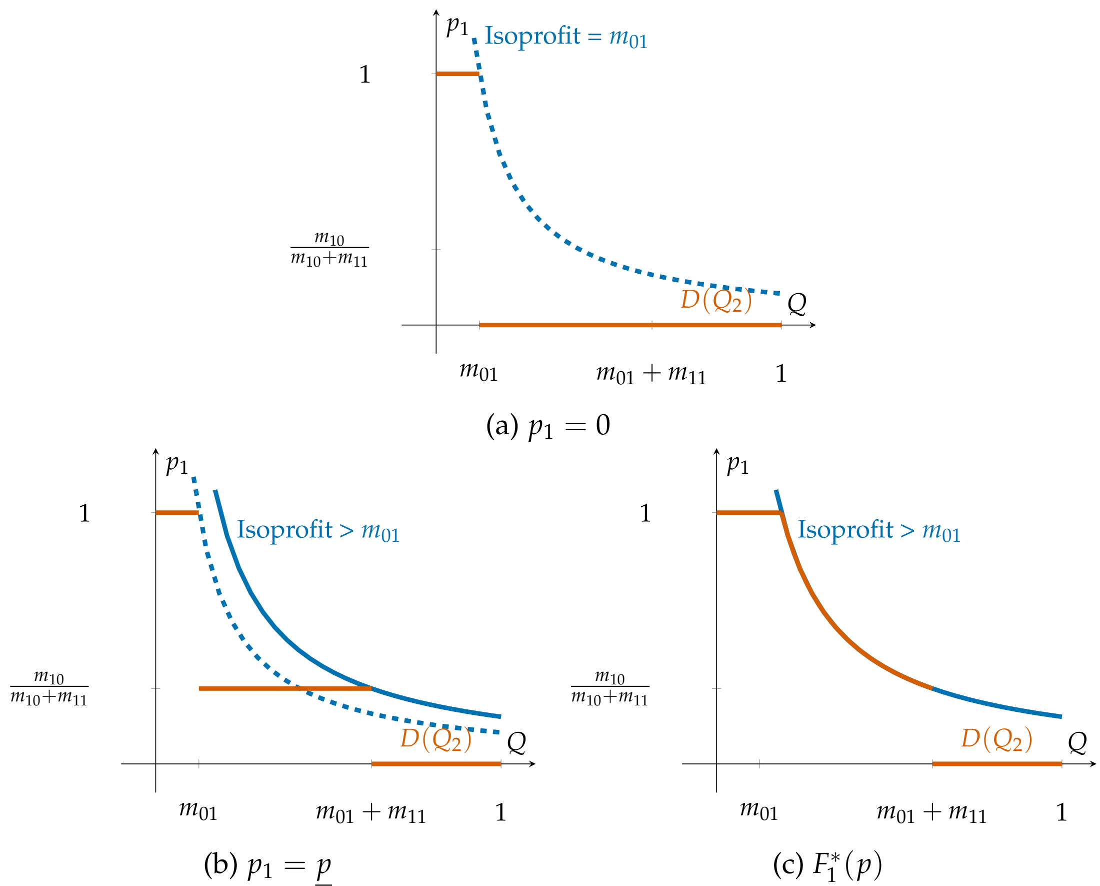

(b) $p_1 = \underline{p}$

(c) $F_1^*(p)$

Figure 4.5: Firm 2's Residual Demand and Indifference Curves

Firm 2 can still do better by randomizing between $\underline{p}$ and one. This is like setting a non-linear price and induces a residual demand curve for firm 1 that is no longer flat in the middle.[9] The demand curve varies with the cumulative distribution of firm 2's randomized strategy. Again, the demand curve for firm 1 is the price set by firm 2, but now in a probabilistic sense. The residual demand curve for the indifferent buyer is simply one minus the CDF of firm 2's randomized price.

The only complication is a specific price is set with positive probability. Better-replying by firm 2 with randomization means increasing firm 1's demand curve for all prices between $\underline{p}$ and one such that firm 1 is still indifferent. Figure 4.4c shows the demand curve where firm 1 is indifferent between a range of prices, which is generated by a particular pricing strategy. This pricing strategy is also firm 2's equilibrium pricing strategy, as explained in Lemma 2.[10]

However, even with all of this mixing, neither can achieve more in equilibrium than winning their respective market with $\underline{p}$. The randomization protects against undercutting by the other seller but in equilibrium does not generate any additional profit, resembling the normal Bertrand force that drives profits to zero.[11]

---

such that MR > MC = 0, instead of the standard MR=MC condition used for continuous functions.

9. Using the indifference curve to trace out the optimal policy resembles one construction of the optimal income tax rate in an optimal income tax model, like Mirrlees (1971).

10. If the parties are symmetric ($m_{01} = m_{10}$), each seller receives her loyal customers and we are back to perfect competition, where each person receives her marginal product (Ostroy 1984), as is standard for symmetric Bertrand competition.

11. This also resembles matching pennies or the information design model in Albrecht (2017) (Chapter 5), where there is political competition with discontinuities and randomization prevents the other party from winning. However, randomization does little to the odds of winning in equilibrium. Both matching pennies and political competition are zero-sum games. With this is technically not zero-sum, it shares some characteristics.

This process of finding the mixed strategy is standard. We search for the strategy of firm 1 that makes firm 2 indifferent (willing to mix) and vice versa. But we are looking for a mixed strategy over a continuum, so it is not immediately easy to visually. That is where the residual demand and iso-profit curves come in. The demand curves immediately allow the economist to picture a continuous variable in a natural way.

4.2.3 Consumer and Firm Optimal Public Data

It is worth pointing out a corollary of the above proposition: no matter the distribution of types, complete data on consumers is always better than no data on consumers.

Corollary 8.2. *1) Consumer surplus under perfect price discrimination is weakly higher than consumer surplus under no price discrimination. 2) Total profits under perfect price discrimination are weakly lower than total profits under no price discrimination. 3) The relationships are strict if* $m_{10} \neq m_{01}$.

Notice that we have already done the work to solve the general model of public information. The fact that this is the "aggregate" market is not important. The equilibrium price distributions will be the same if the market was the result of some consumer data signals. Therefore, we can immediately plot the consumer surplus for any market in the simplex, as shown in Figure 4.6a, or the expected price as shown in Figure 4.6b

In addition, we can search for the buyer optimal (lowest expected price) and seller optimal (highest expected price) information structures. Proposition 9 proves

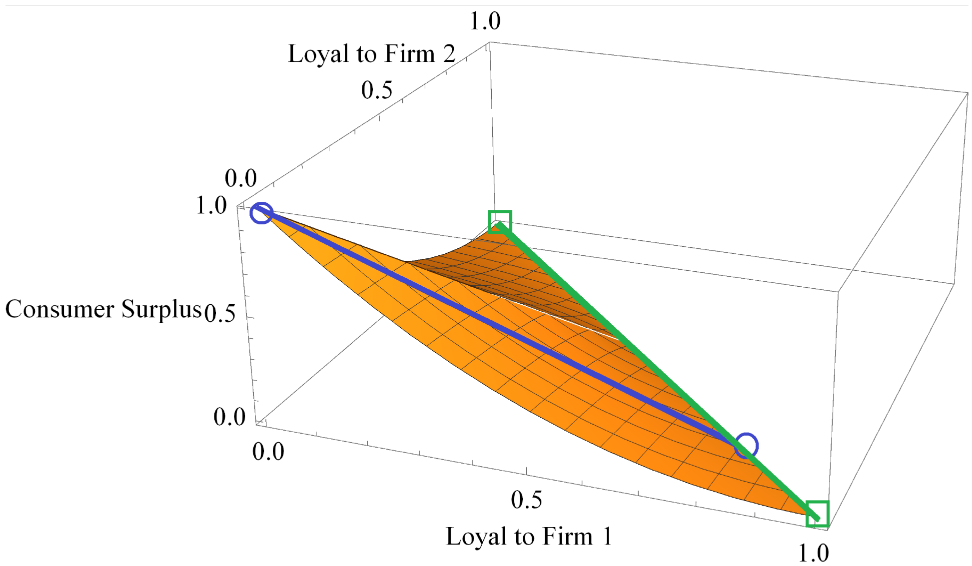

(a) Consumer Surplus

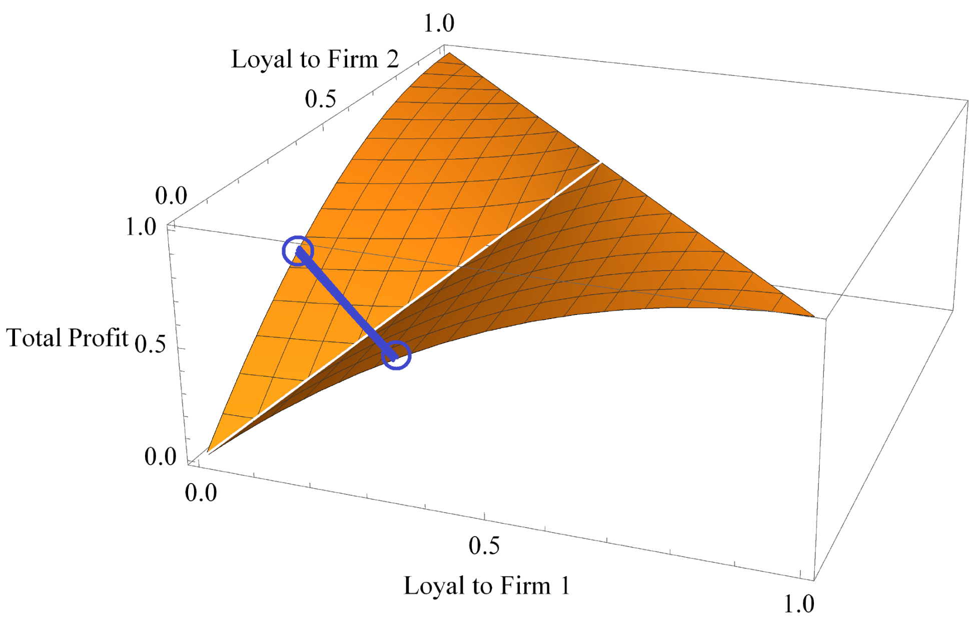

(b) Total Profit

that complete information is consumer optimal.

Proposition 9. *With only public data, consumer surplus is maximized under perfect price discrimination.*

Proof. Bayes' Law only restricts the ex-post distributions of types to sum to the prior. Therefore, we can maximize/minimize over all feasible distributions (markets) by examining the "concavification"(Kamenica and Gentzkow 2011) of the value function of interest: consumer surplus.

The profit minimizing, and therefore consumer surplus maximizing, is given by the concave envelope on the function. Because the total profit is convex on each half of the simplex, the concave envelope touches the convex function at extreme point, with one of the split markets being just indifferent consumers, $(0, 0, 0)$ and one of the markets being the combination of loyal customers $(\frac{m_{10}}{m_{10}+m_{01}}, \frac{m_{01}}{m_{10}+m_{01}}, 0)$. The splitting is given by the blue line and circles. However, because firm 1 can ignore firm 2's loyal customers, that market can be further split according to the green line and squares, so that the effective profit-maximizing markets are complete information: $(1, 0, 0), (0, 1, 0), (0, 0, 1)$.

Therefore, we have shown that complete information is optimal for consumers.

■

The total profit-maximizing data structure can similarly be found by concavification. First, let us follow Armstrong and Vickers (2019a, 2020), and define a *nested market* as a market where one firm's potential customers are a subset of the other firm's. On the simplex, the set of nested markets are the two edges: $(x, 0, 1 - x)$ and $(0, y, 1 - y)$. We then have the following proposition:

Proposition 10. *With only public data, total profits are maximized with nested markets.*

Proof. The proof is immediate from concavification and shown by Figure 4.6b as the blue line and circles. ■

Everything so far has involved public data on consumers. The next section extends the model to allow for private information by each firm.

4.3 Private Consumer Data

Most models of market competition assume firms have public information, as above. However, increases in data collection make this assumption less tenable. I am a different person according to Amazon vs. Target. Amazon knows that I am a South Minneapolis male who buys too many economics books. Target only knows that I am a person shopping in Minneapolis on Tuesday. Still, the firms' pricing strategies will be intertwined and part of an equilibrium; Target still needs to consider Amazon's pricing.

To understand this type of asymmetric consumer data, let us consider a simple case where firm 1 has complete information and firm 2 has only aggregate information. Again, each market for a firm is a point on the simplex. However, now higher-order beliefs matter. There is no longer an objective "market", just overlapping markets. Firm 2 knows that firm 1 knows the true type. Firm 2 only needs to consider competition from firm 1 when the true type is indifferent. For pricing, firm 2's relevant market is made of its loyal consumers and indifferent. The markets are plotted in Figure 4.7.

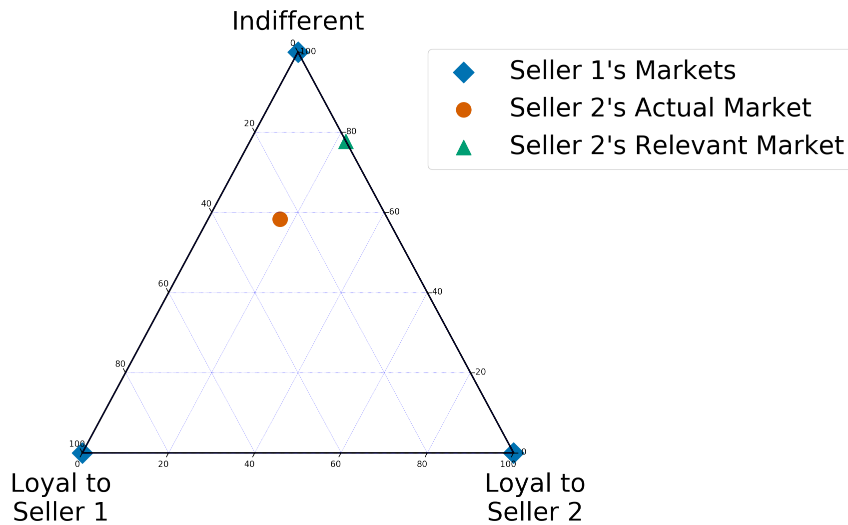

Figure 4.7: Overlapping Markets

Given the profit functions characterized in Lemma 2, it is easy to show that, relative to no consumer data, when firm i has access to consumer data that allows perfect price discrimination, firm i's profit strictly increases, firm j's profit strictly decreases, total profit can either increase or decrease. However, if the firms are symmetric ($m_{10} = m_{01}$), then profits increase and consumer surplus decreases.

Proposition 11. *For any symmetric interior market, perfect price discrimination by only one firm strictly decreases consumer surplus relative to no price discrimination.*

Comparing Proposition 11 to Corollary 8.2 highlights the interplay of competition and information. The powers of competition to drive down consumer prices are harnessed when both firms have access to information. When only one firm has access, the firm uses that information for rent-extraction as in the monopoly case.

But that is still a very simple information structure. In general, overlapping markets can be much more complex and involve correlated information, where the signal that firm i receives is correlated with the signal that firm j receives. The following proposition shows that the worst case for consumers is imperfectly correlated data, which effectively allows the firms to collude and avoid Bertrand competition.

Proposition 12. *Imperfectly correlated data can strictly increase prices relative to any public information structure.*

Proof. Imagine an information designer[12] who reveals consumer data to the firms and recommends an incentive-compatible price. The designer commits the following information for the indifferent consumers, only partially revealing the indifferent consumers.

		Signal to Firm 2: Reveal	Signal to Firm 2: Do Not
Signal to Firm 1	Reveal	$1 - \alpha_1 - \alpha_2$	α_1
	Do Not	α_2	0

For now, fix $\alpha_1 \geq \alpha_2$ and assume that the designer recommends $p = 1$ when a firm sees no signal and recommends a corresponding mixed distribution when it is revealed the customer is indifferent. We will construct an equilibrium that increases profit relative to any public signal.

Consider the case when firm 1 receives a signal and therefore knows the consumer is indifferent. Even though firm 1 knows the consumer is indifferent, prices

12. I stick to the term information designer to highlight the similarities to questions in the information design literature (Bergemann and Morris 2019), but this could also be thought of as mediator, as is used in other models of collusion, *e.g.* Rahman and Obara (2010) and Rahman (2014).

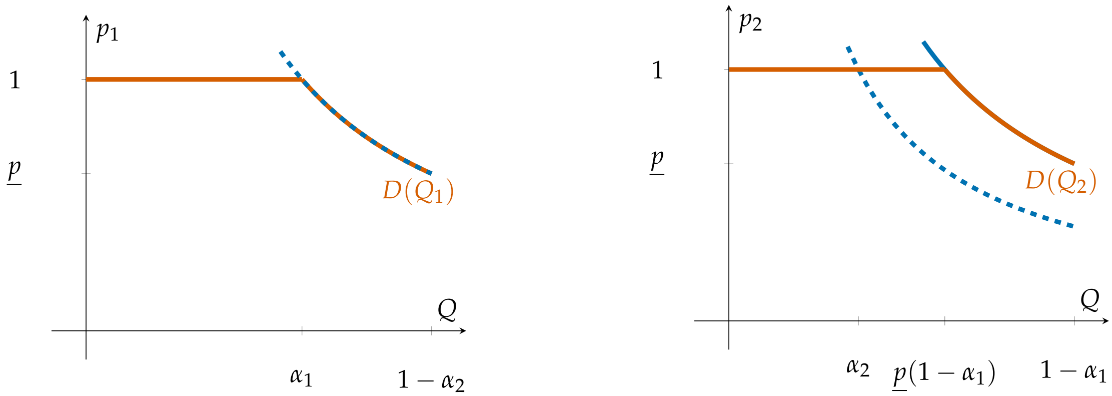

(a) Firm 1

(b) Firm 2

Figure 4.8: Residual Demand after Signal

will not be driven to zero because firm 1 does not know whether firm 2 knows that the consumer is indifferent. With probability α_1, firm 2 did not receive a signal and will, therefore, follow the designer's recommend strategy and set a price of 1. Otherwise, firm 2 will price according to a distribution over $[\underline{p}, 1]$. How high can the information designer raise firm 2's while still keeping firm 1's behavior incentive compatible? Just as Figure 4.3a, firm 2 can price according to a distribution that makes firm 1 indifferent between setting any price between firm 2's lowest price $\underline{p}$ and 1.

We can do the same thing for when firm 2 receives a signal. However, because firm 1 is only willing to reduce the price to $\underline{p}$, firm 2 never needs to price any lower. This means that Firm 2 can receive a higher profit than his maximin. This is shown in Figure 4.8b and mirrors the public information case shown in Figure 4.5b.

Now that we have calculated $\underline{p}$ and the best-response when each firm receives a signal, we only need to verify that setting a price of 1 is a best-response when each firm does not receive a signal. First, consider the case of firm 1 when he sees

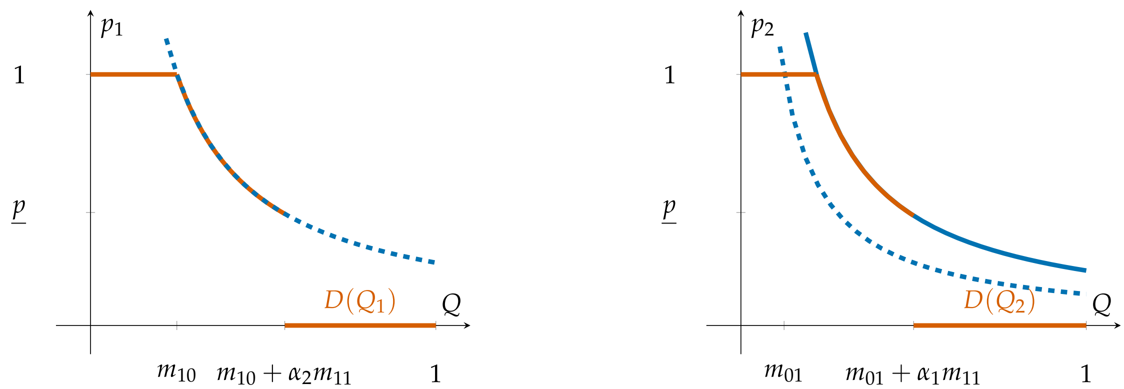

(a) Firm 1

(b) Firm 2

Figure 4.9: Residual Demand after No Signal

no signal. That could be because the consumer is loyal to firm 1 (which happens with probability m_{10}) or because the information simply was not revealed (which happens with probability $\alpha_2 m_{11}$). The residual demand when the firm does not receive a signal is graphed in Figure 4.9. This exact distribution depends on the parameter α_2. We have therefore constructed the strategies so that it is a best-response for firm 1 to keep a price of 1 and firm 2 is also maximizing his profit. The same argument goes through for when firm 2 receives a signal.

To review, fixing α, when the firms receive a signal, they must want to set a distribution of prices. The relevant ICs (from Figure 4.8) are

$$\underline{p} \geq \frac{\alpha_1}{1 - \alpha_2} \quad \text{and} \tag{4.1}$$

$$\underline{p} \geq \frac{\alpha_2}{1 - \alpha_1}, \tag{4.2}$$

where at least the first constraint is binding.

When the firm receives no signal, setting a price of 1 must be a best-response.

In principle, we need to worry about deviating to any other price, but for simplicity let us drop the interior prices, knowing the distribution can be constructed to make sure those are satisfied. The two incentive constraints (seen from Figure 4.9) are therefore

$$m_{10} \geq \underline{p}\,(m_{10} + \alpha_2 m_{11}) \quad \text{and} \tag{4.3}$$

$$m_{01} \geq \underline{p}\,(m_{01} + \alpha_1 m_{11})\,, \tag{4.4}$$

where at least the first constraint is binding.

Therefore, we have solved for the equilibrium as functions of α_1 and α_2. Now we can simply add the profit from each signal to find the total expected profit/price:

$$\text{Expected Price} = \underbrace{\alpha_1}_{\text{Figure 4.8a}} + \underbrace{\frac{\alpha_1}{1-\alpha_2}(1-\alpha_1)}_{\text{Figure 4.8b}} + \underbrace{m_{10}}_{\text{Figure 4.9a}} + \underbrace{\frac{\alpha_1}{1-\alpha_2}(m_{01}+\alpha_1 m_{11})}_{\text{Figure 4.9b}}$$

To maximize prices, the information designer's problem is to choose α to maximize profits, subject to the incentive constraints.

■

With all of these different types of consumer data consider, we can plot the distribution of equilibrium prices for the six cases in Figure 4.10. The expected price and profit is the area above the distribution. Therefore, consumer surplus is the area under the distribution. The vector of surplus for each case is plotted in Figure 4.11. Again, these are for an aggregate market of $\left(\frac{1}{4}, \frac{1}{6}, \frac{7}{12}\right)$.

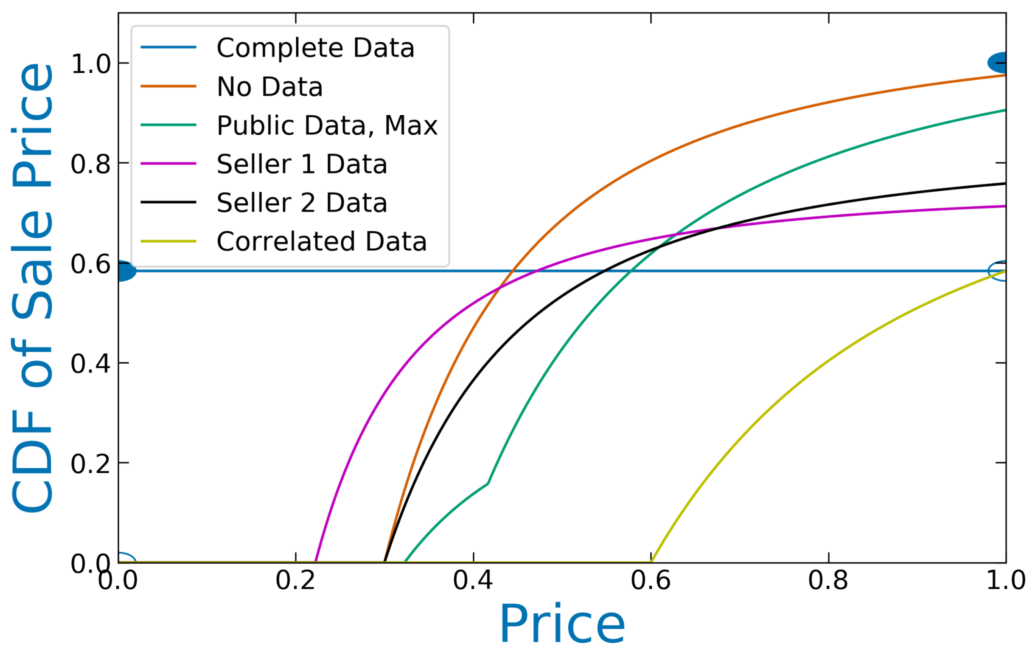

Figure 4.10: Price Distribution

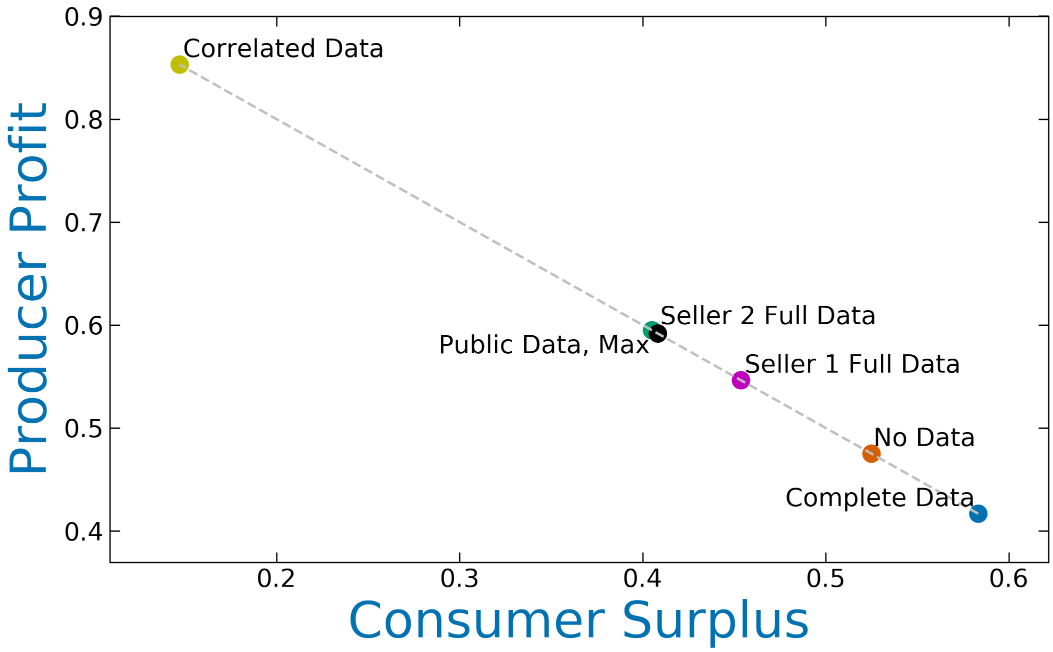

Figure 4.11: Surplus Division

4.4 Conclusion

This chapter studies a version of Bertrand competition where firms may have access to more detailed consumer data. In equilibrium, each firm chooses a strategy which is a distribution of strategies that induces a unit-elastic residual demand curve for the other firm. This makes each firm indifferent over a range of prices and ensures an equilibrium, given some uncertainty.

I then analyze the welfare consequences of the consumer and duopolists having additional information about the consumer's tastes. With additional information, the firms can use that information to set different pricing strategies, i.e. use third-degree price discrimination. I solve for the consumer-optimal information structure: completely revealing data. While firms are able to use data are able to raise its price above marginal cost and thus exploit their market power through price discrimination when the consumer prefers one good, complete information unleashes the power of competition when firms must ruthlessly compete on price. Competition works best in full light.

Chapter 5

Political Persuasion

> "Voice is political action par excellence."
>
> Albert Hirschman 1970

Political campaigns are filled with talking: at rallies, in commercials, or in debates. If there is any information in all that talking, political parties will want to use the information to win. They can talk their way to victory in a combination of two ways. First, a party can convince voters that their candidate will be fantastic at doing what the voter wants. Second, a party can convince voters that their candidate is not as bad as the other party's candidate. The optimal strategy depends on what voters initially believe about both candidates and the other party's campaign strategy.

To better understand how politicians convey information to voters and how the information depends on the competition inherent in political campaigns, I construct a game of information disclosure with two senders and a continuum of receivers. The senders are political parties–Left Party and Right Party–and the receivers are voters. First, in the game, each party decides what information to reveal about its candidate. In practice, this is like committing to a campaign schedule: parties pick debate days, choose the primary schedule, etc. After committing to a campaign schedule, specific information is revealed about the candidate. How the voter perceives the information is beyond the direct control of the party after each party has set up its campaign. The news media reports what comes out and the voter infers something about the candidates. Lastly, the voter uses her information and picks the winner. I model information disclosure formally as a Bayesian persuasion game (Kamenica and Gentzkow 2011).

Since there are two senders competing, each sender has to consider (1) how to influence directly the beliefs of the receiver and (2) the actions of the other sender. For each party, their campaign strategy is equivalent to choosing a distribution of voter beliefs. More than that, the game is zero-sum since any gain that one party generates from persuasion is a direct loss to the other party. I will sometimes refer to one party as the *maximizer* (Left) and the other as the *minimizer* (Right). The maximizer chooses a campaign strategy to maximize the probability that it wins.[1] The minimizer chooses a campaign strategy to minimize the probability that it loses.

The main result of the chapter characterizes the unique equilibrium (Theorem 13). In equilibrium, each party picks a campaign that distributes the voter's

1. Kamenica and Gentzkow (2011) show that the maximized probability is the concave closure of the probability of winning for the realized beliefs. The geometry comes from Aumann and Maschler (1995, 128)

beliefs uniformly. By uniform, I mean that when the voter is choosing the winner after the campaign, the voter is equally likely to believe the candidate is good with a probability of $0.010, 0.011, 0.011$, etc, and everything in between. The uniform distribution, while appearing strange, makes sense after a closer look. First, it cannot be that in equilibrium either party chooses a campaign strategy that generates positive mass on certain beliefs (Lemma 4). If the Right Party generates positive mass on some point, the Left Party can beat that mass of beliefs by ϵ and increase its probability of winning discontinuously. Winning by ϵ means convincing the voter that your candidate is better than the other candidate, even if both candidates are bad. The only place this argument does not hold is at the highest possible belief since there is no ϵ above that (Lemma 5). If the distribution of beliefs about the Right Party's candidate is continuous, the Left Party does not have such an easy target. Therefore, it must be that in equilibrium strategies are continuous distributions almost everywhere.

It is not just discontinuities that the Right Party will guard against. The Left Party benefits from persuasion if the Right Party's strategy is not concave. This *concavification* argument comes directly from Kamenica and Gentzkow (2011). If the Right Party's strategy is not concave, the Left Party can "concavify" and win with a higher probability. The Right Party wants to prevent that option by picking a strategy that is already concave or close to concave. The uniform distribution is concave. A uniform strategy eliminates the ability of the Left Party to win by ϵ or concavify.

The uniform strategy closely resembles–in some cases is identical to–political parties' strategies when they are trying to buy votes, such as in Myerson (1993) and Sahuguet and Persico (2006). In those papers, if the Right Party gives \$1 to a positive mass of the population, the Left Party could give that same group \$1 + ϵ and win all their votes. In equilibrium, the distribution of payments to voters is uniform. Besides the equilibrium, another aspect of the vote-buying literature is worth mentioning. In that literature, the parties commit to paying out campaign promises, even though they may want to renege once in office. The current chapter, by modeling information disclosure as Bayesian persuasion, also assumes that the parties can commit to their campaign schedule. The identical commitment assumption allows me to draw a direct parallel between vote-buying and persuasion.

The chapter highlights a similarity between the concavification argument of Aumann and Maschler (1995) and Kamenica and Gentzkow (2011) and the uniform strategies used in zero-sum games. The similarity should not be surprising. Aumann and Maschler (1995) derive their original concavification result in a repeated zero-sum game. Also, uniform strategies are concave for both players. While political competition through information disclosure is relatively nonstandard, the equilibrium strategies turn out similar to the more standard models.

The rest of the chapter is organized as follows. Section 5.1 goes through an example to relate concavification to uniform strategies. Section 5.2 lays out the formal model. I then turn the information disclosure game, which is between two senders and a receiver, into a zero-sum game between the two senders. I then derive the main theorem. Section 5.3 extends the equilibrium result to more general cases by harnessing concavification and zero-sum game results. Section 5.4 discusses related literature.

5.1 Example

Consider two political parties, creatively named the Left Party and Right Party and indexed by $i \in \{L, R\}$, and a continuum of voters. Each party has two possible types of candidates: good and bad. For this example, one-fourth of Left Party is good, one-half of the Right Party is good, and the party divisions are common knowledge. Each voter's action is whether to vote for the Left or Right Party. Voters are identical at the beginning of the game: they share the same preferences and beliefs. Electing a good candidate gives a utility of one and a bad candidate gives a utility of zero. However, voters do not directly observe each candidate's type. Instead, voters must rely on the campaign to release information about each candidate.

Neither party can choose its candidate, but draws their candidate at random.[2] Instead of choosing their candidate, each party's action is a *campaign schedule*. One can think of the campaign schedule as choosing how many ads to run, debates to have, etc. Before voting, each voter will see a realization of the signal, s. The party's action is a distribution of s that can be conditioned on the type of candidate chosen. More formally, a campaign for party i is a distribution that is conditioned on the type of the candidate, $\pi_i(\cdot|\text{good})$ and $\pi_i(\cdot|\text{bad})$, for some set of possible realizations. Call the strategies L and R for the two parties. For example, suppose that $s \in [0, 1]$.[3] I assume each party chooses one set of conditional distributions of

2. One way to justify the no strategic choice for candidates is to think of the random variable as being defined over a small subset of party insiders that makes up the potential candidates.

3. One possible strategy is a fully informative strategy where

$$\pi(0|\text{good}) = 0 \qquad \pi(0|\text{bad}) = 1$$
$$\pi(1|\text{good}) = 1 \qquad \pi(1|\text{bad}) = 0.$$

s for all the voters. However, each voter receives a private and independent draw from each party's distribution. Because the signal label s is arbitrary, let me label it by the belief it induces that candidate is good. This relabeling of signals is why $s \in [0, 1]$, compared to standard Bayesian persuasion models which only require finite signal realizations.[4]

After seeing π_L, π_R, s_L, and s_R, each voter updates her beliefs. Let q_L denote the probability that the Left candidate is good (his expected quality) and q_R for the Left candidate. If $q_L > q_R$, the voter picks the Left candidate and vice versa. For this example, if she is indifferent, she picks the Right Party.[5] Define $v(q_L, q_R)$ to be the probability that a voter with beliefs q_L and q_R votes for the Left Party. For simplicity of notation, I sometimes suppress the arguments of v, such as using $\mathbb{E}[v]$ for the expected probability that a voter picks for the Left Party. Because there is a continuum of voters, $\mathbb{E}[v]$ is also the expected probability that the Left Party wins the election. There is more on this point in Section 5.2. A party receives a payoff of one if they win the election and zero otherwise. Therefore $\mathbb{E}[v]$ is also the expected value of the election to the Left Party and $\mathbb{E}[1 - v]$ is the expected value for the Right Party.

To see how campaigns can be used, first suppose there is no campaign or equivalently both parties are playing uninformative strategies. A uninformative

---

4. In fact, because of the competition between the two parties and the form of the discontinuity in payoffs, $s \in [0, 1]$ is necessary for equilibrium. Each party needs to be able to randomize using a sufficiently rich space.

5. In principle, when the voter is indifferent she could condition her vote on anything since she is indifferent. To avoid complications in the example, I assume the Left Party wins ties. This turns out to be purely expository and not substantive since ties happen with probability zero in equilibrium. However, that is not true in general. Therefore, the equilibrium I use in Section 5.2 has each voter flipping a coin when she is indifferent.

campaign for the Right Party would be $\pi_R(s|\text{good}) = \pi_R(s|\text{bad})$ for all s. Then every voter chooses the Right Party, because $q_R > q_L$, and $\mathbb{E}[v] = 1$.

However, the Left Party can do better. Fix the uninformative strategy for the Right Party and consider the cumulative distribution function (CDF) of beliefs that the uninformative strategy, R, induces. This is plotted in Figure 5.1a. The Right Party's CDF induces an "interim" value function for the Left Party, $\mathbb{E}[v|R]$, where the conditional is on the Right Party's strategy.[6] If $q_L < q_R^0$, then the Right Party has a higher expected quality and the Left Party loses. That changes when $q_L \geq q_R^0$ when the Left Party wins. Figure 5.1b is useful, because it allows us to find the Left Party's best-response and the value it attains (by concavification), which is given in Figure 5.1c. From Kamenica and Gentzkow (2011), the best that the Left Party can do is the *concavification* of $\mathbb{E}[v|R]$.[7] In this case, the Left Party can choose the following strategy

$$\begin{aligned} \pi_L(0|\text{good}) &= 0 & \pi_L(0|\text{bad}) &= \frac{2}{3} \\ \pi_L\left(\frac{1}{2}|\text{good}\right) &= 1 & \pi_L\left(\frac{1}{2}|\text{bad}\right) &= \frac{1}{3}. \end{aligned} \tag{5.1}$$

Whenever $s = 1/2$, which is for all of the good candidates and one-third of the bad candidates, the voter believes the Left Party's candidate is just a likely to be good as the Right Party's and the Left Party wins (by the assumption in this example that voters break ties toward the Left Party). The Left Party, by best-responding to the uninformative strategy of the Right Party, uses persuasion to increases its

6. Recall that for a zero-sum game, by the Minimax Theorem the Right Party could move first without loss, as long as s is not observed before the Left Party moves (Owen 1982),

7. The concavification of a function f is the smallest concave function that is everywhere as great as f.

probability of winning from zero to one-half.

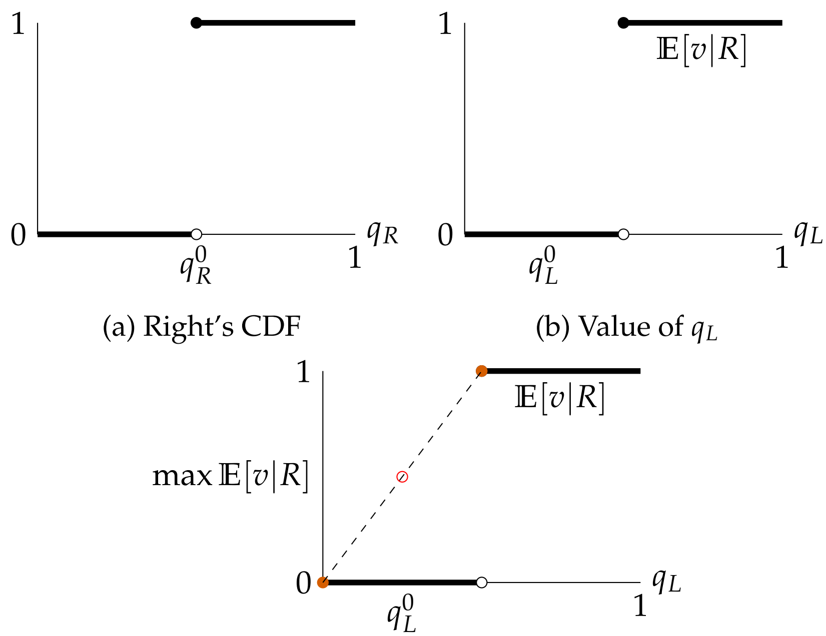

(a) Right's CDF (b) Value of q_L

(c) Left's Best-Response

Figure 5.1: Right plays uninformative, Left Best-Responds

Clearly, an uninformative strategy is not optimal for the Right Party. Now suppose the Right Party is not playing an uninformative strategy, as in Figure 5.1, but is using the following campaign:

$$\begin{aligned} \pi_R\left(\frac{3}{4}|\text{good}\right) &= \frac{1}{2} \quad \pi_R\left(\frac{3}{4}|\text{bad}\right) = \frac{1}{6} \\ \pi_R\left(\frac{1}{2}|\text{good}\right) &= \frac{1}{3} \quad \pi_R\left(\frac{1}{2}|\text{bad}\right) = \frac{1}{3} \\ \pi_R\left(\frac{1}{4}|\text{good}\right) &= \frac{1}{6} \quad \pi_R\left(\frac{1}{4}|\text{bad}\right) = \frac{1}{2}. \end{aligned} \tag{5.2}$$

The CDF that strategy 5.2 induces is given in Figure 5.2a. The Left Party's best-response now can only win with a probability of one-third. Notice that there multiple best-responses for the Left Party. First, the Left Party could be uninformative. Then it wins every time is matched with s_1, which happens with probability 1/3. Similarly, it could play a strategy that randomizes between zero and one half. In either case, the best the Right Party can do is to win at the concavification of the Left Party's strategy.

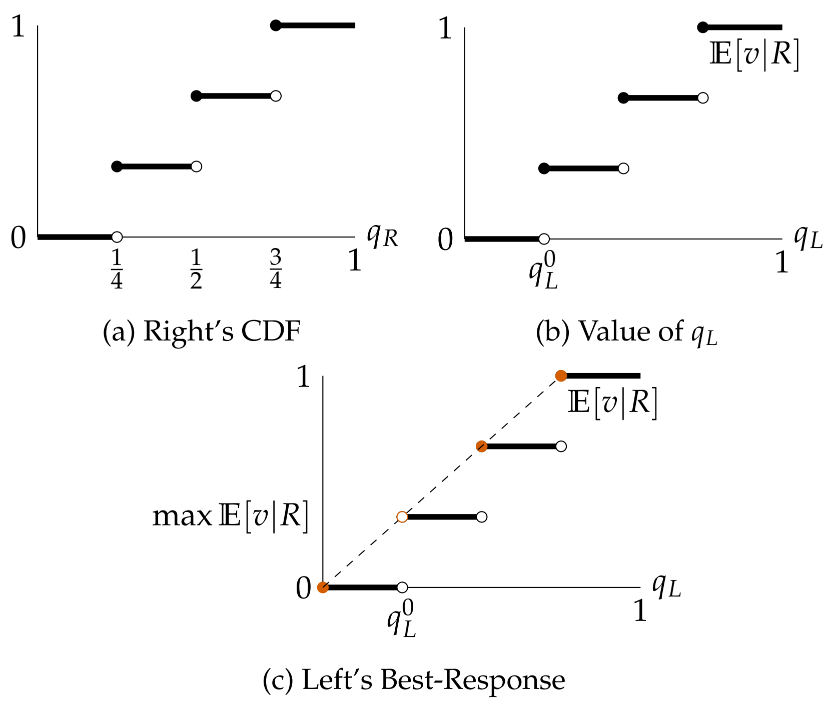

Figure 5.2: Left Party using Persuasion

By using some persuasion and playing a "better-response", the Right Party chooses a CDF that lowers the concavification. The Right Party's best-response is to minimize the concavification, that is, to minimize the value attained by the Left-Party's best-response. In equilibrium, one party minimizes and the other party maximizes.

The unique equilibrium strategies are the campaigns that induce beliefs given in Figure 5.3a for the Right Party and Figure Figure 5.3d for the Left Party. To verify this is an equilibrium, assume the Right Party plays a uniform distribution, generating the beliefs in Figure 5.3a. The best-response of the Left Party leads to the Left Party winning with a probability of one-fourth. Therefore, the Right Party can guarantee that the Left Party wins with probability less than or equal to one-fourth.

Now go to the second row. The Left Party cannot play a pure normal distribution, because it has a low prior probability of being good. Instead, with probability one-half, the Left Party induces a belief of zero, but the other half is a uniform distribution over $(0, 1)$. The Right Party's best-response leads to the Right Party winning with probability three-fourths. Therefore, the Left Party can guarantee that the Right Party wins with probability less than or equal to three-fourths. By the Minimax Theorem, this is an equilibrium. Uniqueness is left for the proof of Theorem 13.

The concavification argument implies the uniform strategy of zero-sum games. Any failure of concavity in the Right Party's strategy, such as when it is discontinuous, can be concavified by the Left Party. The concavification benefits the Left Party. Since the game is zero-sum, concavification by the Left Party hurts the Right Party. Instead of giving that advantage to the Left Party, the Right Party is able to remove any ability to concavify by using the uniform strategy. The next section extends the example to any prior probabilities and characterizes equilibrium.

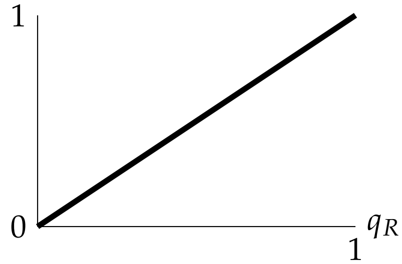

(a) Right's CDF

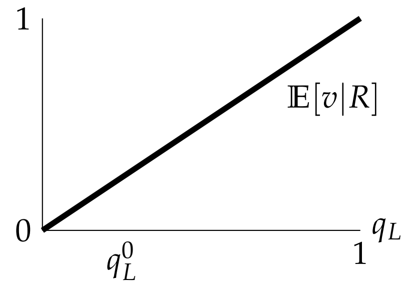

(b) Value of q_L

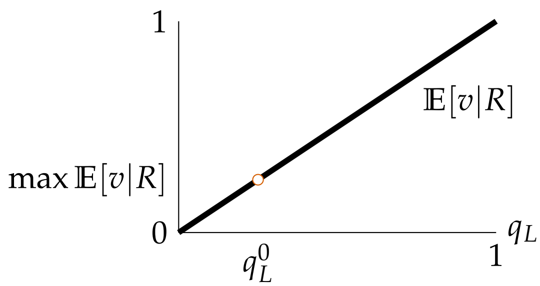

(c) Left's Best-Response

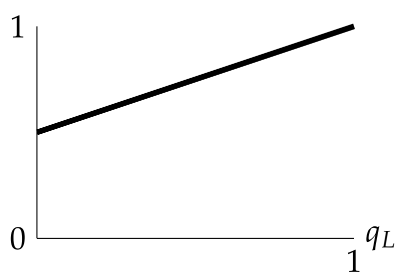

(d) Left's CDF

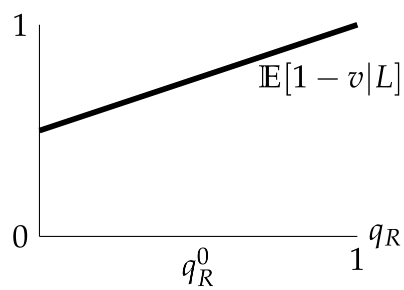

(e) Value of q_R

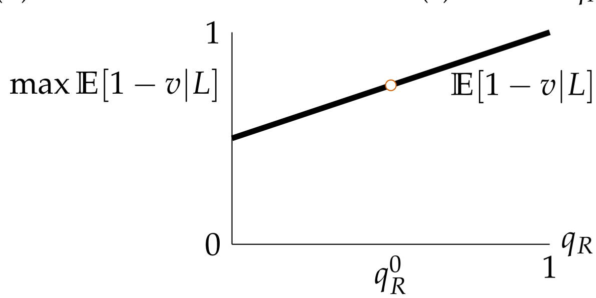

(f) Right's Best-Response

Figure 5.3: Equilibrium

5.2 Model

The previous example suggests that uniform strategies prevent the other party from using persuasion to concavify. I now develop a general model to formalize this intuition and characterize the unique equilibrium.

There are two political parties, indexed by $i \in \{L, R\}$. There is a continuum of voters who each vote for one of the parties. The winner is the party with a majority of votes. Each party receives utility one if they win the election and zero otherwise. Each party has a candidate who competing for office and there is uncertainty about whether each party's candidate will be good for the voter, i.e. the quality of the candidate. Each candidate can either be bad or good quality for the voter. The parties and voter have a shared prior belief about whether each party's candidate will be good. For the Left Party, the candidate is good with a probability of $q_L^0 \in (0,1)$. For the Right Party, the probability is $q_R^0 \in (0,1)$. These realizations of candidate quality are independent across the parties. If the candidate is bad, the voter receives a utility of zero. If the candidate is good, the voter receives a utility of one.[8]

Each party's strategy is a *campaign schedule* and denoted by π_i. A campaign is made up of signal realizations, $s_i \in [0,1]$, and a set of conditional distributions, $\{\pi_i(s_i|q)_{q\in\{0,1\}}\}$.[9] The parties choose their campaigns simultaneously.

After seeing each π_i and s_i, each voter updates her beliefs. Denote a particular

8. The restriction to two types for each party is not without loss. See Gentzkow and Kamenica (2016a). This is elaborated on in Section 5.3.

9. Kamenica and Gentzkow (2011) assume that S is finite. As Lemma 1 shows, because each party's payoff is discontinuous, there will be no equilibrium when S is finite. I need to allow for a larger signal space to smooth out the discontinuities.

pair of realized posteriors $(q_L, q_R) \in [0,1]^2$. From any belief, each voter can calculate the expected candidate quality. While voters could vote based on anything in π_L, π_R, s_L and s_R, I assume–or more precisely study equilibrium where–voters only condition their vote on the expected candidate qualities. Each voter chooses a probability of voting for the Left Party, denoted by $v(q_L, q_R)$.[10]

The timing of the game is as follows: (1) both parties choose their campaign, (2) the voters observe both campaigns and their personal signal realizations, (3) the voters choose a candidate, and (4) payoffs are realized.

Instead of working directly with the set of campaigns, it is more convenient to consider distributions of posterior beliefs that a party chooses to induce. Any strategy π_i and a realization s_i induces a particular posterior belief, q_i. Before knowing any particular realization, any strategy generates a distribution of posteriors. Because each voter acts only on her posterior beliefs, each campaign schedule is equivalent choosing a *belief distribution* for a strategy. Party i chooses a weakly-increasing function $F_i(q_i)$ such that $F_i(0) \geq 0$ and $F_i(1) = 1$. Note that $F_i(0)$ can be greater than zero and the distribution can be discontinuous; mass points are allowed. The only additional restriction that Bayes' Rule imposes on the strategy is that the expected posterior equals the prior:

$$\int_0^1 q_i dF_i(q_i) = q_i^0. \tag{5.3}$$

A Bayes-plausible distribution can also be described using the integral of the

10. In equilibrium, this only restricts the voter when they are indifferent, i.e. $q_L = q_R$.

CDF[11]

$$\int_0^1 F_i(q_i)dq_i = 1 - q_i^0. \tag{5.4}$$

Any distribution of posteriors satisfying Equation 5.4 is called *Bayes-plausible*. The party choose a distribution out of the set of Bayes-plausible distributions. Because Equation 5.4 involves the CDF, it is more natural to work with the CDF when finding equilibrium, as shown in Figure 5.3 . Each party must choose a strategy such that the area under F_i is equal to $1 - q_i^0$.

Equilibrium requires that each party chooses a distribution of posteriors that maximizes the probability of winning the election, given the other party's strategy and the voter's strategy.

Equilibrium: The equilibrium concept I use is *fair subgame perfect equilibrium*[12]. Equilibrium is a strategy of the Left Party, $F_L(q_L)$, Right Party, $F_R(q_R)$, and the voter, $v(q_L, q_R)$, that satisfies four conditions:

1. given Right Party's and the voter's strategies, the Left Party chooses a Bayes-plausible distribution to maximize $\mathbb{E}[v(q_L, q_R)]$;

2. given the Left Party's and the voter's strategy, the Right Party chooses a Bayes-plausible distribution to minimize $\mathbb{E}[v(q_L, q_R)]$;

11. From integration by parts of the Riemann-Stieltjes integral Addison Wesley 2th Edition, 5th printing Page 144:

$$\int_a^b x dF(x) = bF(b) - aF(b) - \int_a^b F(x)dx = 1 - \int_0^1 F(x)dx.$$

12. Kamenica and Gentzkow (2011) use *sender-preferred* subgame perfect equilibrium which means that when the receiver is indifferent, she picks what the sender wants. Since there are two senders, their justification does not apply.

3. the voter chooses the candidate with the highest expected quality;

4. if the voter is indifferent, she votes based on the toss of a fair coin.

The fourth condition is why I refer to the equilibrium as *fair*. Focusing on fair subgame perfect equilibria allows us to directly talk about the probability that the Left Party wins given any posterior beliefs.[13] That is, in any fair subgame perfect equilibrium

$$v^*(q_L, q_R) = \begin{cases} 1 & q_L > q_R \\ \frac{1}{2} & q_L = q_R \\ 0 & q_L < q_R. \end{cases}$$

The probability of voting for the Right Party is $1 - v^*(q_L, q_R)$.[14] Let p denote the probability that a randomly selected voter would vote for the Left Party. In Riemann-Stieltjes integral notation, $p = \int_0^1 F_R(x) dF_L(x)$. With a continuum of voters, the

Lemma 3. *In equilibrium, the CDF is continuous over* $[0, 1)$.

Take parts of Lemma 5 and combine.

Lemma 4. *Fix the strategy for one party that uses a finite set of posteriors. If the other party wins with probability one by using an uninformative signal, the uninformative signal is a best-response. Otherwise, there is no best-response.*

13. If we did not restrict attention to fair subgame perfect equilibria, the voter's decision would not be pinned down by beliefs, but instead it would be an equilibrium object.

14. The asterisk is to denote that I am looking at fair subgame perfect equilibrium. Notice that in other equilibrium concepts, the voter needs to flip a coin. She can do any action when $q_L = q_R$.

Lemma 4 rules out a whole class of strategies from ever being part of an equilibrium. Even though there is no best-response for *some* strategies of the other party, there is a best response to other strategies. If the Right Party chooses a continuous distribution, then the value function for the Left Party is continuous and a best-response does exist. Lemma 5 formalizes this.

Lemma 5. *Fix the strategy for one party that is continuous on $[0, 1)$. Then the other party has a best-response, although it need not be unique. Any best-response achieves the concave closure of the other party's strategy.*

Proof. Fix a $F_R(q_R)$ that is continuous $[0, 1)$. Because $F_R(q_R)$ is a CDF, $F_R(q_R) \leq 1$ and $R(1) = 1$. This implies that $F_R(q_R)$ is upper-semicontinuous on $[0, 1]$. This implies that $v^*(q_L; F_R(q_R))$ is also upper-semicontinuous because of the mapping from $F_R(q_R)$ to $v^*(q_L; F_R(q_R))$. We can therefore invoke Proposition 1 in Kamenica and Gentzkow (2011), which proves that the optimal value is attained at the concave closure of $v^*(q_L; F_R(q_R))$ and it therefore exists. ■

If $F_R(q_R)$ is continuous, there is no longer a discontinuous increase in the probability of winning at any induced belief. The Right Party effectively takes away the "easy" win by an arbitrarily small margin. With a best-response existing for some strategies, equilibria have a chance of existing, even though $v^*(q_L, q_R)$ is discontinuous. One does.

Fixing the strategy of the voter, the model becomes a two-player, zero-sum game where the voter generates the payoff function, $v^*(q_L, q_R)$. Again, the possible strategies for the players are any Bayes-plausible distributions.[15] The full

15. This simplification allows me to harness the zero-sum game literature to solve for the equilibrium. In particular, the persuasion game becomes a variant of Colonel Blotto games, which Hart

strategy is characterized in Theorem 13, which is the main result of the chapter.

Theorem 13. *Let $q_L^0 \geq q_R^0$. There exists a* unique *equilibrium. In equilibrium, the Left Party wins with probability $1 - \frac{1}{2}\frac{q_R^0}{2q_L^0}$.*
If $q_L^0 \leq 1/2$, the strategies are

$$L^*(q_L) = \begin{cases} \frac{1}{2q_L^0} q_L & q_L \leq 2q_L^0 \\ 1 & q_L > 2q_L^0 \end{cases} \qquad R^*(q_R) = \begin{cases} \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right) \frac{1}{2q_L^0} q_R & q_R \leq 2q_L^0 \\ 1 & q_R > 2q_L^0. \end{cases} \tag{5.5}$$

If $q_L^0 > 1/2$, the strategies are

$$L^*(q_L) = \begin{cases} \frac{1}{2q_L^0} q_L & q_L \leq 2(1-q_L^0) \\ \frac{1}{q_L^0} - 1 & q_L \in (2(1-q_L^0), 1) \\ 1 & q_L = 1 \end{cases} \qquad R^*(q_R) = \begin{cases} \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right) \frac{1}{2q_L^0} q_R & q_R \leq 2(1-q_L^0) \\ \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right)\left(\frac{1}{q_L^0} - 1\right) & q_R \in (2(1-q_L^0), 1) \\ 1 & q_R = 1. \end{cases} \tag{5.6}$$

Proof. First consider the case where $q_L^0 \leq 1/2$. This includes the example from Section 5.1, where $q_L^0 = q_R^0 = 1/2$. Suppose the Right Party is using $R^*(q_R)$ from Equation 5.5. With continuous probability distributions, the probability that the

(2008) has named a continuous "General Lotto" game.

Left Party wins is the probability that it generates $q_L \geq q_R$:

$$\begin{aligned}
P(q_L \geq q_R) &= \int_0^1 R^*(x)dL(x) \\
&= \int_0^1 \min\left\{ \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right)\frac{1}{2q_L^0}x, 1\right\} dL(x) \\
&\leq \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right)\frac{1}{2q_L^0}\int_0^1 xdL(x) \qquad (5.7)\\
&= \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right)\frac{1}{2q_L^0}q_L^0 \quad \text{(Bayes-plausible)} \\
&= 1 - \frac{q_R^0}{2q_L^0}.
\end{aligned}$$

With $R^*(q_R)$, the Right Party guarantees that the Left Party wins with probability less than $1 - q_R^0/2q_L^0$. It minimizes the upper-bound on how much the Left Party can win. Geometrically, the uniform distribution minimizes the concavification. Therefore, the value is less than $1 - q_R^0/2q_L^0$.

Now assume the Left Party is following the given strategy. Do the same calculation as Equation 5.7 for $P(q_R \geq q_L)$ to find that it is less than or equal to $q_R^0/2q_L^0$. The Left Party guarantees that the Right Party wins with probability less than $q_R^0/2q_L^0$. Therefore, the value is greater than $1 - q_R^0/2q_L^0$.

By the Minimax Theorem, if the guaranteed payoff for the Left Party equals one minus the guaranteed payoff for the Right Party, there exist strategies to support it. Therefore, the $L^*(q_L)$ and $R^*(q_R)$ make up an equilibrium. This also implies that every equilibrium has the value generated by $L^*(q_L)$ and $R^*(q_R)$.

For uniqueness, recall that by Bayes' Theorem the integral of the CDF is related to the prior by,

$$\int_0^1 dF_L(q_L) = 1 - q_L^0. \tag{5.8}$$

Therefore, any feasible distribution must have the same integral. For a contradiction to uniqueness, suppose an alternative strategy for the Left Party $L'(q_L) \neq L^*(q_L)$ was part of an equilibrium. First suppose $L'(q_L) < L^*(q_L)$ at a $q_L' \leq q_L^0$. By Equation 5.8, there must be an alternative point, $q_L'' \geq q_L^0$, where $L'(q_L) > L^*(q_L)$. An example is given by the solid black line on the left half of Figure 5.4. The equilibrium strategy $L^*(q_L)$ is given by the dashed line for reference.

Given $L'(q_L)$, the probability that the Right Party wins for an induced belief q_R is $1 - v^*(q_R; L'(q_L))$. That function is the dotted line in the right half of Figure 5.4. The Right Party's best-response achieves the concavification at q_R^0 (Kamenica and Gentzkow 2011). The concavification is given by the red line. The value of the concavification at q_R^0 is strictly greater than the concavification against $L^*(q_L)$. That is, the Right Party could use concavification to win with a higher probability than equilibrium using strategies in Equation 5.5. Therefore, if $L'(q_L)$ was part of an equilibrium it would have a different value than the original equilibrium. This is a contraction for any $L'(q_L) < L^*(q_L)$, when $q_L' \leq q_L^0$. The same argument holds for $L'(q_L) > L^*(q_L)$ at a $q_L' \leq q_L^0$. Therefore, it must be that Equation 5.5 is the unique equilibrium for $q_L^0 \leq 1/2$.

For the cases where $q_L^0 > 1/2$, the uniform distribution is still used in equilibrium, but it is slightly changed. It has a flat part for $q_L \in (2(1 - q_L^0), 1)$. The flat part, while changing the shape of the distribution, does not change the algebra.

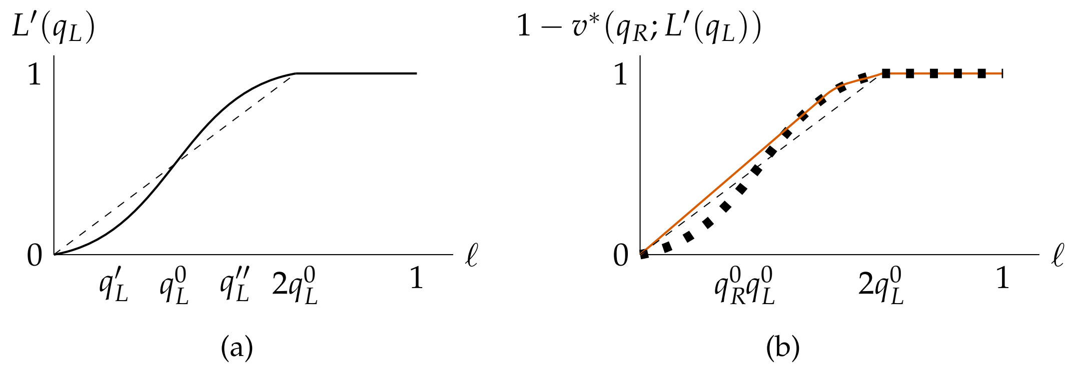

Figure 5.4: Equilibrium Uniqueness

Suppose the Right Party using the strategy in Equation 5.6:

$$\begin{aligned}
P(q_L \geq q_R) &= \int_0^1 R^*(x) dL(x) \\
&= \int_0^1 \min\left\{ \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right) \frac{1}{2q_L^0} x, \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right)\left(\frac{1}{q_L^0} - 1\right) \right\} dL(x) \\
&\leq \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right) \frac{1}{2q_L^0} \int_0^1 x dL(x) \\
&= \left(1 - \frac{q_R^0}{q_L^0}\right) + \left(\frac{q_R^0}{q_L^0}\right) \frac{1}{2q_L^0} q_L^0 \quad \text{(Bayes-plausible)} \\
&= 1 - \frac{q_R^0}{2q_L^0}.
\end{aligned} \tag{5.9}$$

Again, the Right Party has chosen a strategy to minimize the probability that the Left Party wins.

For the uniqueness of the second case, again recall that for zero-sum games all equilibria have the same value. Also, note there exists no concave CDF that is Bayes-plausible when $q_L^0 > 1/2$, because Equation 5.8 must hold. However, with positive probability placed at $q_L = 1$, where the party cannot lose by ϵ, this is the only CDF that is Bayes-plausible and that the concave closure at q_R^0 that equals the value generate by Equation 5.6. ■

To connect the geometric Bayesian persuasion argument from the example with the optimal equilibrium strategies, consider the case where $q_L^0 = q_R^0 > 1/2$. Geometrically, the Left Party will concavify $v^*(q_L; R^*(q_R))$. The Right Party's strategy minimizes the concavification, given by the red line on the right at q_R^0 in Figure 5.5.

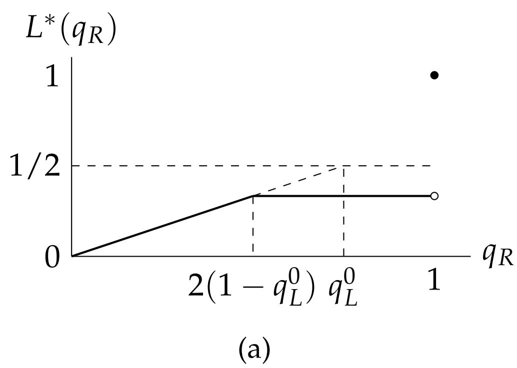

Figure 5.5: Concave Closure

In fact, the Left Party's strategy minimizes the concavification for any prior probability of the Right Party. This leads to the following corollary.

Corollary 13.1. *Let $q_L^0 \geq q_R^0$. Then the Left Party's equilibrium strategy does not depend on the prior probability that the Right Party's candidate is good for the voter.*

Whichever party has a higher prior probability does not need to worry about the strength of the other party. The stronger party is in control of the election. The weaker party induces a posterior belief of zero in an attempt to compete in the other elections.

Compared to complete honesty, the Left Party (the ex-ante favored party) wins more often. To see, compare the probability of winning under persuasion vs. honesty

$$\begin{aligned} 1-\frac{1}{2}\frac{r}{l} &\geq \frac{1}{2}lr+\frac{1}{2}\frac{1-l}{1-r}+l(1-r)=\frac{1}{2}(1-r+l) \\ 2l-r &\geq l+l^2-rl \\ l-r &\geq l^2-rl \\ \frac{l}{r} &\geq (l-r)l \\ 1 &\geq l. \end{aligned}$$

This always holds.

How often is a good candidate elected under the two systems? In all cases, both parties use a partially uniform strategy. For the uniform part of the distribution, each political party designs its strategy to give away as little information as possible. Since any information that the Right Party could give away with its strategy will be used against it by the Left Party, it wants to avoid this. This also means giving as little information to the voter as possible. The uniform distribution does this exactly. One way to see this is to consider the entropy of the distribution chosen. A higher entropy means higher uncertainty.[16] For random variables on [0,1], the uniform distribution is the maximum entropy distribution. It has the highest uncertainty. For the voting example, the uniform distribution

16. Cover and Thomas (2006, 6) state that "entropy is the uncertainty of a single random variable."

of beliefs means that the other party knows as little as possible about the other party's induced beliefs. Each party generates the most uncertainty it can. Caught in the middle is the voter.[17]

The only part of the strategy that has positive mass occurs at zero or one. The mass at one is used because there is a cap on the posterior beliefs, which cannot be above one. For the mass at zero, consider the choice of the Right Party. The Right Party, which has a lower prior probability of having a good candidate, cannot compete in every election. They throw some elections by inducing a belief a zero with probability $1 - q_R^0/q_L^0$. Since the Left Party never induces a belief of zero, the Right Party loses those elections with certainty. However, by strategically throwing some elections, the Right Party is able to compete in the other elections, which occur with probability q_R^0/q_L^0. In those remaining elections, the Right Party follows a uniform like the Left Party.

It is worth noting who benefits from persuasion. Of course, that depends on the answer to *compared to what*? If the comparison is between persuasion and uninformative strategies, the unfavored party benefits. If the comparison is between persuasion and fully informative strategies, the favored party benefits.

5.3 Beyond Bad and Good Types

Much of the simplicity of the strategies comes from there being two types of candidates for each party: bad and good. Instead, suppose the voter has preferences

17. This fits with models like Glazer (1990), where the candidate is purposefully ambiguous.

defined over the candidate quality, $q \in [0, 1]$, given by $u(q) = -|1 - q|$. A candidate of type $q = 0$ is the worst and $q = 1$ is the best. Again, suppose everyone has expected utility. The parties start with a prior probability distribution over $[0, 1]$. Let $L(q)$ and $R(q)$ be the prior CDF for the Left Party and Right Party. It is easier to drop the subscript for the remaining discussion.

When there are more than two types the nature of the signal changes. The parties no longer want to send a signal about the probability the candidate is good, but the expected quality of the candidate. Denote the prior expectations by $\mathbb{E}_L[q]$ and $\mathbb{E}_R[q]$. The parties choose distributions over expected quality for their strategies. However, as elaborated in Gentzkow and Kamenica (2016a), not every distribution of expected quality that equals the prior expected quality can be generated.[18] However, in some cases, there exists a direct analog strategy from the two-type problem is feasible for the general case. By analog strategy, I mean q_L^0 is replaced by $\mathbb{E}_L[q]$, and q_R^0 is replaced by $\mathbb{E}_R[q]$ in Equations 1 and 2. Define the strategies to be $L^*(q)$ and $R^*(q)$. The next two corollaries consider situations when we can directly apply the result of Theorem 13.

The first of these corollaries gives knife-edge cases where the strategy follows exactly from the two-type case.

Corollary 13.2. *Suppose the prior CDFs for both parties are given by $L^*(q)$ and $R^*(q)$ from Theorem 13. Then there exists an equilibrium where both parties use fully-informative strategies.*

18. A simple example should highlight the problem. Suppose there are three types of candidates denoted by their quality, $q \in \{0, 1/2, 1\}$, and they each occur with equal probability. The prior expected position is $1/2$. However, it is impossible to generate a distribution of posterior beliefs with probability $1/2$ both at 0 and 1.

The proof follows immediately from Theorem 13 and the observation that a fully-informative strategy is always feasible. Note that Corollary 13.2 does not allow fully-revealing strategies in an equilibrium with *finite* types. In that case, Lemma 4 still applies and establishes that there is no best-response to a signal with a finite number of realizations.

While Corollary 13.2 shows a knife-edge case where fully-informative signals are used, the equilibrium strategies from Theorem 13 apply to more general distributions, although they are not fully-informative. To say more about the cases where equilibrium is similar in structure to the two cases, I adopt notation from Gentzkow and Kamenica (2016a). Define $c_X(q)$ to be the integral of a CDF X, i.e.

$$c_X(q) = \int_0^q X(t)dt.$$

The X could be the Left Party's or Right Party's prior or posterior distribution. In addition, let $c_{\bar{\pi}}$ be the integral of a fully-informative signal and $c_{\underline{\pi}}$ be the integral of an uninformative signal. The following result from Gentzkow and Kamenica (2016a) explains which posterior distributions of the expected quality are feasible.

Result 1 (Gentzkow and Kamenica 2016a)**.** Given any convex function $c : [0,1] \to \mathbb{R}$ such that $c_{\bar{\pi}} \geq c \geq c_{\underline{\pi}}$, there exists a signal that induces it.

Given the set of feasible distributions of expected quality, one can simply check whether the proposed distribution strategy is feasible using Result 1. The relevant inequality to check is whether $c_{\bar{\pi}} \geq c$ since the other inequality is satisfied for any distribution of expected quality where the expectation equals the prior. Corollary 13.3 incorporates this new constraint to clarify when the equilibrium strategies of Theorem 13 extend to the more general case.

Corollary 13.3. *Let $\mathbb{E}_L[q] \geq \mathbb{E}_R[q]$. Then for every prior distributions $L(q)$ and $R(q)$ such that the integrals of the fully-informative signal are both greater pairwise than the corresponding integrals of $L^*(q)$ and $R^*(q)$, i.e. $c_L[q] \geq c_{L^*}[q]$ and $c_R[q] \geq c_{R^*}[q]$ for every $q \in [0, 1]$, then $L^*(q)$ and $R^*(q)$ are the equilibrium strategies..*

Informally, the prior probability must put enough weight near zero. That ensures that the integral of the equilibrium strategy considered is less than the fully-informative signal and the constraint $c_{\bar{\pi}} \geq c$ is not violated. If that is true, the proposed equilibrium strategies from Theorem 13 are feasible. This fact shows how the uniform strategy is a more common result than just in the two-type case. Therefore, even though Result 1 complicates the equilibrium for some distributions, there is a class of distributions where Theorem 13 directly extends. However, in more general cases, the equilibrium strategies will not directly correspond to the two-type equilibrium. The problematic distribution in footnote 18 comes because the proposed strategy puts too much weight around $q = 0$ and $q = 1$ relative to the prior distribution and is therefore infeasible.

While theoretically clean and tractable, generating a uniform distribution of posteriors is not clear for real campaigns. Such a strategy requires a continuum of outcomes to the campaign. However, as Cover (1974) shows, any distribution of posteriors such that $\mathrm{E}[r] = 1/2$ can be generated by a sequence of fair coin tosses. To take the voting example, a uniform distribution of beliefs can be generated by a sequence of days of information. On one day, the party decides to release a big piece of news. There is a 50% chance that the voter interprets that news as coming from a good candidate. On another day, the party releases "half" as much news.

On a third, half of that. A campaign strategy of this type will generate a uniform distribution of posterior beliefs on [0,1], which is what each party wants.[19]. However, this requires an infinite sequence of news-cycles. That is never actually possible, but maybe 597 days of campaigning for president is a close approximation. Maybe not. However, as said before, the uniform strategy is common in zero-sum games and widely used. The next section elaborates on some of these results and other related literature from voting and Bayesian persuasion.

5.4 Colonel Blotto and Persuasion Literature

The current chapter is closest to the literature of elections that are model as zero-sum Colonel Blotto games (Gross and Wagner 1950; Roberson 2006). I have already mentioned that the closest voting papers are Myerson (1993) and Sahuguet and Persico (2006). In Myerson (1993), candidates compete by offering benefits to voters; the candidates are buying votes. As in this paper, a strategy is a distribution; the probability density is how many voters receive a certain level of benefits. The candidates have budget constraints which mimic the Bayes-plausible restriction in this chapter. The equilibrium strategy is a uniform distribution of benefits.[20] As Roberson (2006) shows, the uniform distribution is a general result for Colonel Blotto games, where the players want to win part of the game, whether

19. The formal argument comes from Bell and Cover (1980, p. 163) Suppose a gambler starts with 1/2 unit of capital. He divides the capital into piles of $\frac{1}{2}\left(\frac{1}{2}\right)^i$, for $i = 1, 2, ...$ There is a sequence of fair coin flips and the gambler bets the ith pile on the ith coin flip. The expected outcome of the sequence of bets is distributed uniformly on [0,1].

20. Bell and Cover (1980) find the same equilibrium strategy for competing investors who have $1 to place on fair bets.

they be battlefields or voters, by as little as possible. In Myerson, each candidate wants to offer ϵ more than the other candidate. Sahuguet and Persico extend Myerson to include asymmetries between the parties. The uniform strategies in Equation 5.5 in Theorem 13 are the same strategies as in Theorem 1 in Sahuguet and Persico (2006, 104). Taking those papers and this chapter together, there are two ways of "buying" votes: actual payments or beliefs. This chapter focuses on beliefs and shows how political parties can go about "outbidding" the other party with beliefs.

The way I chose to model information manipulation comes from the Bayesian persuasion literature, started by Kamenica and Gentzkow (2011).[21] A few papers on Bayesian persuasion and information manipulation are especially worth mentioning and connecting to this chapter. Boleslavsky and Cotton (2015, 2016) consider environments very similar to the present chapter. Boleslavsky and Cotton (2015) study the problem of grading standards and Boleslavsky and Cotton (2016) study the impact of limited capacity (like being allowed to choose one politician). Both papers show that a uniform distribution is the optimal strategy. Au and Kawai (2020) study a general form of competition in persuasion, of which the baseline model in this chapter is a special case. In addition, the discussion in Section 5.3 with more than two types, goes beyond their focus, which is mostly about the two type case.

Two other papers study persuasion in a voting context and are closely related:

21. Some examples include Gentzkow and Kamenica (2014), Ely, Frankel, and Kamenica (2015), and Ely (2017) for single senders. For multiple senders, like this chapter, see Boleslavsky and Cotton (2015, 2016), Gentzkow and Kamenica (2016b, 2017), Koessler, Laclau, and Tomala (2017), Li and Norman (2018), and Au and Kawai (2020).

Aköz and Arbatli (2016) and Alonso and Câmara (2016). In Aköz and Arbatli (2016), the voter does not know his preferred position, although the candidates do. This turns the problem into a situation of information manipulation on a single variable, the voter's preferences. That differs from this chapter where each political party has private information and tries to persuade voters about that information. This makes more sense when seeing the information revelation as designing a campaign schedule, as I do in this chapter. The Republicans choose their schedule. The Democrats choose their schedule. Real campaigns have both elements, but I focus on the schedule design.

Alonso and Câmara (2016) consider a single politician who is trying to get voters to approve certain proposals. It is one sender and a group of receivers. That paper compares different voting rules (say majority vs. unanimity) and compares the outcomes. While that paper and the present both talk of politicians, the papers differ both in the number of senders and the number of receivers, making the problems distinct, though related. An interesting extension would be to analyze different voting rules with competing political parties. There are also many papers that have studied the manipulation of beliefs by politicians outside of elections, such as in Angeletos, Hellwig, and Pavan (2006) and Edmond (2013).

However, most of the voting literature is not touched in this chapter. To isolate the role of persuasion, the model simplifies the decisions by parties and voters. Since the political parties can only use persuasion, parties' strategies are extremely limited. There is no strategic picking of who runs, such as in Besley and Coate (1997), or what policy to run on. Therefore, there is no median voter theorem (Downs 1957). However, this lack of convergence to the median voter is not because of any policy preferences by the parties (Wittman 1977, 1983) or because of signaling (Bernhardt, Duggan, and Squintani 2009; Kartik and Mcafee 2007). To simplify the voter's decision, more extensions were left out. Since there is only one voter, she is basically non-strategic. She votes mechanically, given her beliefs. This avoids complications from insincere voting, such as in Austen-Smith and Banks (1996), or strategic abstention, such as in Feddersen and Pesendorfer (1996).

Related to the role of information in an election, this model abstracts from issues related to costly information acquisition (Persico 2004; Gerardi and Yariv 2008; Gershkov and Szentes 2009; Tyson 2016) and the paradoxes that are involved with that (Martinelli 2006). All these simplifications provide possible voting extensions to the growing Bayesian persuasion literature.

5.5 Conclusion

This chapter studies a model where political parties compete solely through information. Uncertainty about the quality of the candidates opens up the possibility for persuasion. Both parties engage in persuasion in order to win or equivalently prevent the other party from winning. As is common in the Bayesian persuasion literature, a sender benefits if he can concavify the value of the receiver's beliefs. In order to not let the other party concavify, each party chooses a uniform distribution of posteriors, which is already concave. Each party sets up its campaign schedule so that the other party cannot take easy advantage of it. This means each party chooses a campaign schedule that induces a uniform distribution of beliefs.

The uniform strategy is a shared feature of many voting games. In a persuasion model, the uniform strategy means that the political parties introduce the maximum amount of uncertainty into the election.

Bibliography

Acemoglu, Daron. 1996. "A Microfoundation for Social Increasing Returns in Human Capital Accumulation." *The Quarterly Journal of Economics* 111 (3): 779–804. ISSN: 00335533, 15314650. doi:10.2307/2946672. (Cited on page 27).

Aguirre, Iñaki. 2008. "Output and misallocation effects in monopolistic third-degree price discrimination." *Economics Bulletin* 4 (11): 1–11. (Cited on page 79).

Aguirre, Iñaki, and Simon G. Cowan. 2015. "Monopoly price discrimination with constant elasticity demand." *Economic Theory Bulletin* 3, no. 2 (October): 329–340. doi:10.1007/s40505-014-0063-3. http://link.springer.com/10.1007/s40505-014-0063-3. (Cited on page 79).

Aköz, Kemal Kivanc, and Cemal Eren Arbatli. 2016. "Information Manipulation in Election Campaigns." *Economics and Politics* 28 (2): 181–215. ISSN: 14680343. doi:10.1111/ecpo.12076. (Cited on page 132).

Albrecht, Brian C. 2016. "Entrepreneurship as Coordination." https://briancalbrecht.github.io/albrecht-entrepreneurship-coordination.pdf. (Cited on page 10).

———. 2017. "Political Persuasion." https://briancalbrecht.github.io/albrecht-political-persuasion.pdf. (Cited on page 91).

———. 2018. "Some Price Theory of Buyer-Optimal Learning." (Cited on page 77).

Albrecht, Brian C., and Brian Kogelmann. 2018. "Coase, the Austrians, and Models as Foils." https://briancalbrecht.github.io/models-as-foils.pdf. (Cited on page 11).

Alonso, Ricardo, and Odilon Câmara. 2016. "Persuading voters." *The American Economic Review* 106 (11): 3590–3605. ISSN: 00028282. doi:10.1257/aer.20140737. (Cited on page 132).

Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan. 2006. "Signaling in a Global Game: Coordination and Policy Traps." *Journal of Political Economy* 114 (3): 452–484. doi:10.1086/504901. (Cited on page 132).

Armstrong, Mark, and John Vickers. 2019a. "Discriminating against Captive Customers." *AER: Insights* 1 (3): 257–272. doi:10.1257/aeri.20180581. https://doi.org/10.1257/aeri.20180581. (Cited on pages 81, 84, 94).

———. 2019b. "Patterns of Competitive Interaction." Working paper. (Cited on page 57).

———. 2020. "Patterns of Price Competition and the Structure of Consumer Choice." (Cited on pages 81, 94).

Au, Pak Hung, and Keiichi Kawai. 2020. "Competitive information disclosure by multiple senders." *Games and Economic Behavior* 119 (January): 56–78. ISSN: 10902473. doi:10.1016/j.geb.2019.10.002. (Cited on pages 80, 131).

Aumann, Robert J. 1964. "Markets with a Continuum of Traders." *Econometrica* 32 (1-2): 39–50. (Cited on page 40).

Aumann, Robert J., and Michael Maschler. 1995. *Repeated Games with Incomplete Information.* Cambridge, MA: The MIT Press. (Cited on pages 105, 107).

Austen-Smith, David, and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90 (1): 34–45. doi:10.2307/2082796. http://www.jstor.org/stable/2082796. (Cited on page 133).

Bell, Robert M., and Thomas M. Cover. 1980. "Competitive Optimality of Logarithmic Investment." *Mathematics of Operations Research* 5 (2): 161–166. ISSN: 0364-765X. doi:10.1287/moor.5.2.161. (Cited on page 130).

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The Limits of Price Discrimination." *The American Economic Review* 105 (3): 921–957. ISSN: 00028282. doi:10.1257/aer.20130848. http://dx.doi.org/10.1257/aer.20130848. (Cited on pages 78, 80, 81).

Bergemann, Dirk, and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–1308. ISSN: 0012-9682. doi:10.3982/ECTA11105. http://doi.wiley.com/10.3982/ECTA11105. (Cited on page 80).

———. 2016. "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics* 11 (2): 487–522. ISSN: 19336837. doi:10.3982/TE1808. http://doi.wiley.com/10.3982/TE1808. (Cited on page 84).

———. 2019. "Information design: A unified perspective." *Journal of Economic Literature* 57, no. 1 (March): 44–95. ISSN: 00220515. doi:10.1257/jel.20181489. (Cited on page 97).

Bergemann, Dirk, and Karl H. Schlag. 2008. "Pricing without Priors." *Journal of the European Economic Association* 63 (2): 560–569. http://www.jstor.org/stable/40282665. (Cited on page 79).

Bernhardt, Dan, John Duggan, and Francesco Squintani. 2009. "The Case for Responsible Parties." *American Political Science Review* 103 (4): 570–587. doi:10.1017/S0003055409990232. http://www.journals.cambridge.org/abstract_S0003055409990232. (Cited on page 133).

Besley, Timothy, and Stephen Coate. 1997. "An Economic Model of Representative Democracy." *The Quarterly Journal of Economics* 112 (1): 85–114. doi:10.1162/003355397555136. http://www.jstor.org/stable/2951277. (Cited on page 132).

Blackwell, David. 1951. "Comparison of Experiments." In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability,* edited by Jerzy Neyman, 93–102. Berkeley: University of California Press. (Cited on page 78).

———. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* 24 (2): 265–272. doi:10.2307/2236332. http://www.jstor.org/stable/2236332. (Cited on page 78).

Boleslavsky, Raphael, and Christopher Cotton. 2015. "Grading Standards and Education Quality." *American Economic Journal: Microeconomics* 7, no. 2 (May): 248–279. ISSN: 1945-7669. doi:10.1257/mic.20130080. http://pubs.aeaweb.org/doi/10.1257/mic.20130080. (Cited on pages 80, 131).

———. 2016. "Limited capacity in project selection: competition through evidence production." *Economic Theory.* doi:10.1007/s00199-016-1021-0. (Cited on pages 80, 131).

Bolton, Patrick, and Mathias Dewatripont. 2005. *Contract Theory.* Cambridge, United States: MIT Press. (Cited on page 70).

Borenstein, Severin. 1985. "Price discrimination in free-entry markets." *Rand Journal of Economics* 16 (3): 380–397. (Cited on page 79).

Brandenburger, Adam M., and Harborne W. Stuart. 2007. "Biform Games." *Management Science* 53 (4): 537–549. doi:10.1287/mnsc.1060.0591. http://pubsonline.informs.org549.https//doi.org/10.1287/mnsc.1060.0591http://www.informs.org. (Cited on page 12).

Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica* 51 (4): 955–969. http://www.jstor.org/stable/1912045. (Cited on page 80).

Chander, Parkash. 1983. "On the Informational Size of Message Spaces for Efficient Resource Allocation Processes." *Econometrica* 51:919–938. (Cited on page 36).

Chandra, Ambarish, and Mara Lederman. 2018. "Revisiting the Relationship between Competition and Price Discrimination." *American Economic Journal: Microeconomics* 10 (2): 190–224. doi:10.1257/mic.20160252. https://doi.org/10.1257/mic.20160252. (Cited on page 79).

Cole, Harold L., George J. Mailath, and Andrew Postlewaite. 2001a. "Efficient Non-Contractible Investments in Finite Economies." *Advances in Theoretical Economics* 1 (1): 1–32. (Cited on page 12).

———. 2001b. "Efficient non-contractible investments in large economies." *Journal of Economic Theory* 101 (2): 333–373. ISSN: 00220531. doi:10.1006/jeth.2001.2797. (Cited on pages 12, 16, 21).

Condorelli, Daniele, and Balazs Szentes. 2019. "Information Design in the Hold-up Problem." *Journal of Political Economy.* http://personal.lse.ac.uk/szentes/docs/holdup.pdf. (Cited on page 79).

Cooper, Russell, and Andrew John. 1988. "Coordinating Coordination Failures in Keynesian Models." *The Quarterly Journal of Economics* 103 (3): 441–463. doi:10.2307/1885539. https://academic.oup.com/qje/article-lookup/doi/10.2307/1885539. (Cited on page 11).

Cover, Thomas M. 1974. *Universal Gambling Schemes and the Complexity Measures of Kolmogorov and Chaitin.* Technical report. Stanford, CA: Stanford University. https://statistics.stanford.edu/sites/default/files/COV%20NSF%2012.pdf. (Cited on page 129).

Cover, Thomas M., and Joy A Thomas. 2006. *Elements of Information Theory.* 2nd Edition. Hoboken, New Jersey: Wiley-Interscience. (Cited on page 125).

Debreu, Gegard, and Herbert Scarf. 1963. "A Limit Theorem on the Core of an Economy." *International Economic Review* 4 (3): 235–246. (Cited on page 40).

Dekel, Eddie, Barton L. Lipman, and Aldo Rustichini. 1998. "Recent developments in modeling unforeseen contingencies." *European Economic Review* 42, nos. 3-5 (May): 523–542. ISSN: 00142921. doi:10.1016/S0014-2921(97)00114-1. http://www.sciencedirect.com/science/article/pii/S0014292197001141. (Cited on page 41).

Dizdar, Deniz. 2018. "Two-Sided Investment and Matching with Multidimensional Cost Types and Attributes." *American Economic Journal: Microeconomics* 10, no. 3 (August): 86–123. doi:10.1257/mic.20150147. http://www.aeaweb.org/articles?id=10.1257/mic.20150147. (Cited on page 26).

Downs, Anthony. 1957. *An Economic Theory of Democracy.* 1–36. 1957. New York: Harper / Row. ISBN: 9783825229252. (Cited on page 132).

Dubey, Pradeep, and John D. Geanakoplos. 2002. "Competitive Pooling: Rothschild-Stiglitz Reconsidered." *The Quarterly Journal of Economics* 117, no. 4 (November): 1529–1570. ISSN: 0033-5533. doi:10.1162/003355302320935098. https://dx.doi.org/10.1162/003355302320935098. (Cited on pages 14, 33).

Dubey, Pradeep, John D. Geanakoplos, and Martin Shubik. 2005. "Default and Punishment in General Equilibrium." *Econometrica* 73, no. 1 (January): 1–37. doi:10.1111/j.1468-0262.2005.00563.x. https://doi.org/10.1111/j.1468-0262.2005.00563.x. (Cited on pages 14, 33).

Duffie, Darrel, Nicolae Gârleanu, and Lasse H. Pedersen. 2005. "Over the Counter Markets." *Econometrica* 73 (6): 1815–1847. (Cited on page 38).

Edgeworth, Francis Y. 1881. *Mathematical Psychics.* London, United Kingdom: C. Kegan Paul & co. (Cited on page 40).

Edmond, Chris. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80 (4): 1422–1458. ISSN: 00346527. doi:10.1093/restud/rdt020. (Cited on page 132).

Elliott, Matthew, and Andrea Galeotti. 2019. "Market segmentation through information." (Cited on page 81).

Ely, Jeffrey C. 2017. "Beeps." *The American Economic Review* 107, no. 1 (January): 31–53. ISSN: 00028282. doi:10.1257/aer.20150218. http://pubs.aeaweb.org/doi/10.1257/aer.20150218. (Cited on page 131).

Ely, Jeffrey C., Alexander Frankel, and Emir Kamenica. 2015. "Suspense and Surprise." *Journal of Political Economy* 123 (1): 215–260. ISSN: 0048-5772. doi:10.1086/677349. (Cited on page 131).

Feddersen, Timothy J, and Wolfgang Pesendorfer. 1996. "The Swing Voter's Curse." *The American Economic Review* 86 (3): 408–424. http://www.jstor.org/stable/2118204. (Cited on page 133).

Felli, Leonardo, and Kevin Roberts. 2016. "Does Competition Solve the Hold-up Problem?" *Economica* 83 (329): 172–200. doi:10.1111/ecca.12170. (Cited on pages 12, 23).

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?" *American Economic Review* 98 (1): 394–425. (Cited on page 63).

———. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 84 (329): 91–129. (Cited on pages 42, 63).

Gale, Douglas. 1986a. "Bargaining and Competition Part I: Characterization." *Econometrica* 54 (4): 785–806. (Cited on page 40).

———. 1986b. "Bargaining and Competition Part II: Existence." *Econometrica* 54 (4): 807–818. (Cited on page 40).

———. 1987. "Limit Theorems for Markets with Sequential Bargaining." *Journal of Economic Theory* 43 (1): 20–54. (Cited on page 40).

———. 1992. "A Walrasian Theory of Markets with Adverse Selection." *The Review of Economic Studies* 59 (2): 229–255. ISSN: 00346527, 1467937X. doi:10.2307/2297953. (Cited on pages 13, 14, 30).

———. 1996. "Equilibria and Pareto optima of markets with adverse selection." *Economic Theory* 7, no. 2 (June): 207–235. ISSN: 0938-2259. doi:10.1007/BF01213903. http://link.springer.com/10.1007/BF01213903. (Cited on page 14).

———. 2000. *Strategic Foundations of General Equilibrium.* Cambridge, United Kingdom: Cambridge University Press. (Cited on pages 37, 40).

Gentzkow, Matthew, and Emir Kamenica. 2014. "Costly Persuasion." *American Economic Review: Papers and Proceedings* 104 (5): 457–462. doi:10.1257/aer.104.5.457. (Cited on page 131).

———. 2016a. "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review: Papers and Proceedings* 106 (5): 597–601. ISSN: 00028282. doi:10.1257/aer.p20161049. (Cited on pages 115, 127, 128).

———. 2016b. "Competition in Persuasion." *Review of Economic Studies* 84, no. 1 (January): 300–322. ISSN: 1467937X. doi:10.1093/restud/rdw052. (Cited on pages 80, 131).

———. 2017. "Bayesian Persuasion with Multiple Senders and Rich Signal Spaces." *Games and Economic Behavior* 104 (March): 411–429. doi:10.1016/j.geb.2017.05.004. (Cited on page 131).

Gerardi, Dino, and Leeat Yariv. 2008. "Information acquisition in committees." *Games and Economic Behavior* 62 (2): 436–459. ISSN: 08998256. doi:10.1016/j.geb.2007.06.007. (Cited on page 133).

Gershkov, Alex, and Balazs Szentes. 2009. "Optimal voting schemes with costly information acquisition." *Journal of Economic Theory* 144 (1): 36–68. doi:10.1016/j.jet.2008.02.004. (Cited on page 133).

Glazer, Amihai. 1990. "The Strategy of Candidate Ambiguity." *American Political Science Review* 84 (1): 237–241. http://www.jstor.org/stable/1963640. (Cited on page 126).

Gretsky, Neil E., Joseph M. Ostroy, and William R. Zame. 1999. "Perfect Competition in the Continuous Assignment Model." *Journal of Economic Theory* 88:60–118. (Cited on pages 11, 21).

Gross, Oliver, and Robert Wagner. 1950. "A Continuous Colonel Blotto Game." *RAND Corporation RM-408.* http://www.rand.org/pubs/research_memoranda/RM408/. (Cited on page 130).

Guthmann, Rafael R. 2019a. "Price Dispersion as a Result of Unawareness." https://ssrn.com/abstract=3290863. (Cited on page 81).

———. 2019b. "Price Dispersion in Dynamic Competition." https://ssrn.com/abstract=3290853. (Cited on page 81).

Guthmann, Rafael R., and Brian C. Albrecht. 2020. "On the Informational Efficiency of Decentralized Price Formation." (Cited on page 81).

Hammond, Peter J. 1979. "Straightforward Individual Incentive Compatibility in Large Economies." *Review of Economic Studies* 46 (2): 263–282. (Cited on page 36).

Hart, Sergiu. 2008. "Discrete Colonel Blotto and General Lotto games." *International Journal of Game Theory* 36 (3): 441–460. ISSN: 00207276. doi:10.1007/s00182-007-0099-9. (Cited on page 119).

Hayek, Friedrich A. 1937. "Economics and Knowledge." *Economica* 4 (13): 33–54. http://www.jstor.org/stable/2548786. (Cited on pages 25, 35).

———. 1945. "The Use of Knowledge in Society." *The American Economic Review* 35 (4): 519–530. https://www.jstor.org/stable/1809376. (Cited on page 35).

Hirschman, Albert O. 1970. *Exit, Voice, and Loyalty.* Cambridge, MA: Harvard University Press. ISBN: 0-674-27660-4. (Cited on page 104).

Holmes, Thomas J. 1989. "The Effects of Third-Degree Price Discrimination in Oligopoly." *The American Economic Review* 79:244–250. doi:10.2307/1804785. https://www.jstor.org/stable/1804785. (Cited on page 79).

Hurwicz, Leonid. 1977a. "On informationally decentralized systems." In *Studies in Resource Allocation Processes,* edited by K. J. Arrow and L. Hurwicz, 425–459. Cambridge, New York: Cambridge University Press. (Cited on pages 36, 66).

———. 1977b. "On the dimensional requirements of informationally decentralized Pareto-satisfactory processes." In *Studies in Resource Allocation Processes,* edited by K. J. Arrow and L. Hurwicz, 413–424. Cambridge, New York: Cambridge University Press. (Cited on pages 36, 49, 69).

———. 1977c. "Optimality and informational efficiency in resource allocation processes." In *Studies in Resource Allocation Processes,* edited by K. J. Arrow and L. Hurwicz, 393–423. Cambridge, New York: Cambridge University Press. (Cited on page 36).

Jordan, James S. 1982. "The Competitive Allocation Process is Informationally Efficient Uniquely." *Journal of Economic Theory* 28 (1): 1–18. (Cited on pages 36, 49, 69).

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *The American Economic Review* 101 (6): 2590–2615. doi:10.1257/aer.101.6.2590. (Cited on pages 3, 94, 105–107, 110, 115, 117, 119, 122, 131).

Kaplan, Greg, Guido Menzio, Lenna Rudanko, and Nicholas Trachter. 2019. "Relative Price Dispersion: Evidence and Theory." *American Economic Journal: Microeconomics* 11 (3): 68–124. (Cited on page 42).

Kartik, Navin, and R. Preston Mcafee. 2007. "Signaling Character in Electoral Competition." *The American Economic Review* 97 (3): 852–870. doi:10.1257/aer.97.3.852. http://www.jstor.org/stable/30035023. (Cited on page 133).

Kirzner, Israel M. 1973. *Competition and Entrepreneurship.* Chicago: University of Chicago Press. (Cited on page 38).

Klein, Daniel, and Aaron Orsborn. 2009. "Concatenate Coordination and Mutual Coordination." *Journal of Economic Behavior & Organization* 72 (1): 176–187. ISSN: 01672681. doi:http://dx.doi.org/10.1016/j.jebo.2009.05.003. (Cited on page 10).

Koessler, Frederic, Marie Laclau, and Tristan Tomala. 2017. "Competitive Information Design: The Static Case." (Cited on page 131).

Kohlberg, Elon, and Jean-François Mertens. 1986. "On the Strategic Stability of Equilibria." *Econometrica* 54 (5): 1003–1037. ISSN: 00129682, 14680262. doi:10.2307/1912320. http://www.jstor.org/stable/1912320. (Cited on page 9).

Lauermann, Stephan. 2013. "Dynamic Matching and Bargaining Games: A General Approach." *American Economic Review* 103 (2): 663–689. (Cited on page 40).

Lauermann, Stephan, Wolfram Merzyn, and Gábor Virág. 2018. "Learning and Price Discovery in a Search Market." *Review of Economic Studies* 85 (2): 1159–1192. (Cited on page 40).

Li, Fei, and Peter Norman. 2018. "On Bayesian persuasion with multiple senders." *Economics Letters* 170 (September): 66–70. ISSN: 01651765. doi:10.1016/j.econlet.2018.05.023. (Cited on pages 80, 131).

Liu, Qingmin. 2018. "Rational Expectations, Stable Beliefs, and Stable Matching." (Cited on page 14).

Makowski, Louis. 2004. "Pre-Contractual Investment without the Fear of Holdups: The Perfect Competition Connection." (Cited on pages 13, 16, 21).

Makowski, Louis, and Joseph M. Ostroy. 1995. "Appropriation and Efficiency: A Revision of the First Theorem of Welfare Economics." *The American Economic Review* 85 (4): 808–827. http://www.jstor.org/stable/2118233. (Cited on pages 12, 16, 33).

———. 1998. "Arbitrage and the Flattening Effect of Large Numbers." *Journal of Economic Theory* 78:1–31. doi:10.1006/jeth.1997.2350. (Cited on page 38).

———. 2013. "From revealed preference to preference revelation." *Journal of Mathematical Economics* 49, no. 1 (January): 71–81. doi:10.1016/J.JMATECO.2012.10.002. (Cited on page 10).

Mankiw, N. Gregory, and Ricardo Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328. (Cited on page 56).

Martinelli, CÃ©sar. 2006. "Would rational voters acquire costly information?" *Journal of Economic Theory* 129 (1): 225–251. doi:10.1016/j.jet.2005.02.005. (Cited on page 133).

Mas-Colell, Andreu, Michael D. Winston, and Jerry R. Green. 1995. *Microeconomic Theory.* 977. New York: Oxford University Press. ISBN: 0-19-510268-1. (Cited on page 76).

Maschler, Michael, Eilon Solan, and Shmuel Zamir. 2013. *Game Theory.* First. New York: Cambridge University Press. ISBN: 9781139049344. (Cited on page 14).

McAfee, R. Preston. 1994. "Endogenous Availability, Cartels, and Merger in an Equilibrium Price Dispersion." *Journal of Economic Theory* 62:24–47. (Cited on pages 57, 59).

McLennan, Andrew, and Hugo Sonneschein. 1991. "Sequential Bargaining as a Noncooperative Foundation for Walrasian Equilibrium." *Econometrica* 59 (5): 1395–1424. (Cited on page 40).

Menzio, Guido, and Nicholas Trachter. 2015. "Equilibrium price dispersion with sequential search." *Journal of Economic Theory* 160 (December): 188–215. doi:10.1016/J.JET.2015.09.004. (Cited on page 80).

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *The Review of Economic Studies* 38 (2): 175–208. http://www.jstor.org/stable/2296779. (Cited on page 91).

Modica, Salvatore, and Aldo Rustichini. 1999. "Unawareness and Partitional Information Structures." *Games and Economic Behavior* 27, no. 2 (May): 265–298. ISSN: 08998256. doi:10.1006/game.1998.0666. http://www.sciencedirect.com/science/article/pii/S0899825698906662. (Cited on page 41).

Moen, Espen R. 1997. "Competitive Search Equilibrium." *Journal of Political Economy* 105, no. 2 (April): 385–411. ISSN: 0022-3808. doi:10.1086/262077. https://doi.org/10.1086/262077. (Cited on page 13).

Mortensen, Dale T., and Randall Wright. 2002. "Competitive Pricing and Efficiency in Search Equilibrium." *International Economic Review* 43 (1): 1–20. (Cited on pages 40, 46).

Mount, Kenneth, and Stanley Reiter. 1974. "The Informational Size of Message Spaces." *Journal of Economic Theory* 8 (2): 161–192. (Cited on pages 36, 69).

Myerson, Roger B. 1993. "Incentives to Cultivate Favored Minorities Under Alternative Electoral Systems." *American Political Science Review* 87 (4): 856–869. doi:10.2307/2938819. http://www.jstor.org/stable/2938819. (Cited on pages 106, 130, 131).

Narasimhan, Chakravarthi. 1988. "Competitive Promotional Strategies." *The Journal of Business* 61 (4): 427–449. (Cited on pages 84, 85).

Neeman, Zvika. 2003. "The effectiveness of English auctions." *Games and Economic Behavior* 43, no. 2 (May): 214–238. doi:10.1016/S0899-8256(03)00005-8. (Cited on page 79).

Neiman, Brent, and Joseph Vavra. 2019. "The Rise of Niche Consumption." Cambridge, MA, August. doi:10.3386/w26134. http://www.nber.org/papers/w26134.pdf. (Cited on page 76).

Nöldeke, Georg, and Larry Samuelson. 2014. "Investment and Competitive Matching." New Haven, CT. http://cowles.econ.yale.edu/. (Cited on page 32).

———. 2015. "Investment and Competitive Matching." *Econometrica* 83, no. 3 (May): 835–896. doi:10.3982/ECTA12349. http://doi.wiley.com/10.3982/ECTA12349. (Cited on pages 13, 16, 21, 26).

Nosal, Ed, Yuet-Yee Wong, and Randall Wright. 2019. "Intermediation in markets for goods and markets for assets." *Journal of Economic Theory* 183:876–906. (Cited on page 38).

O'Driscoll, Gerald P. 1977. *Economics as a Coordination Problem: The Contributions of Friedrich A. Hayek.* Kansas City: Sheed Andrews / McMeel, Inc. (Cited on page 10).

Osbourne, Martin J., and Ariel Rubinstein. 1990. *Bargaining and Markets.* Cambridge, United States: Academic Press. (Cited on page 40).

Ostroy, Joseph M. 1984. "A Reformulation of the Marginal Productivity Theory of Distribution." *Econometrica* 52 (3): 599–630. (Cited on page 91).

Owen, Guillermo. 1982. *Game Theory.* Second Edition. New York: Academic Press. (Cited on page 110).

Perla, Jesse. 2019. "Product Awareness, Industry Life Cycles, and Aggregate Profits." Working Paper. (Cited on page 56).

Persico, Nicola. 2004. "Committee Design with Endogenous Information." *Review of Economic Studies* 71 (1): 165–191. (Cited on page 133).

Pigou, Arthur Cecil. 1920. *The Economics of Welfare.* 4th Editio. London: Macmillan. 1932. (Cited on page 79).

Rahman, David. 2014. "The Power of Communication." *The American Economic Review* 104, no. 11 (November): 3737–3751. ISSN: 0002-8282. doi:10.1257/aer.104.11.3737. http://pubs.aeaweb.org/doi/10.1257/aer.104.11.3737. (Cited on page 97).

Rahman, David, and Ichiro Obara. 2010. "Mediated Partnerships." *Econometrica* 78 (1): 285–308. ISSN: 0012-9682. doi:10.3982/ecta6131. (Cited on page 97).

Renou, Ludovic, and Karl H. Schlag. 2010. "Minimax regret and strategic uncertainty." *Journal of Economic Theory* 145, no. 1 (January): 264–286. doi:10.1016/J.JET.2009.07.005. (Cited on page 79).

Roberson, Brian. 2006. "The Colonel Blotto game." *Economic Theory* 29:1–24. doi:10.1007/s00199-005-0071-5. (Cited on page 130).

Robinson, Joan. 1933. *The Economics of Imperfect Competition.* London: Macmillan. (Cited on page 79).

Roesler, Anne-Katrin, and Balazs Szentes. 2017. "Buyer-Optimal Learning and Monopoly Pricing." *The American Economic Review* 107 (7): 2072–2080. doi:10.1257/aer.20160145. (Cited on pages 77, 79, 80).

Rosenthal, Robert W. 1980. "A Model in which an Increase in the Number of Sellers Leads to a Higher Price." *Econometrica* 48, no. 6 (September): 1575. doi:10.2307/1912828. (Cited on page 80).

Rubinstein, Ariel, and Asher Wolinsky. 1987. "Middlemen." *Quarterly Journal of Economics* 102 (3): 581–594. (Cited on page 38).

Rust, John, and George Hall. 2003. "Middlemen versus Market Makers: A Theory of Competitive Exchange." *The Journal of Political Economy* 111 (2): 353–403. (Cited on page 37).

Sahuguet, Nicolas, and Nicola Persico. 2006. "Campaign spending regulation in a model of redistributive politics." *Economic Theory* 28, no. 1 (May): 95–124. ISSN: 09382259. doi:10.1007/s00199-005-0610-0. http://link.springer.com/10.1007/s00199-005-0610-0. (Cited on pages 107, 130, 131).

Satterthwaite, Mark, and Artyom Shneyerov. 2007. "Dynamic matching, two-sided incomplete information, and participation costs: existence and convergence to perfect competition." *Econometrica* 75 (1): 155–200. (Cited on pages 40, 53, 54).

———. 2008. "Convergence to perfect competition of a dynamic matching and bargaining market with two-sided incomplete information." *Games and Economic Behavior* 63 (2): 435–467. (Cited on pages 40, 54).

Scheuer, Florian, and Kent Smetters. 2018. "How Initial Conditions Can Have Permanent Effects: The Case of the Affordable Care Act." *American Economic Journal: Economic Policy* 10 (4): 302–343. doi:10.1257/pol.20140204. https://doi.org/10.1257/pol.20140204. (Cited on page 14).

Schipper, Burkhard C. 2014. "Unawareness - A gentle introduction to both the literature and the special issue." *Mathematical Social Sciences* 70:1–9. (Cited on page 41).

Selten, Reinhard. 1975. "Reexamination of the perfectness concept for equilibrium points in extensive games." *International Journal of Game Theory* 4 (1): 25–55. ISSN: 1432-1270. doi:10.1007/BF01766400. https://doi.org/10.1007/BF01766400. (Cited on pages 13, 14).

Shimer, Robert. 1996. "Contracts in Frictional Labor Markets." (Cited on page 13).

Smith, Vernon L. 1982. "Markets as Economizers of Information: Experimental Examination of the "Hayek" Hypothesis." *Economic Inquiry* 20, no. 2 (April): 165–179. doi:10.1111/j.1465-7295.1982.tb01149.x. https://doi.org/10.1111/j.1465-7295.1982.tb01149.x. (Cited on page 22).

Sorensen, Alan T. 2000. "Equilibrium Price Dispersion in Retail Markets for Prescription Drugs." *Journal of Political Economy* 108 (4): 833–850. (Cited on page 42).

Spulber, Daniel F. 1996. "Market Making by Price-Setting Firms." *Review of Economics Studies* 63:559–580. (Cited on pages 37, 54, 63).

Stigler, George J. 1957. "Perfect Competition, Historically Contemplated." *Journal of Political Economy* 65, no. 1 (February): 1–17. ISSN: 0022-3808. doi:10.1086/257878. http://www.jstor.org/stable/1824830. (Cited on page 75).

———. 1961. "The Economics of Information." *Journal of Political Economy* 69 (3): 213–225. (Cited on page 80).

———. 1968. "Competition." In *International encyclopedia of the social sciences,* edited by David L. Sills and Robert K. Merton, 3:181–186. New York: MacMillan. (Cited on page 1).

Stole, Lars A. 2007. "Price Discrimination and Competition." Chap. 34 in *Handbook of Industrial Organization,* edited by Mark Amstrong and Rob Porter, 3:2221–2299. Amsterdam: Elsevier. ISBN: 9780444824356. doi:10.1016/S1573-448X(06)03034-2. (Cited on page 79).

Taneva, Ina. 2019. "Information Design." *American Economic Journal: Microeconomics.* (Cited on page 3).

Tirole, Jean. 1988. *The Theory of Industrial Organization.* MIT Press. (Cited on page 76).

Tyson, Scott A. 2016. "Information Acquisition, Strategic Voting, and Improving the Quality of Democratic Choice." *Journal of Politics* 78 (4): 1016–1030. ISSN: 0022-3816. doi:10.1086/686027. http://www.journals.uchicago.edu/doi/10.1086/686027. (Cited on page 133).

Varian, Hal R. 1980. "A Model of Sales." *The American Economic Review* 70 (4): 651–659. http://www.jstor.org/stable/1803562. (Cited on page 80).

Werner, Jan. 1987. "Arbitrage and the Existence of Competitive Equilibrium." *Econometrica* 55 (6): 1403–1418. ISSN: 00129682, 14680262. doi:10.2307/1913563. http://www.jstor.org/stable/1913563. (Cited on page 38).

Wittman, Donald A. 1977. "Candidates with policy preferences: A dynamic model." *Journal of Economic Theory* 14 (1): 180–189. (Cited on page 133).

———. 1983. "Candidate Motivation: A Synthesis of Alternative Theories." *The American Political Science Review* 77 (1): 142–157. doi:10.2307/1956016. http://www.jstor.org/stable/1956016. (Cited on page 133).

Zame, William R. 2007. "Incentives, Contracts, and Markets: A General Equilibrium Theory of Firms." *Econometrica* 75, no. 5 (September): 1453–1500. ISSN: 0012-9682. doi:10.1111/j.1468-0262.2007.00799.x. http://doi.wiley.com/10.1111/j.1468-0262.2007.00799.x. (Cited on pages 14, 27, 33).

Appendix A

Chapter 3 Omitted Proofs

A.1 Proof of Proposition 4

Proof. First, consider the competitive mechanism. Using the conditions of $\sum_{i=1}^{k} y^i = \mathbf{0}$ and $py_1^i + y_2^i = 0, \forall i$, implies that the function $(p, y) \to (p, \tilde{y}) \in \mathcal{R}_{++} \times \mathcal{R}_{++}^{2k-1}$, where for $1 \leq i \leq 2k-1$, $\tilde{y}^i = y^i$, is a C^∞-diffeomorphism, thus M_c^k is a $(2k-1) + 1 = 2k$-dimensional manifold.

Second, consider the search mechanism, in this case the price vector of the search equilibrium has k^2 dimensions, while Y_k^s has $2k^2$ dimensions, so, by analogous argument as for the competitive mechanism, M_s^k is a $k^2 + 2(k^2) - 1 = 3k^2 - 1$-dimensional manifold. ∎

A.2 Proof of Proposition 5

Proof. To see this is a Nash equilibrium note that posting the competitive price yields zero profits and for any market-maker deviations either imply negative profits (if prices for purchase are higher than p^* and for selling are lower than p^*) or zero profits (in the case the prices for purchase are lower than p^* and for selling are higher than p^*). To see that this is the unique Nash equilibrium note that if market-makers post prices to make strictly positive profits other market-makers could deviate and make profits by capturing the customers of competitor market-maker by posting more attractive bid and ask prices. ■

A.3 Proof of Proposition 6

Proof. Part 1. Existence and characterization:

There is a unique competitive equilibrium price p^*, since A^j satisfies property 3.3.1 and p^* is also the unique competitive equilibrium price for the subset of traders who are aware of a market-maker.

To construct the candidate equilibrium strategy profile $\{P^j\}_{j \in J}$ we consider pricing strategies described by a pair (p_b, p_s) of offers to buy and sell the good by the market-maker where $p_b \leq p^* \leq p_s$. First consider the monopoly prices $\boldsymbol{p}^M = (p_b^M, p_s^M)$ which satisfies the monopolist market-maker problem:

$$\max_{p_b, p_s} \{(p_s - p_b) \min\{sG(p_b), b[1 - F(p_s)]\}\} . \tag{A.1}$$

In the case of existence of multiple profit maximizing pairs of monopoly prices, let (p_b^M, p_s^M) be the pair of monopoly prices with the lowest difference between the buying and selling price, which implies, as $sG(p_b^M) = b[1 - F(p_s^M)]$, that it is the pair with lowest selling price and highest buying price.

Let $\bar{j}$ be the market-maker with the largest awareness parameter ($m^{\bar{j}} = \max\{m^j\}_{j\in J}$). Let

$$\underline{\alpha} = \prod_{h \neq \bar{j}} (1 - m^h)$$

and let Π^M be the monopoly profit normalized in regards to $m^j \in (0, 1]$, that is

$$\Pi^M = (p_s^M - p_b^M)[sG(p_b^M)].$$

Consider a function $\boldsymbol{p} : [0, 1] \to \mathcal{R}_+^2$ such that $\boldsymbol{p}(\alpha) = (p_b(\alpha), p_s(\alpha))$ is a pair of prices that satisfies

$$[p_s(\alpha) - p_b(\alpha)]G[p_b(\alpha)] = \alpha \Pi^M / s, \tag{A.2}$$

and also satisfies market clearing,

$$sG(p_b(\alpha)) = b[1 - F(p_s(\alpha))]. \tag{A.3}$$

That is, $(p_b(\alpha), p_s(\alpha))$ is the pair of prices that implements a feasible net trade for a monopolist market-maker and yield a fraction α of the monopoly profits. In addition if for some $\alpha \in [0, 1]$ there is more than one such pair of prices then $(p_b(\alpha), p_s(\alpha))$ is the pair with smallest difference between the buying and selling prices, formally, for each $\alpha \in [0, 1]$, $(p_b(\alpha), p_s(\alpha))$ satisfies

$$(p_b(\alpha), p_s(\alpha)) = \arg\min_{(b,s)} \{|b - s| : (b, s) \text{ satisfies } A.2, A.3\}.$$

To see that there exists at least one pair of prices that satisfies A.2 and A.3 note that profits for $p_b = p_s = p^*$ are zero and imply $sG(p^*) = b[1 - F(p^*)]$, while profits for $p_b = p_b^M$ and $p_s = p_s^M$ are Π^M and they also satisfy market clearing. As G and F are continuous therefore for any $p_s \in [p^*, p_s^M]$ there exists a (unique) buying price $p_b(p_s)$ that satisfies $sG(p_s) = b[1 - F(p_b(p_s))]$ and continuity of G and F also imply that profits vary continuously from 0 at $p_s = p^*$ to Π^M at $p_s = p_s^M$, by the intermediate value theorem any profit level between 0 and Π^M can be attained by some pair of prices $(p_s, p_b(p_s))$ with $p_s \in [p^*, p_s^M]$.

Note also that if a pair of prices is feasible for a monopolist market-maker then feasibility also holds in the case of competition between pairs of prices given by $\{\boldsymbol{p}(\alpha) : \alpha \in [0,1]\}$. To see that consider a market-maker j, if a competing market-maker j' offers buying and selling prices that are more attractive to sellers (that is $p_s^{j'} > p_s^j$) then market clearing condition for the monopolist implies these prices must also be more attractive to the buyers (that is, $p_s^{j'} > p_s^j$), since the types aware of j' are distributed in the same way as the types in the trader's population $m^{j'}$ which means that the proportion of clients that are buyers or sellers that j loses to j' is the same (that is, supply of the good purchased by the market-maker decreases by a factor of $1 - m^{j'}$ but demand also decreases by a factor of $1 - m^{j'}$) therefore market clearing still holds.

The candidate equilibrium strategy profile $\{P_j\}_{j \in J}$ is a profile of CDF's on $[\underline{\alpha}, 1]$, that is, $P_j(\alpha)$ is the probability that buying (selling) prices higher (lower) than

$p_b(\alpha)(p_s(\alpha))$ that satisfies the equal profit condition

$$\prod_{h \neq j}(1 - P_h(\alpha)m^h)[p_s(\alpha) - p_b(\alpha)]sG[p_b(\alpha)] = \underline{\alpha}\Pi^M, \tag{A.4}$$

where

$$\underbrace{1 - m^h}_{\text{Prob. } h \notin A^i} + \underbrace{[1 - P_h(\alpha)]m^h}_{\text{Prob.} h \in A^i \text{ and } (p_b^h < p_b(\alpha) \text{ or } p_s^h > p_s(\alpha))} = 1 - P_h(\alpha)m^h, \tag{A.5}$$

is the probability that a trader chooses to transact with the market-maker j over competitor h and $\frac{\alpha}{s}\Pi^M$ is the profit margin of a market-maker when posting prices at the minimum profitability level (lower bound for sales, upper bound for purchases), which means its selling (buying) prices undercuts (tops) all competitors. Note that A.4 implies that $P_h(\underline{\alpha}) = 0$ as $[p_s(\underline{\alpha}) - p_b(\underline{\alpha})]sG[p_b(\underline{\alpha})] = \underline{\alpha}\Pi^M$.

To check that this is an equilibrium:

Note that any prices not in the support of equilibrium strategies $\mathcal{P} = \{(p_b(\alpha), p_s(\alpha)) : \alpha \in [\underline{\alpha}, 1]\}$ lead to strictly lower profits: If we consider prices $(p_b(\alpha), p_s(\alpha))$ that satisfy A.2 and A.3 defined for $\alpha < \underline{\alpha}$, profits are strictly lower by construction. For prices $(p_b, p_s) \in [p_b^M, p_b(\underline{\alpha})] \times [p_s(\underline{\alpha}), p_s^M] \cap \mathcal{P}$ they either yield strictly lower profits because they are undercut by prices which would achieve similar profitability in the case the market-maker were a monopolist or they are not feasible (i.e. the market-maker promises to sell more than it purchases).

Given the sharing rule, it is easy to check that profits are constant on the support of $\{P^j\}$ for each j. If there is only one market-maker j with the largest awareness parameter m^j then the equal profit condition A.4 implies that there is atom of probability at $P^j(\boldsymbol{p}^M)$, the monopoly price, in the mixed strategy of the largest market-maker. The sharing rule implies that traders always choose to trade with $h \neq j$ if h posts the monopoly prices $\boldsymbol{p}^M$, hence its profits do not fall discontinuously on the support of the equilibrium strategy $[\underline{\alpha}, 1]$ as $\boldsymbol{p}(\alpha) \to \boldsymbol{p}^M$.

Part 2. Uniqueness:

Note that any feasible pricing strategy for the market-maker is a pair of buying and selling prices that are consistent with market-clearing. Note that pricing strategies on the set of pairs of prices $\{\boldsymbol{p}(a) : a \in [0,1]\}$ are weakly dominant, as any feasible pricing strategy (p_b, p_s) yields the same profits as a strategy $\boldsymbol{p}(a)$ for some $a \in [0,1]$ then sellers and buyers prefer the buying and selling prices $\boldsymbol{p}(a)$. Therefore, for any market-maker posting a pair of prices $(p_b, p_s) \notin \{\boldsymbol{p}(a) : a \in [0,1]\}$ cannot be a best response to a best response. Hence, any candidate for Nash equilibrium consists of distributions over prices in $\{\boldsymbol{p}(a) : a \in [0,1]\}$.

As we are restricted in our candidate equilibrium strategies to prices in $\{\boldsymbol{p}(a) : a \in [0,1]\}$ then the proof of uniqueness of equilibrium is a proof of uniqueness over distributions on $[0,1]$. Consider an equilibrium strategy profile F, undercutting arguments imply that $F = \{F^j\}_{j \in J}$ is non-degenerate and clearly the upper bound of the support for at least a pair of market-makers must include the monopoly price (otherwise its a profitable deviation to post the monopoly price).

The union of the supports for the strategies must be convex; otherwise market-makers could increase profits by posting prices in the complement of the support. Additionally, the supports for the mixed strategies of individual market-makers must be convex; otherwise the equal-profit condition will be violated.

Further, equal-profit conditions are required to hold in a mixed strategy equilibrium and these conditions imply that when the interior of the supports overlap, Equation A.4 holds. This implies that, assuming no atoms at the lower bound of the support of the distribution, the lower bound of the supports for any pair of seller types must be the same if the interior of the supports overlap. Which implies that equilibrium strategies for all seller types have convex supports (that is, the union of the supports is a its own convex hull). In addition, note that there cannot be atoms at a lower bound of the support of equilibrium price distributions; otherwise other market-makers have the incentive to post a more attractive pair of prices $\boldsymbol{p}(\alpha)$ for α in the ϵ-neighborhood of the lower bound of the support.

It remains to show that any equilibrium profile of distribution of prices is such that the interior of the supports overlap. That is, for any pair of market-makers the interior of the support of mixed pricing strategies must overlap.

To see that suppose, without loss of generality, that there is an equilibrium strategy profile $\boldsymbol{P}$ such that there is a pair of market-makers j, j' with the same awareness parameter $0 < a^j = a^{j'} < a^k$ who compete against each other posting prices according to a strategy that is described by pair of distributions $P^j, P^{j'}$ on $[0, 1]$ and the price posting function $\boldsymbol{p}$ that maps $[0, 1]$ into pairs of buying and selling prices. The distributions $P^j, P^{j'}$ have the same support $[\underline{\alpha}^*, \overline{\alpha}^*]$ while all other market-makers post prices according to distributions that have their supports in $[\underline{\alpha}, 1]$ with $\underline{\alpha} = \overline{\alpha}^*$. In words, market-makers j and j' compete by posting strictly more attractive prices to buyers and sellers than all others. Let $\overline{K}$ be the market-maker with largest awareness parameter in the subset of market-makers $\hat{J} = J - \{j, j'\}$.

Let $\Pi^o(\alpha)$ be the equilibrium profits per unit of awareness of market-maker o in posting prices $\boldsymbol{p}(\alpha)$ we call that o's equilibrium "profitability". Since $\overline{\alpha}^* = \underline{\alpha}$, the profitability of market-maker j of posting prices $\boldsymbol{p}(\overline{\alpha}^*)$ is

$$\Pi^j(\overline{\alpha}^*)/a^j = (1 - a^{j'})\overline{\alpha}^*\Pi^M. \tag{A.6}$$

If market-maker $o, o \neq j, j'$ post prices $\boldsymbol{p}(\underline{\alpha})$, his or her equilibrium profitability is

$$\Pi^o(\underline{\alpha}) = (1 - a^{j'})(1 - a^j)\underline{\alpha}\Pi^M. \tag{A.7}$$

Note that $\underline{\alpha} = \overline{\alpha}^*$, and therefore the equal-profit condition for j, and A.6, implies that

$$\underline{\alpha}^* = (1 - a^{j'})\underline{\alpha}. \tag{A.8}$$

Finally, equations A.7, and A.8 together imply that for market-maker $o \neq j, j'$ that his or her profitability in posting prices $\boldsymbol{p}(\underline{\alpha}^*)$ satisfies

$$\Pi^o(\underline{\alpha}^*) = \underline{\alpha}^*\Pi^M \tag{A.9}$$

$$= (1 - a^{j'})\underline{\alpha}\Pi^M > (1 - a^{j'})(1 - a^j)\underline{\alpha}\Pi^M \tag{A.10}$$

$$= \Pi^o(\underline{\alpha}). \tag{A.11}$$

The inequality A.10 is a contradiction with $\boldsymbol{P}$ being an equilibrium. Therefore, in equilibrium the interior of the supports must overlap.

Hence, the pricing strategy described by $\{P_j\}_{j\in J}$ and $\{(p_b(\alpha), p_s(\alpha)) : \alpha \in [\underline{\alpha}, 1]\}$ is the unique Nash equilibrium given the sharing rule that the market-maker with the smaller awareness parameter m^j has priority in transactions to buyers and sellers in the case of a tie in prices. ■

A.4 Proof of Proposition 8

Proof. The proof is divided into two parts.

Part 1: For simplicity of the argument first let us consider the case where agents are myopic (only care about present payoffs, that is the discount factor $\beta = \frac{1}{1+r} = 0$) so, in this case, the deterrence game has only one period.

As in the proof of proposition 6 let $\pi^M = (p_s^M - p_b^M)sG(p_b^M)$ be the monopoly profit rate. Consider a profit rate $\pi \in (0, \pi^M)$, and suppose the incumbent 1 considers posting a price $\boldsymbol{p}(\pi) = (p_b(\pi), p_s(\pi))$ that is the pair of buying and selling prices with smallest difference that satisfies $[p_s(\pi) - p_b(\pi)]sG[p_b(\pi)] =$

π. That is, $\boldsymbol{p}(\pi)$ is the pricing strategy that yields a payoff of $\pi \times m^j$ if j is a monopolist.

For $\pi \leq E/m_e$, then if the incumbent posts $\boldsymbol{p}(\pi)$ the cost of entry E is higher than the profits 2 can make after entry by undercutting 1 (as the discount rate is $r = 1$ the market-makers only care about present profits). Therefore 2 does not enter. Therefore, if the incumbent posts $\boldsymbol{p}(\pi)$ for $\pi = E/m_e$ (the highest profit margin that deters entry) and the entrant playing "no entry", it is an equilibrium if 1 has no incentive to deviate.

Suppose that 1 deviates from $\boldsymbol{p}(\pi)$ and posts prices according to the mixed strategy equilibrium described in Proposition 6 (for instance), suppose E is low enough so that 2's profits in mixed strategy equilibrium (given by $m_e(1 - m_e)\pi^M$) are higher than the entry cost. Then, 1's profits are $(1 - m_e)\pi^M$. If E is larger than $m_e\underline{\alpha}\pi^M = m_e(1 - m_e)\pi^M$ then, for the monopolist a pure strategy$\boldsymbol{p}(\pi)$ for $\pi = E/m_e$ that deters entry yields strictly higher profits than if the monopolist plays the mixed strategy P^j which yields profits of $(1 - m_e)\pi^M$, since deterrence profits π satisfy

$$\pi = E/m_e > \frac{m_e\underline{\alpha}\pi^M}{m_e} = (1 - m_e)\pi^M.$$

That is if $m_e\underline{\alpha}\pi^M < E$ then the only equilibrium is for the monopolist to post $\boldsymbol{p}(\pi)$ for $\pi = E/m_e$, as even higher profit margin encourages undercutting by the entrant and playing the mixed strategy. If $E < m_e\underline{\alpha}\pi^M$, then $\pi < (1 - m_e)\pi^M$ and the only equilibrium is entry of 2 and the mixed strategies (P^1, P^2) described in the proof of Proposition 6.

Part 2: This logic can be extended to the environment where agents are not myopic and instead have a common discount factor $\beta \in (0, 1)$. Then payoffs of 1 and 2 in the Markov perfect equilibrium after entry are respectively

$$U_e^1 = \sum_{t=0}^{\infty} \beta^t (1 - m_t^2) \pi^M, \tag{A.12}$$

$$U_e^2 = \sum_{t=0}^{\infty} \beta^t m_t^2 (1 - m_t^2) \pi^M, \tag{A.13}$$

where t is the number of periods after entry so $\{m_t^2\}_t$ is the sequence of awareness parameters for 2 that satisfies 3.9 for $t > 0$ and $m_0^2 = m_e$.

To implement a strategy of entry deterrence in this setting the monopolist should be able to commit to a pricing strategy in period 0, that is, choose a pricing schedule in period 0 that is valid for all future periods. Otherwise, after entry of market-maker 2, if market-maker 1 cannot commit to a pricing strategy they will play the Markov perfect equilibrium with payoffs U_e^1, U_e^2 for 1 and 2 starting in period 1. If discount rates are low enough, so payoffs in period 0 do not matter much, then for low E it is easy to see that $\beta U_e^2 > E$ so without commitment there does not exist a Markov perfect equilibrium that deters entry in this case.

Therefore, suppose the monopolist commits to the strategy of posting $\boldsymbol{p}(\pi)$ for profits $\pi \in [0, \pi^M]$ for every period. The present value of 2's profits conditional on entry when 1 is following its commitment $\boldsymbol{p}(\pi)$ is bounded above by

$$U_d^2(\pi) = \sum_{t=0}^{\infty} \beta^t m_t^2 \pi.$$

To deter entry π must imply that $U_e^2(\pi) \leq E$ therefore it must satisfy

$$\pi \leq E/\left(\sum_{t=0}^{\infty} \beta^t m_t^2\right).$$

The present value of the payoffs for the strategy of entry deterrence for 1 are

$$U_d^1 = \frac{\pi}{(1-\beta)} = \frac{E}{(1-\beta)\left(\sum_{t=0}^{\infty} \beta^t m_t^2\right)} \tag{A.14}$$

To allow for the equilibrium with entry deterrence the monopolist must find deterring entry profitable: $U_d^1 \geq U_e^1$. Clearly, the equation A.14 implies that for an entry cost E high enough we have $U_d^1 > U_e^1$.

Take $\{\beta_n\}$ such that $\lim \beta_n = 1$ and let $\{E_n\}$ be a sequence of entry costs such that

$$U_d^1(E_n, \beta) \geq U_e^1(\beta). \tag{A.15}$$

There conditions such that $\exists\{E_n\}$ such that $\pi \to 0$ (which means there exists a $\{E_n\}$ that is bounded above by some $\overline{E}$). To see that note that since $m_t^2 \in (0, 1], \forall t$ with some $T > 1$ such that $m_t^2 < 1$ for all $t \leq T$ therefore A.14 implies that $U_d^1 > E$. On the other side, $U_e^1(\beta) = \Delta(\beta)\pi^M/(1-\beta)$ for some $\Delta \in (0, 1)$, as $\beta \to 1$ since $\lim m_t^2 = 1$ implies that $\lim_{\beta\to 1} \Delta(\beta) = 0$. Want to find conditions so that $U_e^1(\beta)$ is bounded above. One such condition is that m_t^2 converges to 1 fast enough so that $\sum_{t=0}^{\infty}(1 - m_t^2) \leq C$, then $U_e^1(\beta) \leq C \times \pi^M$.

For $E_n \geq C \times \pi^M$ then deterrence is the only equilibrium, fix a sequence $\{E_n\}$ with $E_n = \overline{E} = C, \forall n$. Then $\pi \to 0$ as $\beta \to 1$ and therefore $\boldsymbol{p}(\pi)$ converges to

$p_s = p_b = p^*$ as the discount rate r falls to zero and the equilibrium allocation must converge to the competitive equilibrium.

Finally, consider the case where the monopolist can choose a sequence of pairs of prices to post instead of being restricted to a single profit margin. The overall situation is similar but with added tedious notation. The monopolist chooses a sequence of profit shares $\{\pi_t\}_{t=0}^{\infty}$ with corresponding sequence of pairs of bid and ask prices $\boldsymbol{p}(\pi_t)$. To deter entry the sequence $\{\pi_t\}$ must satisfy

$$\sum \beta^t m_t^2 \pi_t \geq U_e^2, \tag{A.16}$$

the profits of the monopolist under this strategy are

$$U_d^1 = \sum \beta^t \pi_t. \tag{A.17}$$

Hence the profit maximizing strategy for the monopolist is to choose among the sequences that satisfy the deterrence condition A.16 the one that maximizes A.17. Given that $m_t^2 \to 1$ and is strictly increasing there is a unique profit maximizing sequence $\{\pi_t\}$ where the monopolist "frontloads" by extracting the highest profits in the early periods as the entrant is relatively constrained by m_t^2 being smaller than in later periods from taking advantage of these higher margins. These profits are strictly higher than the profits from the strategy to commit to constant prices $(\pi/(1-\beta))$ and therefore the previous arguments also apply in this case. ■