THREE PAPERS ON STRATEGIC DECISION MAKING IN DIGITALLY MEDIATED MARKETS

by

Anna V. Osepayshvili

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Information)
in The University of Michigan
2009

Doctoral Committee:

Professor Jeffrey K. MacKie-Mason, Chair
Professor Yan Chen
Professor Michael P. Wellman
Assistant Professor Rahul Sami

To the love of my life and best friend, Dmitri

ACKNOWLEDGMENTS

This dissertation consists of three research projects on which I have worked in collaboration with different teams of co-authors. Chapter 1 is joint work with Michael P. Wellman, Jeffrey K. MacKie-Mason, and Daniel Reeves. Rahul Suri contributed to the equilibrium analysis in Section 1.6. My major contribution is the analysis of price-predicting strategies. This work appeared as a journal publication in 2008 (Wellman, Osepayshvili, MacKie-Mason, and Reeves (2008)) and is reproduced here without modifications.

Chapter 2 is based on work done in collaboration with Scott A. Fay and Jeffrey K. MacKie-Mason (Fay, MacKie-Mason, and Osepayshvili (2006), working paper). In particular, Fay proposed the model of heterogeneous consumer preferences (Section 2.3.2). Fay and MacKie-Mason also contributed, among other things, to the choice of parameters of the preference distributions described in Section 2.4.1. Some of the work originally appeared elsewhere. In particular, I relied on Fay and MacKie-Mason (1999) for the material in Section 2.2, which we incorporated in the working paper (Fay et al., 2006).

Finally, Chapter 3 was written under the supervision of Yan Chen and Jeffrey K. MacKie-Mason. Yan Chen supervised every stage of the study and greatly contributed to the experimental design. MacKie-Mason proposed the research question for this work and also provided his expertise in the statistical evaluation of the data.

My work on Chapters 1 and 2 was supported in part by grant IIS-0414710 from the National Science Foundation awarded to Jeffrey K. MacKie-Mason and Michael P. Wellman. My work on Chapter 3 was supported by grant IIS-0325837 from the National Science Foundation awarded to Paul Resnick and Yan Chen and 2008 Spring/Summer Research Grant from Horace H. Rackham School of Graduate Studies awarded to Yan Chen. I also thank the School of Information at the University of Michigan, for the generous funding and for giving me the opportunity to learn from the best teachers and work with the best researchers.

I am very grateful to my co-authors and committee members, for their collaboration, academic guidance, and financial support. Jeff MacKie-Mason deserves a special mention. Like many of his students, I will always remember the day when Jeff

agreed to be my advisor. Much to my relief, there were no background requirements in Economics or Computer Science. All I was asked to do was to come back in a week with sixty[1] open research questions that could extend his work, learn two programming languages in the summer, and earn a Master's degree in Economics along the way. Thank you for giving me a chance and for your immeasurable confidence in me.

Needless to say, I wouldn't stand a chance without the support of my School of Information cohort, the Incentive-Centered Design group, and especially Lian Jian, Tapan Khopkar, Tracy Liu, John Lin, Rick Wash, and Ben Stearns. Thank you all for being my study buddies, coming to my practice talks, and giving me rides home. I also thank my family and friends, for helping me remember that there just might be life before graduation.

Finally, I thank my fiancé Dmitri. A talented researcher, a caring friend, and an amazing cook, Dmitri has been my muse and life support. Thank you for your love and encouragement.

[1]Meeting notes are available upon request.

PREFACE

Each chapter of this dissertation addresses some of the challenges that commonly characterize strategic interaction through digitally mediated markets. Examples of such markets include online auctions, financial markets, matching systems, recommender systems, music recording production, online archives of academic journals, etc. In a strategic interaction, the outcome of one's decision depends on someone else's actions, and therefore one has to anticipate others' behavior to make better choices. The task may be challenging, especially if the set of choices is large or if one knows little about others to be able to anticipate their likely choices. The development of information technology has exacerbated these challenges in many markets, at the same time offering tools to support decision making in complex settings. The online auction site eBay,[2] which relies on information technology for its routine operations, presents an example of the shift toward higher strategic complexity and uncertainty about other market participants. In particular, eBay concurrent independent auctions may offer related items, but the bidding rules do not allow bidders to express such relationships. Consider, for example, computer components. Bidders usually have a good idea of how much a particular configuration is worth to them, but for a variety of technical and computational reasons, the option of bidding on multiple components is oftentimes not available. Therefore, bidders need to anticipate the final price of each component and think about all likely contingencies. To make a good prediction about final prices, it helps to know how much others are willing to pay. However, because online marketplaces bring together many participants from geographically dispersed locations, an average bidder may find the task daunting. Such uncertainty and the need to keep track of multiple auctions and perform calculations is what makes eBay concurrent markets particularly complex. On the other hand, information technology can be used to automate some computationally intensive tasks that are important for making informed decisions. Automated stock trading tools (e.g., SmartQuant,[3] Interactive

[2] `http://www.ebay.com`
[3] `http://www.smartquant.com/products.php`

Brokers,[4] and RightEdge,[5] among others) are examples of software that has been designed to assist traders in decision making. Such software includes tools for graphical representation and performing quantitative analysis of financial data, computing trends, and developing and running new trading strategies.

Each of my chapters is motivated by some of the challenges that technology has introduced for market participants, as well as by new opportunities it offers. In Chapter 1, my co-authors and I apply novel heuristic methods to search for bidding strategies for Simultaneous Ascending Auctions, in which exhaustive consideration of the full strategy space is intractable. In Chapter 2, I apply the methodology of Chapter 1 to a different problem: that of competing firms that bundle information goods. This chapter is motivated by the drastic changes that digitization of information goods has introduced into the cost structure and by new opportunities to sell information items individually as well as in bundles (such as CD recordings or scholarly journals). Finally, Chapter 3 is motivated by the development of online mechanisms and, in particular, by the question as to whether Bayesian Nash equilibrium is a reasonable solution concept for such mechanisms. The key issue here is that extremely limited information conditions, which often characterize online settings, do not meet the assumptions of Bayesian games. I conduct a human-subject experiment to understand under what conditions Bayesian Nash equilibrium can arise as a result of learning.

[4] http://www.interactivebrokers.com/ibg/main.php

[5] http://www.rightedgesystems.com/

TABLE OF CONTENTS

LIST OF FIGURES

Figure

LIST OF TABLES

Table

LIST OF APPENDICES

Appendix

ABSTRACT

Technological developments have been reshaping existing markets and giving rise to new ones. In my dissertation, I address several questions that emerged as a result of these changes. In the first two chapters, I study strategies in two different markets: a simultaneous ascending auction (SAA) and an information-good market. The challenge in both cases is analytical intractability of the problem of finding equilibrium strategies. I therefore apply heuristic methods and computer simulation to incrementally search for improvements relative to known leading strategies. For SAAs with complementary goods, in which bidders often face an *exposure risk*, the result of this search procedure is a price-predicting strategy with *self-confirming predictions*. If the goods are substitutes, the results suggest that a simple *demand-reduction* strategy yields the highest expected surplus. In chapter two, I apply similar methods to study bundling strategies in an information-good market. I restrict my attention to three bundling schemes: *pure unbundling*, *pure bundling*, and *mixed bundling*, where the latter is a choice between the two pure schemes. I find that when consumers' tastes are highly diverse, the relative profitability of the schemes is defined by the consumer preference distribution, even when the marginal costs are zero. This holds for a duopoly as well as a monopoly. Under duopoly, restricting firms to pure schemes may result in higher equilibrium profits relative to mixed bundling. I also provide comparison of the market efficiency, profits and consumer welfare under a monopoly and a duopoly. In chapter three, I take the perspective of a market designer and conduct a human-subject experiment to study learning of Bayesian Nash equilibrium. The study is motivated by the increasing need to test game-theoretic predictions in the relatively new online environment. I focus on three treatment variables: available information, existence of a *potential*, and *supermodularity* of the game. The findings are largely consistent with the theory of supermodular and potential games. In addition, the evidence suggests that there is a strong interaction effect between supermodularity and information, which I attribute to differences in the nature of irrationality that prevails under the different information conditions.

CHAPTER 1

Bidding Strategies for Simultaneous Ascending Auctions

1.1 Introduction

A *simultaneous ascending auction* (SAA) (Cramton, 2005) allocates a set of m related goods among n agents via separate, concurrent English auctions for each good. This is characteristic of a variety of related but not identical real-world auction mechanisms, such as concurrent independent auctions on eBay, power markets, spectrum auctions in many countries, and other explicitly designed trading environments (Milgrom, 2003). Some of the key strategic issues presented by SAAs apply whenever there are concurrent markets for interrelated goods, even if those markets are not formal auctions.

Simultaneity is significant only if demands (or supplies) for the various goods are interrelated. We address here some of the challenges bidders with such demands face when formulating their strategies for participation in SAAs. Interrelated demands generally exhibit *complementarity* or *substitutability* (or both), each of which induces characteristic bidding problems.

To study bidding strategies in the face of the strategic challenges presented by complementarity or substitutability, we intentionally abstract from any single application. There are features specific to spectrum auctions, for example, that we do not address, just as there are unaddressed features specific to simultaneous eBay auctions and other particular SAA environments. In hope of producing results generalizable to a range of applications, we analyze a generic SAA exhibiting a few important characteristics that are common across most specific settings.

Complementarity manifests when an agent's value for a good is greater if it also obtains one or more other goods (Lehmann, Lehmann, & Nisan, 2006). For example, an airline passenger may wish to obtain two connecting segments to complete a trip. Goods exhibit complementarity from the perspective of an agent when her valuation for those goods is *superadditive*. Let X, Y, and Z be sets of goods such that $Y \cup Z = X$ and $Y \cap Z = \emptyset$. Given a quasi-linear valuation function, $v \colon 2^{|X|} \rightarrow \mathbb{R}$, that assigns value to possible subsets of X, superadditive preference for Y and Z means that $v(X) > v(Y) + v(Z)$. In other words, the combined bundle X is worth more than the sum of its parts. As a special case, if goods in a set are each worthless without the others, they are *perfect complements*. We say that a valuation exhibits complementarities if there are some subsets of goods for which preference is superadditive.

When the inequality is reversed, the valuation is *subadditive*, which occurs for example when goods are substitutes. Flights on the same route by different airlines would typically be considered substitutes, as would flights to two candidate vacation destinations. Technically, goods are substitutes when raising the price of one does not decrease demand for others – that is, for any optimal bundle before the price increase there is an optimal bundle post-increase that includes at least as much demand for all goods that did not increase in price. Substitutability is a strictly stronger condition than subadditivity (Lehmann et al., 2006). An important extreme case of substitutability is perfect substitutes or *single-unit demand* (Gul & Stacchetti, 1999), where for all $Y \subseteq X$, $v(Y) = \max_{i \in Y} v(\{i\})$. If, in addition, goods are (for this agent) *homogeneous* then they are 1:1 perfect substitutes.

Concurrent auctions with interdependent goods are strategically challenging because agents bid separately in auctions for each item, but willingness-to-pay depends nontrivially on which combination of items the agent ultimately wins. When bids represent non-repudiable offers, submitting bids to separate auctions entails an *exposure problem*. With complementarities, if an agent bids on a set of items based on her willingness-to-pay for the set, she may pay more than her valuation for the subset she actually wins. With substitutes, an agent bidding based on willingness-to-pay for individual goods risks paying more for a set than it is worth. The SAA mechanism makes it easy for agents to avoid exposure in the case of substitutes. Since a price increase for one good cannot decrease demand for others, the agents can manage their bids to ensure they are never winning more goods than they want at the current prices. With any violation of substitutability, however, a bidder cannot in general obtain a desired package without incurring some exposure risk.

The exposure problem motivates mechanisms that take complementarities directly into account, such as *combinatorial auctions* (Cramton, Shoham, & Steinberg, 2005; de Vries & Vohra, 2003), in which the auction mechanism determines optimal packages based on agent bids over bundles. Although such mechanisms may provide an effective solution in many cases, there are often significant barriers to their application (MacKie-Mason & Wellman, 2005). Indeed, SAA-based auctions are often deliberately adopted, despite awareness of strategic complications (Milgrom, 2003; McAfee & McMillan, 1996).

A second strategic problem for bidders is accounting for *own price effects*: the impact of their own bids on resulting prices. For example, a bidder winning q units may find that bidding her incremental value for the $q+1$st unit results in an increase in price paid for the first q units. The strategy of shading bids to take account of this inframarginal surplus loss is known as *demand reduction* (Ausubel & Cramton, 2002; Weber, 1997).

Given exposure and own price effects, it is clear that bidding willingness-to-pay is generally not optimal. Worse for designers, researchers, and bidders, auction theory to date (Krishna, 2002) has little to say about how one should bid in simultaneous markets with substitutes or complements. There exists no useful analytical characterization of equilibria for SAA games. Moreover, the best-response strategies to even simple specified bidding policies can be surprisingly complex (Reeves et al., 2005). Simulation studies shed light on some strategic issues (Csirik, Littman, Singh, & Stone, 2001), as have accounts of strategies employed in specific auctions (Cramton, 1995; Weber, 1997), but the game is too complex to admit definitive strategic recommendations. We cannot emphasize this point enough: there is a striking gap in the literature, and the main motivation for the novel empirical methods we employ is that analytically deriving equilibrium strategies appears intractable for nontrivial SAA games.

We employ a different approach to analyze bidding strategies, which we elsewhere describe as a *computational reasoning* (MacKie-Mason & Wellman, 2005) or *empirical game-theoretic* methodology (Wellman, 2006) for analyzing mechanisms and strategies. We begin with an explicit formulation of the resource allocation problem, generate a set of candidate parametrized strategies, then simulate the game for various profiles of strategy parameters. Through simulation, we in effect convert an extensive-form game of incomplete information with high-dimensional strategy space into a normal-form game over the restricted set of strategies defined by the instances of strategy parameters explored. We then use standard tools to solve the restricted-strategy (yet often still quite large) normal-form games, and analyze the results. For the families of candidate strategies we study, we are able to characterize those which participate in equilibria of the transformed game, and the quality of the resulting outcomes.

One advantage of this fundamentally empirical method is that if others believe they have superior strategies, it is straightforward to apply the method *incrementally* to evaluate the new candidates with respect the best-performing strategies known to date. This is important because the SAA environment is so complex, and in any

SAA the specific rules may call for variations on the basic strategy family we study. For example, some auctions impose activity rules, which introduces an eligibility management problem into the design of bidding strategies. Budget constraints may also affect the design of bidding strategies. We do not claim that our present analysis covers the entire space of bidding-strategy design for SAAs. We do claim that we provide some of the first systematic evidence for a successful family of strategies in a generic SAA. We also claim that strategic lessons from this generic environment will be a useful starting point for those designing strategies for the rules of particular SAA environments they face.

We next proceed with a formal specification of the problem and of the generic SAA mechanism we study.

1.2 The Simultaneous Ascending Auctions (SAA) Domain

The formal specification of the *SAA game* includes a number of agents, n, a number of goods, m, a type distribution that yields valuation functions v_j for the agents $j \in \{1, \ldots, n\}$,[1] and a specification of the SAA mechanism rules. In general the SAA mechanism comprises m separate auctions, one for each good, that operate over multiple rounds of bidding. In the generic SAA version we study, bidding is synchronized so that in each round each agent submits a bid in every auction in which it chooses to bid. At any given time, the *bid price* on good i is β_i, defined to be the highest bid b_i received thus far, or zero if there have been no bids. The bid price along with the current winner in every auction is announced at the beginning of each new round. To be admissible, a new bid must meet the *ask price*, i.e., the bid price plus a bid increment (which we take to be one w.l.o.g., allowing for scaling of the agent values): $b_i^{new} \geq \beta_i + 1$. If an auction receives multiple admissible bids in a given round, it admits the highest, breaking ties randomly. An auction is *quiescent* when a round passes with no new admissible bids, i.e., the new bid prices $\boldsymbol{\beta}^{new} = \boldsymbol{\beta}$ which become the final prices $\boldsymbol{p}$. When every auction is simultaneously quiescent they all close, allocating their respective goods per the last admitted bids.

An agent's *current information state*, $\boldsymbol{B}$, comprises the current bid prices, $\boldsymbol{\beta}$, along with a bit vector indicating which goods the agent is currently winning. Let $\boldsymbol{\mathcal{B}}$ denote the set of possible current information states. A *local bidding strategy* is a mapping $\boldsymbol{\mathcal{B}} \rightarrow \boldsymbol{b}$, where the bid vector $\boldsymbol{b}$ specifies a bid for each of the m auctions. More generally, an agent's bidding strategy maps the *history* of information states to bids. For the present work, we limit consideration to local bidding strategies.

[1]We may include in the type distribution *Nature's type* which determines the random tie-breaking when agents place identical bids.

This is a substantive limitation, ruling out, for example, methods that infer other agents' types from dynamic price patterns, or strategies that punish others' behavior. Nevertheless, the strategic issues we consider primary can be addressed at the level of local bidding strategies, and thus we take the simplification achieved through ignoring history to be worthwhile.[2]

Submitting an inadmissible bid (e.g., $b_i = 0$) is equivalent to not bidding. An agent's payoff – also referred to as its *surplus* – is defined by the auction outcomes, namely, the set of goods it wins, X, and the final prices, $\boldsymbol{p}$:

$$\sigma(X, \boldsymbol{p}) \equiv v(X) - \sum_{i \in X} p_i. \tag{1.1}$$

In the next section we describe a broad characterization that encompasses many SAA bidding strategies in the prior literature, as well as a new category of strategies we propose subsequently. Then we describe two bidding methods from the prior literature that are special cases of the general class. One of these extant methods is itself a strategy family, with bidding behavior that varies dramatically depending on the choice of a continuous parameter, so these prior strategies we analyze represent substantial variation. In Section 1.4 we propose a new bidding approach, which falls into the same broad characterization, but represents yet another substantial variation on the range of strategies we evaluate. This new, price prediction-based strategy family itself encompasses a wide range of methods for generating and using price predictions, which we explore in subsequent sections.

Heuristic strategies are sometimes motivated by bounded rationality (Gigerenzer & Selten, 2001), in which case the primary concern is behavioral realism. Our appeal to heuristics is in the spirit of Rosenthal's approach to defining games over "rule-of-thumb" strategies (Rosenthal, 1993b, 1993a). Given the intractability of exhaustive consideration of the full strategy space, we rely on heuristics to represent key strategic ideas in our domain. Behavioral models may be one source of heuristic elements, though explicit optimization procedures or other sophisticated reasoning may be incorporated in our heuristics as well when they are motivated by potential performance gains.

1.3 Perceived-Price Bidding Strategies

If an agent knew the final prices of all m goods and if those prices did not depend on its own bidding strategy, then its optimal strategy would be clear: bid on a subset

[2]Assuming that agents submit bids for a subset of goods at the minimum increment, the number of states in the strategy space is $|\boldsymbol{\mathcal{B}}|2^m$. Conditioning on a history of length t would expand this space by a factor of 2^{mt}.

of goods that maximizes its surplus at known prices. When prices are uncertain or bid-dependent, this is not optimal, but may nevertheless serve as a useful starting point. In this section, we define a class of bidding strategies that generalizes this approach by selecting a subset of goods that maximizes surplus at *perceived* prices.

Definition 1.1 (Perceived-Price Bidder)**.** *A* perceived-price bidder *is parametrized by a* perceived-price function $\rho\colon \boldsymbol{\mathcal{B}} \to \mathbb{Z}_*^m$ *which maps the agent's information state,* $\boldsymbol{B}$*, to a (nonnegative, integer) perceived-price vector,* $\rho(\boldsymbol{B})$*. It computes the subset of goods*

$$X^* = \arg\max_X \sigma(X, \rho(\boldsymbol{B}))$$

breaking ties in favor of smaller subsets and lower-numbered goods.[3] *Then, given* X^**, the agent bids* $b_i = \beta_i + 1$ *(the ask price) for the* $i \in X^*$ *that it is not already winning.*

A perceived-price bidding strategy is defined by how the agent constructs the perceived price from its information state. We now define two versions of the function ρ, corresponding to perceived-price bidding strategies well-studied in prior literature. In Section 1.4 we define the newer price-prediction perceived-price strategies we analyze in this article. Our discussion focuses on the particularly challenging case of superadditive preference – complementary goods. We return to address the case of substitutable goods in Section 1.7.

1.3.1 Straightforward Bidding

One example of a perceived-price bidder is the widely studied *straightforward bidding* (SB) strategy.[4] An SB agent sets $\rho(\boldsymbol{B})$ to *myopically perceived prices*: the bid price for goods it was winning in the previous round and the ask price for the others:

$$\rho_i(\boldsymbol{B}) = \begin{cases} \beta_i & \text{if winning good } i \\ \beta_i + 1 & \text{otherwise,} \end{cases} \tag{1.2}$$

where $\boldsymbol{\beta}$ is the current bid prices.

Straightforward bidding is a reasonable strategy in some environments. When all agents have single-unit demand, and value every good equally (i.e., the goods are all

[3] More precisely: when multiple subsets tie for the highest surplus, the agent chooses the smallest. If the smallest subset is not unique it picks the subset whose bit-vector representation is lexicographically greatest. (The bit-vector representation $\boldsymbol{\omega}$ of $X \subseteq \{1, \ldots, m\}$ has $\omega_i = 1$ if $i \in X$ and 0 otherwise. For example, the bit-vector representation of $\{1, 3\} \subseteq \{1, 2, 3\}$ is $\langle 1, 0, 1\rangle$.) This tie-breaking scheme is somewhat arbitrary, and we expect alternative choices would be inconsequential. We describe our version here in detail to facilitate replication of our experimental results.

[4] We adopt the terminology introduced by Milgrom (2000). The same concept is also referred to as "myopic best response", "myopically optimal", and "myoptimal" (Kephart, Hanson, & Sairamesh, 1998).

1:1 perfect substitutes), the situation is equivalent to a problem in which all buyers have an inelastic demand for a single unit of a homogeneous commodity. For this problem, Peters and Severinov (2006) show that straightforward bidding is a perfect Bayes-Nash equilibrium.

If agents have additive utility, i.e., $v(Y) = \sum_{i \in Y} v(\{i\})$, then they can treat the auctions as independent and in this case too, SB is in equilibrium. To see this, consider the case that all other agents are playing SB with additive preference. Then your bid in one auction does not affect your surplus in another. This implies the auctions can be treated independently and SB is a best response.

The degenerate SAA with $m = 1$, i.e., a single ascending auction, is strategically equivalent to a second-price sealed-bid auction (Vickrey, 1961). In other words, SB is a weakly dominant strategy in a single ascending auction, similarly to "truth-telling" in a second-price sealed-bid auction.[5] For $m > 1$, however, the joint strategy space allows *threats* such as "if you raise the price on my good I will raise it on yours." These will then support demand-reduction equilibria, even in the additive case. Thus, although SB is a good strategy and is in equilibrium for some special-case environments without complementarities, it is not (even weakly) dominant.

Up to a discretization error, the allocation in an SAA with single-unit demand is efficient when agents follow straightforward bidding. It can also be shown (Bertsekas, 1992; Wellman, Walsh, Wurman, & MacKie-Mason, 2001) that the final prices will differ from the minimum unique equilibrium prices by at most $\min(m, n)$ times the bid increment. The value of the allocation, defined to be the sum of the bidder surpluses, will differ from the optimal by at most the bid increment times $\min(m, n)(1 + \min(m, n))$.

Unfortunately, none of these properties hold for general preferences. The final SAA prices can differ from the minimum equilibrium price vector, and the allocation value can differ from the optimal, by arbitrarily large amounts (Wellman et al., 2001). And most importantly, SB need not be a Nash equilibrium.

Example 1.3.1. *There are two agents, with values for two goods as shown in Table 1.1. One admissible straightforward bidding path*[6] *leads to a state in which agent 2 is winning both goods at prices (15,14). Then, in the next round, agent 1 would bid 15 for good 2. The auction would end at this point, with agent 1 receiving*

[5] Technically, this equivalence applies to a strategically restricted version of the ascending auction which does not allow arbitrary bids above the ask price (and raises the ask price continuously rather than discretely). Otherwise, there exist strategies (albeit pathological) to which SB is not a best response. For example, suppose my policy is to not bid more than $100 unless the bidding starts lower, in which case I will keep bidding indefinitely. The best response to such a strategy requires *jump bidding*.

[6] The realized progression of the SAA protocol depends on tie-breaking.

	$v(\{1\})$	$v(\{2\})$	$v(\{1,2\})$
Agent 1	20	20	20
Agent 2	0	0	30

Table 1.1: A simple problem illustrating the pitfalls of SB (Example 1.3.1).

good 2 and agent 2 receiving good 1, both at a price of 15.

In this example, SB leads to a result with total allocation value 20, whereas the optimal allocation would produce a value of 30. We can construct slightly more complex examples by adding goods and agents, enabling us to magnify the suboptimality to an arbitrary degree.

We see that straightforward bidding fails to guarantee high quality allocations. It is also easy to show that straightforward bidding is not an equilibrium strategy in general. Consider again Example 1.3.1. If agents follow the SB strategy, the mechanism reaches quiescence at prices $\{15, 15\}$. However, it is not rational for agent 2 to stop at this point. If, for example, agent 2 continued bidding, prices would reach $\{21, 20\}$ with agent 2 winning both goods, and the auction would end (assuming agent 1 plays SB). Agent 2 would be better off, with a surplus of -11 rather than -15.

It is clear that SB is not a reasonable candidate for a general strategy in SAA. We show next how a simple parametric generalization to SB can address a key strategic shortfall.

1.3.2 Sunk-Aware Bidding

Another example of perceived-price bidding is the sunk-aware family of bidding strategies. We showed in Example 1.3.1 that in some problems agents following a straightforward bidding strategy may stop bidding prematurely. To motivate the alternative sunk-aware approach, we consider why SB is failing in this situation. In a given round, agents following SB bid on the set of goods that maximizes their surplus at myopically perceived prices (current bid or ask prices). If none of the nonempty subsets of goods appear to yield positive net surplus, the agent chooses the empty set, i.e., it does not bid at all, because the alternative is to earn negative surplus. However, this behavior ignores outstanding commitments: the agent may already be winning one or more goods. If the agent drops out of the bidding, and others do not bid away the goods the agent already is winning, then its alternative surplus could be much worse than if it continued to bid despite preferring the empty bundle at current prices. In the case of an agent dropping out of the bidding on some goods in a bundle of perfect complements, its surplus is negative the sum of the bid prices for the goods in the bundle it gets stuck with. This failure of straightforward bidding is due to ignoring the true opportunity cost of not bidding.

We refer to this property of straightforward bidding as "sunk-unawareness" (Reeves et al., 2005). SB agents bid as if the incremental cost for goods they are currently winning is the full price, β_i. However, if the probability that someone else will outbid the agent for this good is α, then the agent is already committed to an expected payment of $(1-\alpha)\beta_i$. This represents a sunk cost that should not affect rational continuation bidding. We can think of the difference, $\alpha\beta_i$, as a rough measure of the incremental cost the agent incurs by deciding to stay with this good.

To address this limitation of straightforward bidding, we parametrize an alternative family of perceived-price bidding strategies (Definition 1.1) that permits agents to account to a greater or lesser extent for the true incremental cost of goods they are currently winning. We call this strategy "sunk aware". A sunk-aware agent bids as if the incremental cost for goods it is currently winning is somewhere on the interval of zero and the current bid price.

Our sunk-aware strategies generalize SB's method for choosing the perceived-price vector (1.2) through the parameter $k \in [0,1]$:

$$\rho_i(\boldsymbol{B}) = \begin{cases} k\beta_i & \text{if winning good } i \\ \beta_i + 1 & \text{otherwise.} \end{cases}$$

Using this perceived-price vector to define sunk-aware bidders, Definition 1.1 above gives us a complete specification of the agent's bidding strategy. If $k = 1$ the strategy is identical to straightforward bidding. At $k = 0$ the agent is fully sunk aware, bidding as if it would retain the goods it is currently winning with certainty. Intermediate values are akin to bidding as if the agent puts an intermediate probability on the likelihood of retaining the goods it is currently winning. We treat as a special case agents with single-unit demand: their sunk-aware strategy is to bid straightforwardly ($k = 1$) since for such agents SB is a no-regret strategy.

The sunk-awareness parameter provides a heuristic for a complex tradeoff: the agent's bidding behavior changes after it finds itself exposed to the underlying problem (owning goods for which the agent has lower value if not part of a larger package). In our previous study we experimentally determined good settings of the sunk-awareness parameter in various environments (Reeves et al., 2005).

1.4 Prediction-Based Perceived-Price Bidding

Straightforward and sunk-aware bidding represent alternative SAA bidding strategies, distinguished by the way in which bidders formulate "perceived prices" that determine the items on which they bid. In this section we propose yet another class of bidding strategies: price prediction bidding. In this heuristic approach, bidders form predictions of final prices in order to select the items on which they will bid in a given round.

Whenever an agent has non-substitutes preference and chooses to bid on a bundle of size greater than one, it may face exposure. Exposure in SAA is a direct tradeoff: bidding on a needed good increases the prospects for completing a bundle, but also increases the expected loss in case the full set of required goods cannot be acquired. A decision-theoretic approach would account for these expected costs and benefits, bidding when the benefits prevail, and cutting losses in the alternative.

Re-consider agent 2's plight in Example 1.3.1: following SB it is caught by the exposure problem, stuck with a useless good and a surplus of -15. (Other tie-breaking choices result in different outcomes but all of them leave agent 2 exposed and with negative surplus.) If the agent instead plays a fully sunk-aware strategy the result could be an outcome in which it purchases both goods at prices $\{21, 20\}$ for a net surplus of $30 - 41 = -11$. This is better than using SB, but the agent would fare better still by not bidding at all.

The effectiveness of a particular strategy will in general be highly dependent on the characteristics of other agents in the environment. This observation motivates the use of *price prediction*. We would prefer strategies that employ type-distribution beliefs to guide bidding behavior, rather than relying only on current price information as in the sunk-aware strategies (including SB). Forming price predictions for the goods in SAA is a natural use for type-distribution beliefs. In Example 1.3.1, suppose agent 2 could predict with certainty before the auctions start that the prices would total at least 30. Then it could conclude that bidding is futile, not participate, and avoid the exposure problem altogether. Of course, agents will not in general make perfect predictions. However, we find that even modestly informed predictions can significantly improve performance.

We now propose to improve on SB and sunk-aware bidding by using explicit price predictions for perceived prices. Let $F \equiv F(\boldsymbol{B})$ denote a joint cumulative distribution function over final prices, representing the agent's belief given its current information state $\boldsymbol{B}$. We assume that prices are bounded above by a known constant, V. Thus, F associates probabilities with price vectors in $\{1, \ldots, V\}^m$.

We next consider two ways to use prediction information to generate perceived prices. We first define a *point prediction* $\boldsymbol{\pi}$, which anticipates possible exposure risks. Then we define a *distribution prediction* F that explicitly models uncertainty about the exposure prospects. The distribution prediction generates perceived incremental prices, $\boldsymbol{\Delta}$, which account for the likelihood that the agent's current winning bids are sunk costs. As with sunk-awareness, price-prediction strategies for agents with single-unit demand ignore the predictions and play SB.

Before we define our price-prediction strategies we want to make two points. First, we are not (initially) concerned with *how* the agent formulates her beliefs (price predictions), nor the optimality of the prediction method. Rather, we propose strategies that use *some* beliefs. In our experiments we investigate several different predictors.[7] Second, since these are strategies for bidding in iterative auctions, we face the question of how to update the initial price predictions based on information revealed during the bidding process. We do not insist here on identifying the optimal updating procedure; again, we define strategies that incorporate *some* belief updating procedure. Whereas we experiment with different initial price predictors, in this paper we employ only one specific, simple, updating procedure.

1.4.1 Point Price Prediction

Suppose the agent has (at least) point beliefs about the final prices that will be realized for each good. Let $\boldsymbol{\pi}(\boldsymbol{B})$ be a vector of predicted final prices. Before the auctions begin the price prediction is $\boldsymbol{\pi}(\varnothing)$, where $\varnothing$ is the null information state available pre-auctions.

The auctions in SAA reveal the bid prices each round. Since the auctions are ascending, once the current bid price for good i reaches β_i, there is zero probability that the final price p_i will be less than β_i. We define a simple updating rule using this fact: the current price prediction for good i is the maximum of the initial prediction and the myopically perceived price:

$$\pi_i(\boldsymbol{B}) \equiv \begin{cases} \max(\pi_i(\varnothing), \beta_i) & \text{if winning good } i \\ \max(\pi_i(\varnothing), \beta_i + 1) & \text{otherwise.} \end{cases} \tag{1.3}$$

Armed with these predictions, the agent plays the perceived-price bidding strategy (Definition 1.1) with $\rho(\boldsymbol{B}) \equiv \boldsymbol{\pi}(\boldsymbol{B})$. We denote a specific point price-prediction strategy in this family by $PP(\boldsymbol{\pi}^x)$, where x labels particular initial prediction vectors, $\boldsymbol{\pi}(\varnothing)$. Note that straightforward bidding is the special case of price prediction with the predictions all equal to zero: SB $= PP(\mathbf{0})$. If the

[7] We believe that finding the *optimal* predictor to use in a particular strategy is likely to be as computationally infeasible as the problem of finding an optimal strategy.

agent underestimates the final prices, it will behave identically to SB after the prices exceed the prediction. If the agent overestimates the final prices, it may stop bidding prematurely.

1.4.2 Distribution Price Prediction

We generalize the class of price-prediction strategies by taking into account the entire distribution F, rather than just a nominal point estimate (e.g., the expectation of F). We assume the agent generates $F(\varnothing)$, an initial, pre-auction probabilistic belief about the final prices.

As with the point predictor, we restrict the updating in our distribution predictor to conditioning the distribution on the fact that prices are bounded below by $\boldsymbol{\beta}$. Let $\Pr(\boldsymbol{p} \mid \boldsymbol{B})$ be the probability, according to F, that the final price vector will be $\boldsymbol{p}$, conditioned on the information revealed by the auction, $\boldsymbol{B}$. Then, with $\Pr(\boldsymbol{p} \mid \varnothing)$ as the pre-auction initial prediction, we define:

$$\Pr(\boldsymbol{p} \mid \boldsymbol{B}) \equiv \begin{cases} \dfrac{\Pr(\boldsymbol{p} \mid \varnothing)}{\sum\limits_{\boldsymbol{q} \geq \boldsymbol{\beta}} \Pr(\boldsymbol{q} \mid \varnothing)} & \text{if } \boldsymbol{p} \geq \boldsymbol{\beta} \\ 0 & \text{otherwise.} \end{cases} \tag{1.4}$$

(By $\boldsymbol{x} \geq \boldsymbol{y}$ we mean $x_i \geq y_i$ for all i.) For (1.4) to be well defined for all possible $\boldsymbol{\beta}$ we define the price upper bounds such that $\Pr(V, \ldots, V \mid \varnothing) > 0$.

We now use the distribution information to implement a further enhancement to take sunk costs into account in a more decision-theoretic way than the sunk-aware agent. If an agent is currently not winning a good and bids on it, then the expected incremental cost of winning the good is the expected final price, with the expectation calculated with respect to the distribution F. If the agent is currently winning a good, however, then the expected incremental cost of winning that good depends on the likelihood that the current bid price will be increased by another agent, so that the first agent has to bid again to obtain the good. If, to the contrary, it keeps the good at the current bid, the full price is sunk (already committed) and thus should not affect incremental bidding. Based on this logic we define $\Delta_i(\boldsymbol{B})$, the expected *incremental* price for good i.

First, for simplicity, we use only the information contained in the vector of marginal distributions, $(F_1, \ldots, F_m)$, as if the final prices were independent across goods. Define the expected final price conditional on the most recent vector of bid prices, $\boldsymbol{\beta}$:

$$E_F(p_i \mid \boldsymbol{\beta}) = \sum_{q_i=0}^{V} \Pr(q_i \mid \beta_i) q_i = \sum_{q_i=\beta_i}^{V} \Pr(q_i \mid \beta_i) q_i.$$

The expected incremental price depends on whether the agent is currently winning good i. If not, then the lowest final price at which it could win is $\beta_i + 1$, and the expected incremental price is simply the expected price conditional on $p_i \geq \beta_i + 1$,

$$\Delta_i^{\mathrm{L}}(\boldsymbol{B}) \equiv E_F(p_i \mid p_i \geq \beta_i + 1) = \sum_{q_i=\beta_i+1}^{V} \Pr(q_i \mid p_i \geq \beta_i + 1) q_i. \tag{1.5}$$

If the agent is winning good i, then the incremental price is zero if no one outbids the agent. With probability $1 - \Pr(\beta_i \mid \beta_i)$ the final price is higher than the current price, and the agent is outbid with a new bid price $\beta_i + 1$. Then, to obtain the good to complete a bundle, the agent will need to bid at least $\beta_i + 2$, and the expected incremental price is

$$\Delta_i^{\mathrm{W}}(\boldsymbol{B}) = (1 - \Pr(\beta_i \mid \beta_i)) \sum_{q_i=\beta_i+2}^{V} \Pr(q_i \mid \beta_i + 2) q_i.$$

The vector of expected incremental prices is then defined by

$$\Delta_i(\boldsymbol{B}) = \begin{cases} \Delta_i^{\mathrm{W}}(\boldsymbol{B}) & \text{if winning good } i \\ \Delta_i^{\mathrm{L}}(\boldsymbol{B}) & \text{otherwise.} \end{cases}$$

The agent then plays the perceived-price bidding strategy (Definition 1.1) with $\rho(\boldsymbol{B}) \equiv \boldsymbol{\Delta}(\boldsymbol{B})$. We denote the strategy of bidding based on a particular distribution prediction by $PP(F^x)$, where x labels various pre-auction distribution predictions, $F(\varnothing)$.

To recapitulate, in this and the previous section we have formally specified three categories of bidding strategies, encompassed within a single broad but flexible class we call "perceived price" bidding. The first two, straightforward and sunk-aware bidding, have been explored in prior literature; the third, price-prediction bidding, is new. Both sunk-aware and price-prediction are families that admit a wide range of specific strategies, and thus represent a variety of actual bidding behaviors.[8] For the generic SAA we study, when bidders have non-substitute preferences but face exposure (but not other problems such as budget constraints), we think this broad set of strategy candidates captures most of the existing wisdom about SAA strategy design.

In Section 1.6 we intensively analyze and compare the performance of this broad set of bidding strategies in a series of SAA environments with non-substitute preferences. After that (Section 1.7) we specify yet another family of strategies, based on prior literature, and perform a strategic analysis over this enlarged set for environments with substitute goods.

[8] In the next section we define a variety of different price prediction methods, each leading to different bidding behavior.

1.5 Some Methods for Predicting Prices in SAA

In Section 1.4 we define bidding strategies based on point price and distribution-based predictions. These are classes of strategies parametrized by the choice of initial prediction: a vector of predicted final prices in the case of the point predictor, or a distribution of final prices for the distribution predictor. We now present several ways to obtain an initial prediction. Each different prediction method generates a different bidding strategy, with potentially different bidding behavior. Furthermore, these methods all take as input the problem's type distribution, and so (unlike a particular sunk-awareness setting, for example) are potentially appropriate to apply across different environments.

1.5.1 Predictions from Simulated Data

One natural method for generating an initial prediction is to fix a particular strategy profile, and simulate the play of this profile for a large number of games by sampling agent valuations from the underlying type distribution. From this set of simulated games, we can observe the resulting prices, and use these as a basis for prediction. For a point price prediction, we simply compute the average final prices over the simulation experience, and for a distribution-based prediction we compute final price histograms. This yields a large family of prediction-based strategies, each member distinguished by its form of predictor (point vs. distribution), and by the strategy profile employed in simulation to generate the price data.

As a noteworthy special case of the above, our baseline prediction is the distribution of final prices resulting when all agents follow the SB strategy. We denote the baseline point predictor $PP(\boldsymbol{\pi}^{\mathrm{SB}})$ and the baseline distribution predictor $PP(F^{\mathrm{SB}})$.

As noted above, these prediction strategies, and others presented below, take the type distribution as input to the simulation process. Thus, in order to use this method in practice, agents (who know only their own valuation function) need to employ probabilistic beliefs over the valuation functions for other agents. The advantage of this approach is that the methods themselves can be applied to a range of environments by modifying this type-distribution input, without any further parameter tuning required.

1.5.2 Walrasian Equilibrium for Point and Distribution Prediction

An alternative, competitive-analysis approach is to use as predictions the prices that would obtain *if* the market were to reach a *Walrasian price equilibrium* with respect to the m goods and agent valuation functions over those goods. In another complex bidding setting, we have found that predictions based on competitive equilibrium can be surprisingly effective, achieving accuracy comparable to sophisticated machine learning approaches (Wellman, Reeves, Lochner, & Vorobeychik, 2004). We emphasize that our appeal to competitive equilibrium in this context is purely heuristic; we are not assuming as analysts that the equilibrium is realized, but rather employing the well-defined equilibrium concept as a means to generate price predictions for use in bidding.

One immediate complication that underscores the distinction is that Walrasian prices need not actually exist in our setting. Consider once again the $m = n = 2$ configuration of Example 1.3.1 (Table 1.1). Versions of this example, in which one agent views the goods as complements and the other as substitutes, are commonly employed to illustrate the absence of a competitive equilibrium (Cramton, 2005; McAfee & McMillan, 1996). There exist no prices for goods 1 and 2 such that both agents optimize their demands at the specified prices and the markets clear. General conditions for existence of price equilibria given discrete goods are provided by Bikhchandani and Mamer (1997).

To deal with this problem, we specify the Walrasian prediction strategy operationally, just as we do the simulation-based predictors presented in Section 1.5.1. That is, we define the "equilibrium" prices to be those produced after a specified number of iterations of some designated price-adjustment protocol, applied to the environment corresponding to the given SAA game. Although this construction does not guarantee the prices employed are actually in equilibrium (indeed, such guarantee is not possible), it does ensure that the prediction strategy is well-defined.

Another complication is that the standard definition of Walrasian equilibria presumes deterministic demand functions, whereas in our setting we are faced with probability distributions over agent valuations. We can generalize the price-equilibrium calculation in two ways to allow for probabilistic knowledge of the aggregate demand function. The first is to find the *expected price equilibrium* (EPE): the expectation (over the type distribution) of the Walrasian price-equilibrium vector. The most straightforward way to estimate this is Monte Carlo simulation, sampling from the type distribution. A particular sampled type determines the demand function $\boldsymbol{x}$, which we can then employ in a tâtonnement protocol. Let $\boldsymbol{p}^t$ denote the price vector at iteration t, and α^t an adjustment parameter that decays with t. The standard tâtonnement procedure (Arrow & Hahn, 1971) applied to the SAA setting (one unit of each good available) iteratively revises the price vector according to the following difference equation:

$$\boldsymbol{p}^{t+1} = \boldsymbol{p}^t + \alpha^t[\boldsymbol{x}(\boldsymbol{p}^t) - \mathbf{1}]. \tag{1.6}$$

Repeated sampling of types and application of (1.6) yields a crude Monte Carlo estimate of the expected price equilibrium.

An alternative (which may sometimes be preferred for computational reasons) to estimating a price equilibrium in the face of probabilistic demand is the *expected-demand price equilibrium* (EDPE): the Walrasian price equilibrium with respect to expected aggregate demand. In other words, we calculate or estimate the expected demand function and then apply tâtonnement once to find an equilibrium as if realized demand were in fact equal to expected demand. We calculate expected demand analytically when possible (Cheng, Leung, Lochner, O'Malley, Reeves, & Wellman, 2005); otherwise, we can estimate it by Monte Carlo simulation, again sampling from the type distribution.

Either of these generalized Walrasian price-equilibrium methods can be applied to generate point predictions. We denote the expected price-equilibrium point predictor by $PP(\boldsymbol{\pi}^{\mathrm{EPE}})$ and the expected-demand price-equilibrium point predictor by $PP(\boldsymbol{\pi}^{\mathrm{EDPE}})$. The method of expected price equilibrium can also be straightforwardly generalized – by tracking the empirical distribution of price equilibria instead of just average prices – to the case of distribution predictor, yielding $PP(F^{\mathrm{EPE}})$.[9]

1.5.3 Self-Confirming Price Predictions

Our final class of prediction methods combines the spirit of simulation-based and equilibrium-based approaches. The basic idea is to estimate prices under the assumption that agents will follow prediction strategies with accurate price predictions. We refer to these as *self-confirming* predictions. We begin with the simpler case of point predictions.

Definition 1.2 (Self-Confirming Point Price Prediction)**.** *Let* Γ *be an instance of an SAA game. The prediction* $\boldsymbol{\pi}$ *is a* self-confirming prediction for Γ *iff* $\boldsymbol{\pi}$ *is equal to the expectation (over the type distribution) of the final prices when all agents play* $PP(\boldsymbol{\pi})$*.*

In other words, if all agents use a point price-prediction strategy, then the self-confirming predictions are those that on average *are correct* at the end of the auction.[10] We denote the self-confirming prediction vector by $\boldsymbol{\pi}^{\mathrm{SC}}$ and the self-

[9] Unlike the EPE method, which produces a price vector for each sample from the type distribution, the EDPE-method price data is always a single price vector, because tâtonnement is applied only once at the last step. Therefore, we did not construct *distribution* price-prediction strategies based on the latter.

[10] As described above, our price-prediction strategies perform simple updating based on price-quote information as the auction proceeds. Our self-confirmation notion, however, applies only to initial predictions and final prices – we do not insist that the intermediate updated predictions are also confirmed.

confirming point prediction strategy by $PP(\boldsymbol{\pi}^{\mathrm{SC}})$.

The key feature of self-confirming predictions is that agents make decisions based on predictions that turn out to be correct with respect to the type distribution and the assumption that all agents play this particular prediction strategy.[11] Since agents are employing these predictions strategically, we might reasonably expect the strategy to perform well in an environment where its predictions are confirmed.

We next define the concept of a self-confirming *distribution* of final prices in SAA.

Definition 1.3 (Self-Confirming Price Distribution). *Let Γ be an instance of an SAA game. The prediction F is a* self-confirming price distribution for Γ *iff F is the distribution of prices resulting when all agents play bidding strategy $PP(F)$.*

The actual joint distribution will in general have dependencies across prices for different goods. We are also interested in the situation in which if the agents play a strategy based just on marginal distributions, that resulting distribution has the same marginals, despite dependencies.

Definition 1.4 (Self-Confirming Marginal Distribution). *Let Γ be an instance of an SAA game. The prediction $F = (F_1, \ldots, F_m)$ is a vector of* self-confirming marginal price distributions for Γ *iff for all i, F_i is the marginal distribution of prices for good i resulting when all agents play bidding strategy $PP(F)$ in Γ.*

Existence of Self-Confirming Predictions

We demonstrate in Section 1.5.3 that we can often find approximately self-confirming point and distribution predictions. However, we first observe that they do not always exist, for the same reason that Walrasian prices may not exist.

Proposition 1.5. *There exist SAA games for which no self-confirming point price prediction exists, nor do any self-confirming or marginally self-confirming price distributions.*

Proof. Define an SAA game corresponding to the configuration of Table 1.1. Recall the argument in Section 1.5.2 that there are no Walrasian prices for this example. Given a deterministic SAA mechanism (one without asynchrony or random tie-breaking), for fixed value functions the outcome from playing any profile of deterministic trading strategies is a constant. Thus, the only possible self-confirming distributions (which were defined for agents playing the deterministic $PP(F)$ strategies) must assign probability one to the actual resulting prices. But

[11] An equilibrium with this feature is sometimes called a "fulfilled expectations equilibrium" (Novshek & Sonnenschein, 1982).

given such a prediction, our trading strategy will pursue the agent's best bundle at those prices, and must actually get it since the prices are correct if the distribution is indeed self-confirming. But then the markets would all clear, contrary to the fact that the predicted prices cannot constitute an equilibrium, since such prices do not exist in this instance. □

Despite this negative finding, we conjecture that price distributions that are self-confirming to a reasonable degree of approximation exist for a large class of nondegenerate preference distributions, and can be computed given a specification of the preference distribution. For instance, in Example 1.5.1 below we demonstrate that if the preference distribution is such that any particular preference profile is a different variation of Example 1.3.1 rather than a fixed configuration, approximate self-confirming point price predictions may exist even though Walrasian prices do not exist in any game instance. We also show in Section 1.5.3 that approximate marginal self-confirming price distributions may also exist in games with such preferences. We now present a procedure for deriving self-confirming distributions, and some evidence for its effectiveness.

Deriving Self-Confirming Price Predictions

To find approximate self-confirming point predictions, we follow a simple iterative procedure. First, we initialize the predicting agents with some prediction vector (e.g., all zero) and simulate many game instances with the all-predict profile. When average prices obtained by these agents are determined, we replace the initial prediction vector with the average prices and repeat. When this process reaches a fixed point, we have the self-confirming prediction, $\boldsymbol{\pi}^{\mathrm{SC}}$.

Example 1.5.1. *There are $n \geq 2$ agents and $m \geq 2$ goods. Agent 1 has single-unit demand, and the value of each good is $v_1 > 0$. The rest of the agents each need all the goods in order to obtain any value. For $i \neq 1$, agent i's value for the m-good set is v_i, such that*

$$v_1 < v_i < m v_1. \tag{1.7}$$

Let $v_1, \ldots, v_i, \ldots$ be chosen probabilistically from a given type distribution satisfying condition (1.7)*.*

The values in Table 1.1 are an example of preferences satisfying condition (1.7). As for that example, no Walrasian equilibrium prices exist for any combination of v_1 and v_i consistent with the condition. However, because self-confirming predictions are *expectations* over the type distribution, non-existence of equilibrium prices for specific preferences does not imply non-existence of self-confirming prices with respect to the *ex ante* distribution. In Figure 1.1 we show the convergence to a self-confirming price-prediction vector for Example 1.5.1, for a particular distribution and various numbers of agents and goods. In three out of four such SAA games we analyzed, the prices converged within 10 iterations. In the game with two agents and five goods (Panel (c)), there is some persistent oscillation, but the prices stay within 0.5% of the upper bound on a single good value V.

In cases of price oscillation, we found that by resetting the vector of predicted prices to equal the averages around which the prices are fluctuating, the process often immediately converges to a more precise fixed point. We used this method to construct $\boldsymbol{\pi}^{\mathrm{SC}}$ for the SAA game presented in Section 1.6.2.

A similar approach can be applied to derive distribution predictions. Starting from an arbitrary prediction F^0, we run many SAA game instances (sampling from the given preference distributions) with all agents playing strategy $PP(F^0)$.[12] We record the resulting prices from each instance, and designate the sample distribution observed by F^1.[13] We repeat the process using the new distribution F^t for iteration $t+1$ for some further series of iterations. If it ever reaches an approximate fixed point, with $F^t \approx F^{t+1}$ for some t, then we have statistically identified an approximate self-confirming price distribution for this environment.

We employ the Kolmogorov-Smirnov (KS) statistic as one reasonable measure of similarity of probability distributions, defined as the maximal distance between any two corresponding points in the CDFs:

$$KS(F, F') = \max_x |F(x) - F'(x)|.$$

For self-confirming marginal distributions, we take the maximum of the KS distances measured separately for each good: $KS_{\mathrm{marg}} = \max_i KS(F_i, F_i')$.

Specifying our procedure requires (i) a number of samples per iteration, (ii) a threshold on KS or KS_{marg} on which to halt the iterations and return a result, (iii) a maximum number of iterations in case the threshold is not met, and (iv) a smoothing parameter designating a number of iterations to average over when the procedure reaches the maximum iterations without meeting the threshold. The bound on the number of iterations ensures the procedure terminates and returns a price distribution, which may or may not be self-confirming. When this occurs, the smoothing parameter avoids returning a distribution that is known to cause oscillation. We do not, of course, expect the bidding strategy to perform as well

[12] In most of our experiments, the initial prediction is uniformly distributed prices between 0 and the upper bound, V, on the value of a single good, but our results do not appear sensitive to this.

[13] In order to ensure that conditioning on information state during bidding is always well defined, we modify the observed distribution to add uniform infinitesimal probability for all price vectors greater than those observed in samples.

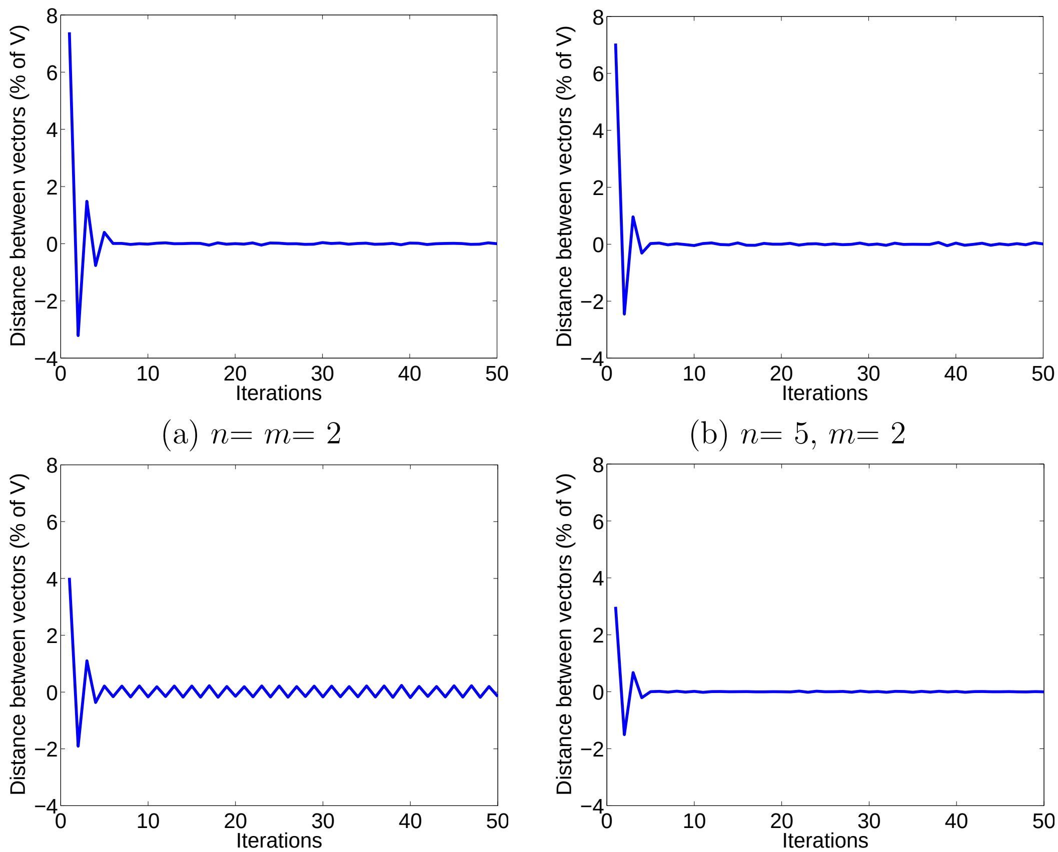

(a) $n= m= 2$ (b) $n= 5,\ m= 2$

(c) $n= 2,\ m= 5$ (d) $n= 5,\ m= 5$

Figure 1.1: Convergence of iterative estimation of self-confirming price-prediction vectors in environments in which no Walrasian equilibrium prices exist. The scenarios are instances of Example 1.5.1, with $v_1 \sim U[3, V-1]$, $V = 50$, and each $v_i \sim U[v_1+1, \min(m(v_1-1), V)]$. The initial prediction is that all prices would be zero, with prices at subsequent iterations determined by a million simulated games at the previous predicted price. The graph plots the distance between the price vectors in consecutive iterations. We define vector distance as the maximum over pointwise distances, measured as a percentage of the upper bound, V, on the value of a single good.

when we cannot find a convergent self-confirming distribution and the underlying oscillations are large.

For our empirical analyses, we specify an SAA game based on a scheduling problem in which there are m units (called *time slots*) of a single schedulable resource, indexed $1, \ldots, m$. Each of n agents has a single job that can be accomplished using the resource. Agent j's job requires λ_j time slots to complete, and by accomplishing this job it obtains some value depending on the time it completes. Specifically, if j acquires λ_j time slots by deadline t, it accrues value $v_j(t)$. Deadline values are nonincreasing: $t < t'$ implies $v_j(t) \geq v_j(t')$.

To illustrate, we consider such a scheduling problem with five agents competing for five time slots. We draw job lengths randomly from $U[1, 5]$. We choose deadline values randomly from $U[1, 50]$ then prune to impose monotonicity (Reeves et al., 2005). The initial prediction is the baseline distribution prediction F^{SB}. The initial prediction is the baseline distribution prediction F^{SB}. We set the algorithm parameters at one million games per iteration, and a 0.01-KS_{marg} convergence criterion. The predicted and empirical distributions quickly converge, with a KS_{marg} distance of 0.007 after only six iterations.

To see if our method produces useful results with some regularity, we applied it to 26 additional instances of the scheduling problem, varying the numbers of agents and goods, and the preference distributions. The initial prediction in all 26 additional instances is that all prices are uniformly distributed. We again drew deadline values from $U[1, 50]$ and pruned them for monotonicity. We used two probability models for job lengths in the first 21 instances. In the *uniform* model, they are drawn from $U[1, m]$. In the *exponential* model job length λ has probability $2^{-\lambda}$, for $\lambda = 1, \ldots, m-1$, and probability $2^{-(m-1)}$ when $\lambda = m$.

We constructed 10 instances of the uniform model, comprising various combination of $3 \leq n \leq 9$ and $3 \leq m \leq 7$. In each case, our procedure found self-confirming marginal price distributions (KS_{marg} threshold 0.01) within 11 iterations. Similarly, for 11 instances of the exponential model, with the number of agents and goods varying over the same range, we found SC distributions within 7 iterations. We plot the distribution of KS_{marg} values from these 21 instances in Figure 1.2.

The 22nd instance was designed to be more challenging: we used the $n = m = 2$ example with fixed preferences described in Table 1.1. Since there exists no SC distribution, our algorithm did not find one, and as expected, after a small number of iterations it began to oscillate among a few states indefinitely. After reaching the limit of 100 iterations, our algorithm returned as its smoothed prediction distribution the average over the last 10.

Finally, we also tested the procedure on the four cases of Example 1.5.1, employed

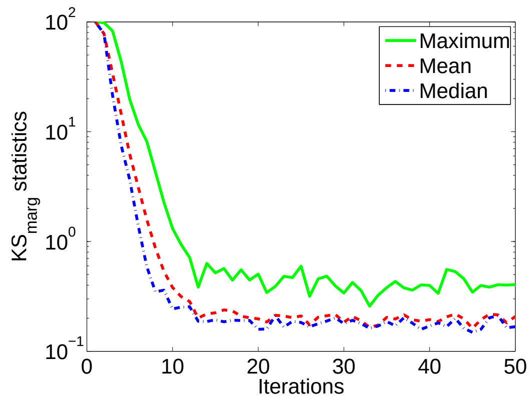

Figure 1.2: Convergence of iterative estimation of self-confirming marginal price distributions. The initial prediction is that all prices are uniformly distributed. The prices at each iteration are determined by a million simulated games. The graph plots the maximum, mean, and median KS_{marg} from 21 instances of the scheduling problem with uniform and exponential preference models.

above (see Figure 1.1) to evaluate search for self-confirming point predictions. Recall that in these environments, no Walrasian equilibrium prices exist in any instance of the preference distribution. For the cases with the number of goods $m = 5$ (both $n = 2$ and $n = 5$), price distributions immediately converged to a self-confirming distribution with the price of the first good equal to one with probability one and the rest of the prices equal to zero. To see that this is self-confirming, note that the agents who need all m goods will calculate their incremental cost by conditioning on the price being positive (see Equation (1.5)). The posterior probability is uniform over all prices (since the condition is greater than any observed prices), which deters all such agents from bidding at all. This is the correct decision from their perspective, as they could not profitably obtain all the goods in the presence of an agent who needs only one and obtains a greater per-good value. For the cases with $m = 2$, in contrast, the uniform belief is not always sufficient to keep these agents from bidding initially, and so the prediction of $p_1 = 1$ and $p_2 = 0$ is not self-confirming. For these examples, we found that the KS_{marg} value during iterative search oscillates in a range of up to 30% of V.

Though we do not expect self-confirming predictors to always provide excellent predictions (they are, after all, heuristics), these examples indicate that even in a variety of challenging environments (in which Walrasian price equilibria do not exist) they often provide reasonable predictions. Of course, the real test of their value as a method for use in bidding comes from their performance against other bidding strategies. We now turn to the computational evaluation of over 50 strategies in various SAA games.

1.6 Empirical Game Analysis: Complementary Preferences

We now analyze the performance of self-confirming price distribution predictors in a variety of SAA games, against a variety of other strategies. We use Monte Carlo simulation to estimate the payoff function for an *empirical game*, which maps profiles of agent strategies to expected payoffs for each agent. This approach converts a game in extensive form to normal form in the expected payoffs. We then analyze equilibria in these normal forms. Our methods extend the approach developed in our prior work (MacKie-Mason et al., 2004; Reeves et al., 2005; Wellman, 2006), and build on ideas from other recent studies in a similar empirical vein (Armantier, Florens, & Richard, 2000; Kephart et al., 1998; Walsh, Das, Tesauro, & Kephart, 2002). We emphasize here that all of the analysis below applies directly to the *estimated* empirical game. These correspond to statistical claims about the actual restricted-strategy game, and lead to arguments generalizing the observations to related games.

1.6.1 Environments and Strategy Space

We studied SAAs applied to market-based scheduling problems, as described in Section 1.5.3. Particular environments are defined by specifying the number m of goods, the number n of agents, and a preference model comprising probability distributions over job lengths and deadline values. The bulk of our computational effort went into an extensive analysis of one particular environment: the $m = n = 5$ uniform model presented above. As described in Section 1.6.2, the empirical game for this setting provides much evidence supporting the unique strategic stability of $PP(F^{SC})$. We complement this most detailed trial with smaller empirical games for a range of other scheduling-based SAA environments. Altogether, we have studied selected environments with uniform, exponential, and fixed distributions for job lengths; a modified uniform distribution for deadline values; and agents in $3 \leq n \leq 8$; goods in $3 \leq m \leq 7$.

To varying degrees, we have analyzed the interacting performance of 53 different strategies. These were drawn from four strategy families described above: SB, 20 sunk-aware agents with varying sunk-awareness parameters k, 13 point predictors, and 19 distribution predictors based on various prediction methods. The price prediction methods include variations of Walrasian equilibrium prediction (nine point and one distribution predictor), historical-data predictions (two point and one distribution), self-confirming predictions (one point and one distribution), and other

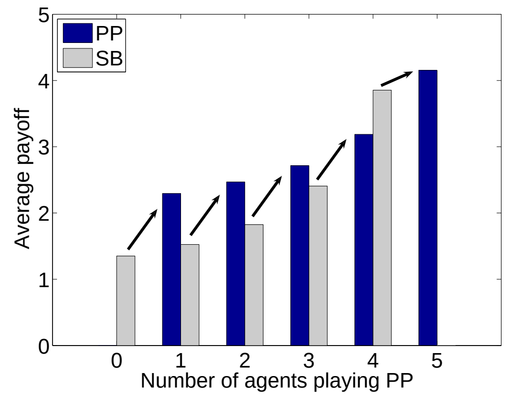

Figure 1.3: Normal-form payoffs for a 5-player game with 2 strategies. The arrows indicate best responses. All-PP is the unique Nash equilibrium.

methods.[14]

The choice of strategies was based on prior experience. We believe that the set includes the best strategy candidates from the prior literature, though we make no claim to have covered all reasonable variations. Naturally, our emphasis is on evaluating the performance of $PP(F^{SC})$ in combination with the other strategies.

Given n agents and S possible strategies, the corresponding symmetric normal-form game comprises $\binom{n+S-1}{n}$ distinct strategy profiles. The game size thus grows exponentially in n and S; for the $n = 5$, $S = 53$ game we estimate below, there are over four million different strategy profiles to evaluate. We first illustrate the process for a simpler game, with five agents, each choosing between SB or the baseline point price-prediction strategy $PP(\boldsymbol{\pi}^{\mathrm{SB}})$ (abbreviated PP). There are six possible profiles which can be described as profiles with j agents playing PP (and the rest SB) for $j = 0, \ldots, 5$. We simulate a large number of games for each profile and average the payoffs for a player of each type (PP, SB). We present the resulting empirical game in Figure 1.3. For this simple game, we can solve the normal form for a unique pure-strategy Nash equilibrium by inspection, illustrated by the arrows. If all five players choose SB, any one can get a higher expected payoff by deviating to PP. If only one plays PP, a second can beneficially deviate to PP. Likewise for each profile except all playing PP, from which none can gain by deviating to SB, establishing a unique Nash equilibrium.

[14] Space considerations preclude a full description of the 53 strategies here. An appendix with specification of all parameters, including complete description of all the prediction methods used for point and distribution predictors, is available at `http://hdl.handle.net/2027.42/57741`.

1.6.2 5×5 Uniform Environment

By far the largest empirical SAA game we have constructed is for the SAA scheduling environment discussed in Section 1.5.3, with five agents, five goods, and uniform distributions over job lengths and deadline values. We estimate payoffs empirically for each profile by running millions of simulations of the auction protocol, so estimating the entire payoff function for over 4.2 million strategy profiles is infeasible. However, we can estimate the payoff matrix for subsets of all profiles, and as we describe below, with well-chosen subsets we can reach useful conclusions about equilibria in the 53-strategy game.

Our results are based on estimated payoffs for 4457 strategy profiles, calculated from an average of 7 million samples per profile (with some profiles simulated for as few as 200 thousand games, and some for as many as 200 million, depending on sampling variances). Despite the sparseness of the estimated payoff function (covering only 0.1% of possible profiles), we have been able to obtain several results.

First, as discussed above, we conjectured that the self-confirming distribution-prediction strategy, $PP(F^{SC})$, would perform well. We have directly verified this: *the profile where all five agents play a pure $PP(F^{SC})$ strategy is a Nash equilibrium of the empirical game.* That is, we verified that no unilateral deviation to any of the other 52 pure strategies is profitable. Note that in order to verify a pure-strategy symmetric equilibrium (all agents playing a strategy s) for n players and S strategies, one needs only S profiles: one for each strategy playing against $n-1$ copies of s. Similarly, to refute the possibility of a particular profile being in Nash equilibrium, we need to find only one profitable deviation profile (i.e., obtained by changing the strategy of one player to earn a higher payoff given the others' strategies).

The fact that $PP(F^{SC})$ is pure symmetric Nash for this game does not of course rule out the existence of other Nash equilibria. Indeed, without evaluating any particular profile, we cannot eliminate the possibility that it represents a (non-symmetric) pure-strategy equilibrium itself. However, the profiles we did estimate provide significant additional evidence, including the elimination of broad classes of potential symmetric mixed equilibria.

Let us define a strategy *clique* as a set of strategies for which we have estimated payoffs for all combinations.[15] Each clique defines a subgame, for which we have complete payoff information. Within our 4457 profiles we have eight maximal cliques that include strategy $PP(F^{SC})$. For each of these subgames, $PP(F^{SC})$ is the only strategy that survives iterated elimination of (strictly) dominated strategies. It follows that $PP(F^{SC})$ is the *unique* (pure- or mixed-strategy) Nash equilibrium in

[15] Thus we have a 2-strategy clique if we have estimated all six profiles that five agents can form from these two strategies.

each of these clique games. We can further conclude that in the full 53-strategy game there are no equilibria with support contained within any of the cliques, other than the special case of the pure-strategy $PP(F^{SC})$ equilibrium.

Analysis of the available two-strategy cliques (not generally maximal) provides further evidence about potential alternative equilibria. Of the $\binom{52}{2} = 1326$ pairs of strategies not including $PP(F^{SC})$, we have all profile combinations for 49. Based on profiles estimated, we have determined that for any symmetric profile defined by a mixture of one of these pairs, an agent can improve its payoff by a minimum of 0.32 through deviating to some other pure strategy. For reference, the average payoff for the all-$PP(F^{SC})$ profile is 4.51, so this represents a nontrivial difference.

That is, none of the two-strategy mixtures for which we have data comes very close to equilibrium, further strengthening our confidence in $PP(F^{SC})$.

Finally, for each of the 4457 evaluated profiles, we can derive a bound on the ϵ rendering the profile itself an ϵ-Nash pure-strategy equilibrium. The three most strategically stable profiles by this measure are:

1. all $PP(F^{SC})$: $\epsilon = 0$ (confirmed Nash equilibrium of the empirical game);
2. one $PP(F^{SB})$, four $PP(F^{SC})$: $\epsilon > 0.13$;
3. two $PP(F^{SB})$, three $PP(F^{SC})$: $\epsilon > 0.19$.

All the remaining profiles have $\epsilon > 0.25$ based on confirmed deviations.

Our conclusion from these observations is that $PP(F^{SC})$ is a highly stable strategy within this strategic environment, and likely uniquely so. Of course, only limited inference can be drawn from even an extensive analysis of only one particular distribution of preferences, so we now consider other environments.

1.6.3 Self-Confirming Prediction in Other Environments

To test whether the strong performance of $PP(F^{SC})$ generalizes across other SAA games, we undertook smaller versions of this analysis on variations of the model above. We explored 17 additional instances of the market-based scheduling problem: eight with the uniform (U), eight with the exponential (E) preference models (3–8 agents, 3–7 goods), and one with fixed preferences, corresponding to the counterexample model of Table 1.1. For each we derived self-confirming price distributions (failing in the last case, of course), as reported in Section 1.5.3. We also derived price vectors and distributions for the other prediction-based strategies. We ran between two and ten million games per profile in all of these environments.

For the non-symmetric game with fixed preferences, we evaluated all 53 profiles with at least one agent playing $PP(F^{SC})$.

For eleven of the symmetric games (eight U and three E models), we started by evaluating 27 profiles: one with all $PP(F^{SC})$, and for each of 26 other strategies s, one profile with $n-1$ $PP(F^{SC})$ and one s. In eight of these games, $PP(F^{SC})$ and $PP(F^{SB})$ were among top three unilateral deviations from $PP(F^{SC})$ in the all-$PP(F^{SC})$ profile. For each of the eleven games, we identified five (additional) top-ranking deviations from $PP(F^{SC})$ and evaluated complete 7-cliques involving these five strategies, $PP(F^{SC})$ and $PP(F^{SB})$ in the respective environments (at least 340,000 samples per profile).

For the five additional E models, we evaluated all profiles over seven selected strategies.[16]

Our results for U and E models are summarized in Table 1.2. For each case, we report the ϵ that, for the estimated payoff matrix, renders all-$PP(F^{SC})$ an ϵ-Nash equilibrium. The next two columns report sensitivity information about this figure, given its basis in payoffs estimated from samples. First, since our payoff matrix is estimated (and thus each payoff has a sampling variance), we calculate the expected value $\bar{\epsilon}$ of ϵ with respect to the empirical distributions of the estimated payoffs (assuming that the errors in our payoff estimates are independent, and using the sample variances as population variances). Thus, for example, the environment $E(3,5)$ has a pure Nash equilibrium of all-$PP(F^{SC})$ for the estimated payoff matrix, but taking into account sampling variation, on average that profile has an ϵ of 0.005.

Under the same independence assumption, "$\Pr(\epsilon = 0)$" represents the probability that all-$PP(F^{SC})$ is actually an equilibrium. Finally, for each empirical game with $n \leq 6$ we also obtained a symmetric mixed-strategy Nash equilibrium using replicator dynamics.[17] The rightmost column reports the probability of playing $PP(F^{SC})$ in the resulting mixture, to evaluate its significance when it does not constitute a pure-strategy equilibrium.

In 14 out of these 16 environments, $PP(F^{SC})$ was verified to be an ϵ-Nash equilibrium for $\epsilon < 0.1$. Twelve have $\epsilon < 0.05$, and in six of these (one U and five E) it was an exact equilibrium. The two worst environments were $U(5,3)$ and $U(7,8)$. In the last case, expected payoff for all-$PP(F^{SC})$ was 2.67, so ϵ represents

[16]For these models we did not incur the additional computational cost of evaluating all 27 profiles to select best deviations from $PP(F^{SC})$, which is a somewhat arbitrary procedure for selecting strategies for a clique in any case. Rather, we selected the seven candidate strategies based on regularities in the results from the other eleven games described above.

[17] By replicator dynamics we mean an iterative (evolutionary) algorithm for finding symmetric mixed-strategy equilibria in symmetric games. Our implementation is based on the replicator dynamics formalism introduced by Taylor and Jonker (1978) and Schuster and Sigmund (1983) and is described in detail in our earlier work (Reeves et al., 2005). Though the method is not guaranteed to generate all Nash equilibria, we have found it particularly useful for finding sample Nash equilibria.

Env(m,n)	ϵ-gain from one-player deviation	$\bar{\epsilon}$-gain adjusted for sampling error	$\Pr(\epsilon=0)$: Probability of exact Nash equilibrium	Probability of play in rep. dyn. solution
$E(3,3)$	0	0	1.00	1.00
$E(3,5)$	0	.005	.600	.996
$E(3,8)$	.031	.032	0	—
$E(5,3)$	0	0	1.00	.999
$E(5,5)$	0	.001	.900	.998
$E(5,8)$	.029	.031	0	—
$E(7,3)$	0	.007	.667	.992
$E(7,6)$	.003	.007	.567	.549
$U(3,3)$	.097	.099	0	.725
$U(3,5)$	0	0	1.00	1.00
$U(3,8)$	.017	.016	0	—
$U(5,3)$	.103	.103	0	.809
$U(5,8)$	.047	.048	0	—
$U(7,3)$	.058	.060	0	.942
$U(7,6)$	.018	.018	0	.929
$U(7,8)$	.133	.132	0	—

Table 1.2: Evaluations of all-$PP(F^{SC})$ profile for U and E models.

about 5% of the value. For no other case did it reach 2%. Moreover, the results are quite insensitive to statistical variation. The $\bar{\epsilon}$ values never exceed ϵ by much, and in every environment for which we produced an equilibrium with replicator dynamics, $PP(F^{SC})$ appears in this symmetric mixed-strategy profile with substantial if not overwhelming probability.

Overall, we regard this as favorable evidence for the $PP(F^{SC})$ strategy across the range of market-based scheduling environments. Not surprisingly, the environment with fixed preferences is an entirely different story. Recall that in this case the iterative procedure failed to find a self-confirming price distribution. The distribution it settled on was quite inaccurate, and the trading strategy based on this performed poorly – generally obtaining negative payoffs regardless of other strategies. Since one of the available strategies simply does not trade, $PP(F^{SC})$ is clearly not a best-response player in this environment.

1.7 Strategies for Environments with Substitutes

In the previous sections we focused on the exposure problem when there are complementarities in preferences. We found that strategies based on price prediction can be quite effective in mitigating the problem. In this section we extend our analysis of bidding strategies to the case of substitutable goods. The strategic challenge in this environment is bidding when there are significant own price effects: bidding below willingness-to-pay for the marginal unit may lower the price sufficiently on inframarginal units to be a profitable strategy (Ausubel & Cramton, 2002). We now expand the space of bidding strategies we evaluate to include simple demand-reduction strategies as well as a sophisticated approach to predicting own price effects inspired by the success of self-confirming price prediction for environments with complementarities. In the environment with substitutes we study, we find that the simple demand-reduction strategies clearly outperform this price predictor.

To analyze bidding strategies in an SAA game with substitutes, we assume that each auction sells one unit of a homogeneous indivisible good, and the bidders' marginal value for units of this good is weakly decreasing. We implemented such preferences by randomly drawing marginal values v_k for the kth good from $U[0, v_{k-1}]$, with $v_0 = V$ a uniform upper bound on the marginal value of one unit.

In homogeneous-good environments bidders derive the same value from any bundle of q goods regardless of their labels. The definitions of strategies in this section rely on this assumption, though it would not be difficult to generalize their approaches to apply to environments with a more general type of substitutability. The assumption of homogeneous goods is convenient for computational implementation and analysis, however, we believe that it is not essential to our main results.

1.7.1 Demand-Reduction Strategy

Consider an SAA game with m auctions, each selling one unit of an identical (homogeneous) good. If all agents follow SB, the outcome is that the bidders for the m most highly valued units win them, at a uniform price equal to the value of the most highly valued losing unit (possibly plus the bid increment). This is virtually equivalent to truth-telling in an $m+1$st sealed-bid uniform-price auction. Like the truth-telling/sealed-bid case, the all-play-SB outcome is efficient (modulo the bid increment), but it is not an equilibrium. In fact, efficient equilibria in the $m+1$st sealed-bid uniform-price auction do not exist (Ausubel & Cramton, 2002). To motivate a possibly better strategy, consider the intuition for the non-existence of an efficient equilibrium: if a bidder has a positive probability of influencing price in a situation in which the bidder wins a positive quantity, then the bidder has an incentive to shade her bid in a sealed-bid uniform-price auction. Bid-shading leads to inefficient outcomes. This intuition and the failure of SB motivates considering strategies that suppress demand.[18]

We introduce a relatively simple demand-reduction strategy, DR. Let us modify SB by introducing a parameter $\kappa \in [0, V]$ defining the degree of the agent's demand reduction. An agent playing strategy $DR(\kappa)$ bids the ask price on the lth cheapest good as long as it is not winning that good, and its marginal surplus is at least $\kappa(l-1)$. In other words, the agent considers the goods in order of price, adding the lth good to its bundle until the marginal value v_l drops below the ask price plus $\kappa(l-1)$. The DR strategy family is a simple way of capturing the intuitions of the demand-reduction literature: bidders should shade their bids, and the amount of shading increases with the number of winning goods (Ausubel & Cramton, 2002).

Formally, define $DR(\kappa)$'s perceived price of the good with the lth lowest myopically perceived price (defined in Section 1.3):

$$\rho_l(\boldsymbol{B}) \equiv \begin{cases} \beta_l + \kappa(l-1) & \text{if winning the good} \\ \beta_l + 1 + \kappa(l-1) & \text{otherwise,} \end{cases} \tag{1.8}$$

where $\boldsymbol{\beta}$ is the vector of current bid prices. Agent $DR(\kappa)$ plays the perceived-price bidding strategy using this $\rho(\boldsymbol{B})$. Note that $\rho(\boldsymbol{B})$ as defined by (1.8) assumes that

[18] Note that the sunk-awareness modification of SB we introduced in Section 1.3.2 to address the exposure problem leads to overbidding, as opposed to bid-shading, in this environment. Using the terminology of Definition 1.1, the perceived-price vector of a sunk-aware strategy is equal to or below the myopic perceived-price vector used by SB, which results in more aggressive bidding. The perceived price of the demand-reduction strategy we introduce in this section is always at least as high as the myopic perceived-price vector.

the goods are indistinguishable. We use the subscript l instead of i to emphasize that each good is labeled by its myopic price rank order rather than by the auction selling it.

1.7.2 Predicting Own Price Effects

The ability of a single agent to affect final prices is strategically central when goods are substitutes. Therefore, the focus of price prediction in the substitutes case is to model this relationship. Specifically, for the homogeneous-good environment, price predictions take the form of a mapping from purchase sizes (i.e., the agent's chosen demand) to final prices. The main role of this prediction is to guide the agent as to when it is beneficial to refrain from bidding on potentially valuable goods.

The assumption that final prices depend on the number of goods the agent is trying to win implies that the agent's prediction of the final price of good i can no longer be represented by a scalar. Let $\pi_{iq}(\boldsymbol{B})$ be the predicted final price of good i given that the agent tries to win q goods and its information state at the current round is $\boldsymbol{B}$. We can think of the agent's *predicted own-effect prices* as an $m \times m$ matrix, in which the rows are auction labels and the columns are the intended purchase sizes. We define an updating rule for π_{iq}, $i, q \in \{1, \ldots, m\}$, similar to the point price-prediction rule described in Section 1.4.1. The current price prediction for good i when the agent plans to bid on q goods is the maximum of the initial prediction and the myopically perceived price:

$$\pi_{iq}(\boldsymbol{B}) \equiv \begin{cases} \max(\pi_{iq}(\varnothing), \beta_i) & \text{if winning good } i \\ \max(\pi_{iq}(\varnothing), \beta_i + 1) & \text{otherwise.} \end{cases} \tag{1.9}$$

There is no apparent reason why an agent should believe that the final price of a homogeneous good on one auction will be higher than the price on another auction. Therefore, we construct the initial price prediction to be equal across auctions: $\pi_{iq}(\varnothing) = \pi_{jq}(\varnothing)$ for all i and j for all purchase sizes q. In other words, the elements in a column are identical in the agent's initial prediction matrix. We label the initial prediction matrix of predicted own-effect prices by π^x, in which the subscript x labels particular initial predictions.

In the homogeneous-good environment, agents are indifferent between item subsets of equal sizes. Thus, in our strategy, the agent uses price prediction to determine the number q^* of units to buy, but not to identify specific auctions in which to participate in the current round. Formally,

$$q^* = \arg\max_q \max_{|Y|=q} \sigma(Y, \boldsymbol{\pi}_{\cdot q}(\boldsymbol{B})),$$

where $\sigma(Y, \boldsymbol{p})$ is the agent's surplus for goods Y defined by (1.1), and $|Y|$ refers to the number of goods in set Y.

Given q^*, the choice of goods X^* on which to actually bid is based on the current myopically perceived prices, $\rho(\boldsymbol{B})$ as defined by (1.2). Using myopically perceived prices ensures that the agent never regrets the composition of its bid set (conditional on size) even if its predicted own-effect prices are wrong.

$$X^* = \arg \max_{|X|=q^*} \sigma(X, \rho(\boldsymbol{B}))$$

The agent breaks ties as in Definition 1.1. Given X^*, the agent bids $b_i = \beta_i + 1$ (the ask price) for the $i \in X^*$ that it is not already winning. We call this strategy family the *own-effect price predictor* ($OEPP$) and denote a specific strategy in this family by $OEPP(\boldsymbol{\pi}^x)$.

Similar to the point price predictor defined for complementary goods, the $OEPP$ family includes SB as a special case when the predicted own-effect prices are a matrix of zeros: SB $= OEPP(0)$. As mentioned in Section 1.7.1, if all players follow SB, the allocation is efficient. Perceived prices based on an own-effect price matrix with positive elements are weakly higher than the myopic perceived prices SB uses. Therefore, an $OEPP$ agent using positive predictions tends to bid on fewer items than is efficient given the others' bids, and never bids on more goods than SB would.

1.7.3 Self-Confirming Own-Effect Prices

We define the concept of self-confirming own-effect price prediction similarly to self-confirming point price prediction for complementary environments.

Definition 1.6 (Self-Confirming Own-Effect Prices)**.** *Let Γ be an instance of an SAA game with homogeneous goods. Matrix $\boldsymbol{\pi}$ is a* self-confirming own-effect price matrix for Γ*, if for all $i, q \in \{1, \ldots, m\}$, $\pi_{iq}(\varnothing)$ is equal to the expectation (with respect to the type distribution) of the final price when one agent tries to win q goods and all the other agents follow $OEPP(\boldsymbol{\pi})$.*

In other words, self-confirming own-effect prices satisfy the condition that if one of the agents bids to win q goods and the other agents "exploit" their own-effect price predictions, that prediction on average *is correct* for all q. We denote the self-confirming own-effect price matrix by $\boldsymbol{\pi}^{SC}$ and the self-confirming own-effect price-prediction strategy by $OEPP(\boldsymbol{\pi}^{SC})$.

To find approximate self-confirming own-effect prices, we follow an iterative procedure similar to that described in Section 1.5.3. First, we initialize the own-effect predictors with some own-effect price matrix (e.g., all zero) and, sampling

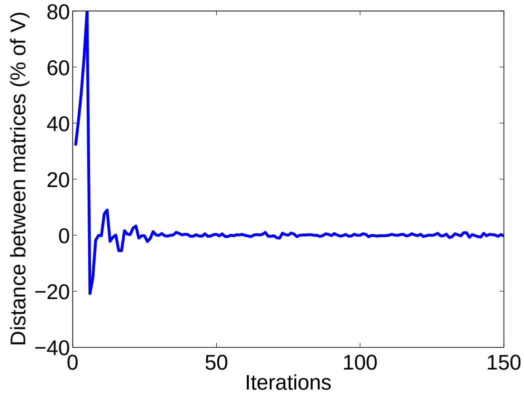

Figure 1.4: Convergence to a self-confirming own-effect price matrix, starting with an initial prediction that all prices would be zero regardless of the size of the agent's purchase. The prices at each iteration are determined by 10 thousand simulated games. The graph plots the distance between the own-effect prices in consecutive iterations. We define distance between matrices as the maximum over pointwise distances, measured as a percentage of the upper bound on the marginal value, V, of a single unit of the good. The bound V equals 127 in all of our SAA games with substitutes.

from the homogeneous-good type distribution, run many SAA game instances with a profile in which one agent (the *explorer*) ignores its preferences and tries to win a single good, while the others follow $OEPP$. When average prices obtained by these agents are determined, we replace the first column in the own-effect price matrix with a column vector with all elements equal to the average price, reset the explorer to win two goods and repeat. After the second batch of simulations, we replace all elements in the second column of the own-effect matrix with the average price and increase the explorer's target number of goods by one. We repeat the process, recycling back to a single unit after the exploration target reaches m. When this process reaches a fixed point, we have the matrix of self-confirming own-effect prices, $\boldsymbol{\pi}^{SC}$. We have not investigated whether a fixed point necessarily exists in homogeneous-good environments, but the price predictions converged in this environment within 30 iterations in all of our experiments (see Figure 1.4).

1.7.4 Empirical Game Analysis

We perform analyses similar to, but less extensive than, those reported in Section 1.6. We analyzed the $m = n = 5$ environment with uniform preferences introduced at the beginning of Section 1.7. We set the upper bound V to 127. In Figure 1.5 we display the agents' average valuations as a function of the number

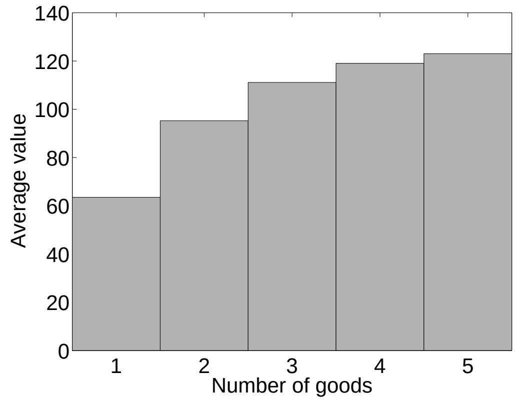

Figure 1.5: Preference distribution in the homogeneous-good environment.

of goods. As before, our goal is to evaluate the performance of a self-confirming price-prediction strategy, $OEPP(\boldsymbol{\pi}^{SC})$ in this instance. Since the literature predicts that agents suppress demand in equilibrium, we include many instances of our demand-reduction strategy family. We analyzed 51 strategies: SB, 47 $DR(\kappa)$ with $1 \leq \kappa \leq 120$, one sunk-aware strategy with parameter $k = 0.5$, a self-confirming own-effect price predictor $OEPP(\boldsymbol{\pi}^{SC})$, and the baseline distribution predictor $PP(F^{\mathrm{SB}})$ (defined in Section 1.5).

We estimated payoffs for 16542 strategy profiles (out of 3.48 million possible), based on an average of 986 thousand samples per profile. Some profiles are simulated for as few as 40 thousand samples; near-Nash-equilibrium profiles were simulated for up to 205 million game instances per profile. Despite the high-quality information $OEPP(\boldsymbol{\pi}^{SC})$ employs about own effect on final prices, the strategy's use of this information did not provide any advantage over the simpler information-free demand-reduction agents. In the majority of profile settings where it was tested, $OEPP(\boldsymbol{\pi}^{SC})$ can be refuted with a $DR(\kappa)$ strategy. Indeed, for 96% of the 16542 profiles analyzed, we found the best deviation in our data set to be an instance of $DR(\kappa)$ with $10 \leq \kappa \leq 22$. For 108 profiles our data set includes estimated payoffs for all deviations from all strategies. The best deviations for these profiles are always an instance of $DR(\kappa)$ with $16 \leq \kappa \leq 19$.

We provide more evidence in Figure 1.6 by displaying the number of times a strategy was a best deviation (dark bars) relative to the number of estimated profiles in which that strategy appeared (light bars). The latter is proportional to the approximate number of opportunities for that strategy to be a best deviation from some other profile. The dark bars reflect the preponderance of situations in which agents prefer moving toward a $DR(\kappa)$ strategy with κ near 15. The light bars document our decision, as this evidence was emerging, to focus our finite

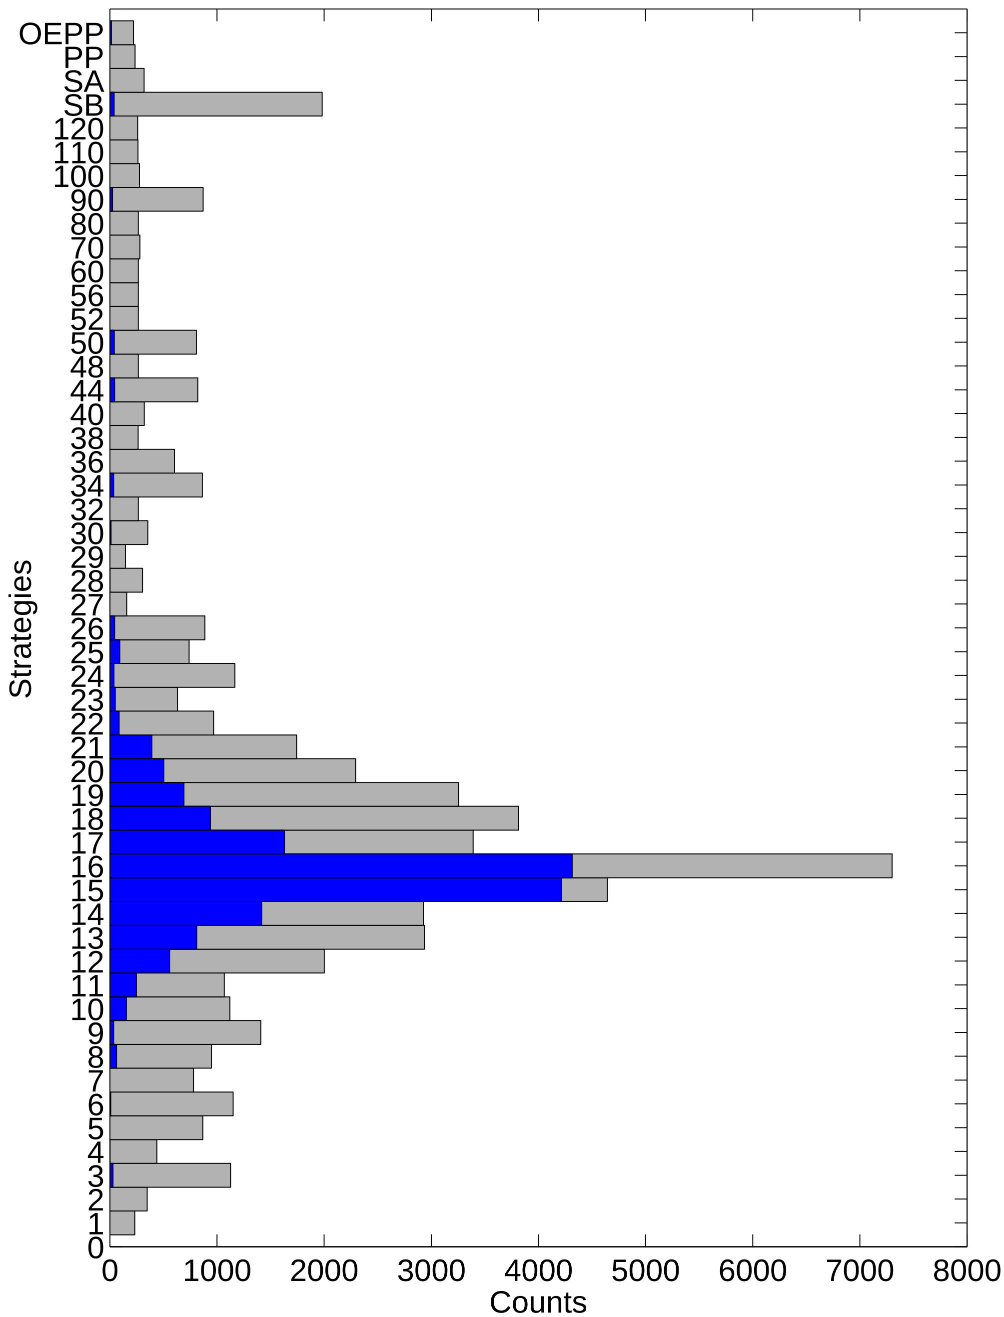

Figure 1.6: Distribution of best deviations. The light bars reflect the number of estimated profiles in which the corresponding strategy appeared. The dark bars reflect how many times the strategy in fact *was* a best deviation. We index demand-reduction strategies $DR(\kappa)$ by their corresponding κ-values. $OEPP$ refers to $OEPP(\boldsymbol{\pi}^{SC})$, PP to $PP(F^{\mathrm{SB}})$, and SA refers to the sunk-aware strategy with $k = 0.5$.

computational resources on estimating regions of the payoff matrix most important for (near-)equilibrium play.

We found only 14 profiles for which the highest gain can be obtained by deviating to $OEPP(\boldsymbol{\pi}^{SC})$. This is 0.085% of all estimated profiles and 6.36% of all profiles containing at least one $OEPP(\boldsymbol{\pi}^{SC})$ player. We found 40 profiles for which SB is the best deviation. The sunk-aware and $PP(F^{\mathrm{SB}})$ strategies are never the most attractive deviations in our data.

We found many pure-strategy asymmetric ϵ-Nash equilibria in this environment. Those with the lowest ϵ are profiles of $DR(\kappa)$ with $14 \leq \kappa \leq 17$. To give a sense of the magnitude of demand (bid) suppression, these κs correspond to 33–40% of the average final unit price if all players follow SB. In Table 1.3 we present all ϵ-Nash equilibria for which $\epsilon \leq 0.015$[19] and two of our benchmark profiles: all-SB and all-$OEPP(\boldsymbol{\pi}^{SC})$ (for which the ϵ is rather large). The probability that the profile is an exact Nash equilibrium was estimated empirically as described in Section 1.6.3. The profiles are listed in the order of increasing ϵ. We have estimated payoffs of all unilateral deviations from the strategies in the near-Nash-equilibrium profiles to all of the other 50 pure strategies. These ϵ-equilibria all consist of $DR(\kappa)$ with κs in a narrow range; the best deviations are to nearby κs (column 2). If all agents follow $OEPP(\boldsymbol{\pi}^{SC})$, a single agent can improve her payoff by at least 2.86 (5.5% of the average payoff) by deviating to $DR(24)$.

As expected, equilibrium outcomes are inefficient in this environment. However, the efficiency loss is small: all-16, the symmetric profile with the smallest ϵ, achieves 98.55% efficiency. We present efficiency results for a few symmetric near-Nash-equilibrium profiles and our benchmark profiles in Table 1.4.

Our results suggest that $OEPP(\boldsymbol{\pi}^{SC})$ is a weak competitor against $DR(\kappa)$. The weakness of $OEPP(\boldsymbol{\pi}^{SC})$ may lie in its failure to adjust its bidding to its opponents' behavior: having good information does not guarantee strategic advantage. We observe that $OEPP(\boldsymbol{\pi}^{SC})$ bids like an aggressive demand-reduction agent. As a consequence, it earns high profits when playing against other predictors: essentially, in a profile of all-$OEPP$, players are tacitly colluding to reduce demand and thus prices. Payoffs would be higher if all agents could commit to this behavior. However, when collusion is unenforceable, the usual motive to deviate unilaterally is strong.

[19]For reference, the payoffs range from 30 to 69 in our empirical payoff matrix. Thus, the near-equilibrium profiles in Table 1.3 are quite close to equilibria: the ϵ of 0.015 constitutes at most 0.05% of the payoff.

ϵ-Nash-equilibrium profile	Best deviation	ϵ-gain from one-player deviation	$\bar{\epsilon}$-gain adjusted for sampling error	Probability the profile is exact Nash equilibrium
15 15 16 16 16	15 → 16	0	0.001	0.58
all-16	16 → 15	0.001	0.004	0.25
15 16 16 16 16	16 → 15	0.001	0.006	0.14
15 15 15 16 16	15 → 16	0.004	0.005	0.11
14 16 16 16 16	16 → 15	0.004	0.008	0.02
15 15 15 15 16	15 → 16	0.005	0.009	0
15 16 16 16 17	15 → 16	0.006	0.007	0.11
14 14 15 15 16	14 → 15	0.006	0.008	0.02
14 14 14 16 16	14 → 15	0.007	0.008	0.09
14 15 15 16 16	14 → 16	0.008	0.009	0.02
15 15 17 17 17	17 → 16	0.008	0.012	0
14 14 15 15 15	14 → 15	0.009	0.008	0.07
15 15 15 17 17	17 → 15	0.009	0.010	0.02
14 14 14 15 15	14 → 15	0.010	0.010	0.05
14 15 15 15 16	16 → 15	0.011	0.010	0.02
16 16 16 16 17	17 → 15	0.011	0.010	0.02
all-15	15 → 16	0.012	0.012	0.04
15 15 16 17 17	17 → 16	0.012	0.012	0.01
15 16 16 17 17	17 → 16	0.012	0.013	0
14 14 16 16 16	14 → 15	0.012	0.013	0
15 17 17 17 17	15 → 16	0.012	0.014	0
14 14 14 15 16	16 → 15	0.013	0.015	0
14 15 16 16 16	14 → 15	0.013	0.014	0.01
15 15 16 16 17	17 → 16	0.013	0.013	0
14 14 15 16 16	14 → 15	0.014	0.014	0
all-17	17 → 16	0.014	0.015	0
15 16 17 17 17	15 → 16	0.015	0.015	0
16 17 17 17 17	17 → 16	0.015	0.015	0
all-SB	SB → 14	1.450	1.469	0
all-$OEPP$	$OEPP$ → 24	2.857	2.905	0

Table 1.3: ϵ-Nash equilibria for the substitutes environment. The profiles are listed in order of increasing ϵ.

ϵ-Nash-equilibrium profile	Best deviation	ϵ-gain from one-player deviation	Average payoff	Efficiency (%)
all-SB	SB $\to$ 14	1.450	34.266	100
all-14	14 $\to$ 15	0.020	44.665	98.82
all-15	15 $\to$ 16	0.012	45.230	98.69
all-16	16 $\to$ 15	0.001	45.773	98.55
all-17	17 $\to$ 16	0.014	46.307	98.40
all-18	18 $\to$ 17	0.035	46.810	98.26
all-$OEPP$	$OEPP \to 24$	2.857	52.063	93.75

Table 1.4: Efficiency of some symmetric ϵ-Nash equilibria in the substitutes environment. The profiles are listed in order of decreasing efficiency.

1.8 Discussion

Our investigation of bidding strategies for simultaneous auctions leads to qualitatively different conclusions for environments characterized by complementary and substitutable preferences. For the case of complements, we find strong support for a bidding strategy based on probabilistic price prediction, with *self-confirming predictions* derived through an equilibration process. Like other decision-theoretic approaches to bidding (Greenwald & Boyan, 2004), this strategy tackles the exposure problem head-on, by explicitly weighing the risks and benefits of placing bids on alternative bundles, or no bundle at all. The fact that the predictions are self-confirming suggests that this cost-benefit analysis will be accurate when other agents are following the same strategy.

Given the analytic and computational intractability of the SAA game, we evaluated our self-confirming probabilistic price-prediction strategy, $PP(F^{SC})$, using an empirical game-theoretic methodology. We explored a restricted strategy space including $PP(F^{SC})$ along with a range of candidate strategies identified in prior work. Despite the infeasibility of exhaustively exploring the profile spaces, our analyses support several game-theoretic conclusions. The results provide favorable evidence for our new strategy – very strong evidence in one environment we investigated intensely, and somewhat less categorical evidence for a range of variant environments.

For the case of substitutes, the driving strategic issue is demand reduction rather than exposure risk, and thus it is necessary to predict own price effects as well as exogenous price levels. We defined a bidding strategy, $OEPP$, based on such predictions, and a concept of self-confirming prices analogous to the approach that proved so successful in complementary environments. In this domain, however, the strategy $OEPP(\boldsymbol{\pi}^{SC})$ based on explicit self-confirming predictions did not fare well, proving in our empirical experiments significantly inferior to an approach based on simple across-the-board demand reduction.

There are several possible explanations for the relative lack of success of explicit price prediction in substitutes environments. One is that the particular $OEPP$ method we investigated measures own price effects under unrealistic assumptions. Specifically, the strategy predicts the effect of selecting a demand level (number of goods to go for), and sticking with that choice thereafter. In actuality, the agent can and does reconsider its choice at each round conditional on the current auction information. This myopic assumption about the agent's own behavior would tend to overestimate the effect of its immediate decision about demand at the current prices, and thus cause it to reduce demand more aggressively than warranted.

The simple demand-reduction strategy, $DR(\kappa)$, can pursue an appropriate degree of demand reduction in a particular environment by tuning the free parameter κ. This approach was successful in our experimental environment, but would presumably need to be retuned for a different configuration of goods and preferences. It remains for future work to identify a general approach for deriving robust demand-reduction strategies directly from specification of preference distributions.

Returning to environments with complementarities, our results establish the self-confirming price-prediction strategy as the leading contender for dealing broadly with the exposure problem. If agents make optimal decisions with respect to prices that turn out to be right, there may not be room for performing a lot better. On the other hand, there are certainly areas where improvement should be possible, for example:

- incorporating price dependencies (but with reasonable computational effort);
- more graceful handling of instances when self-confirming price distributions do not exist;
- more sophisticated prediction updates given price quotes, including possible incorporation of history; and,
- timing of bids: trading off the risk of premature quiescence with the cost of pushing prices up.

Dealing with combinations of complementarity and substitutability, by combining considerations of exposure and demand reduction, is perhaps the most obvious direction for extending the scope of bidding-strategy ideas developed here.

Finally, an indirect contribution of this work is to demonstrate an empirical methodology for game-theoretic analysis when strategy determination is analytically intractable (MacKie-Mason & Wellman, 2005; Wellman, 2006). We find that even when strategy spaces are enormous, much can be learned by empirically converting an extensive-form game into a normal form in expected payoffs for strategy choices, combined with thoughtful selection of payoff-matrix regions to estimate, and carefully targeted analyses of results.

CHAPTER 2

Bundling Information Goods: A Study of Competing Firms Facing Heterogeneous Consumers

2.1 Introduction

Information goods such as books, news stories, scholarly articles, music recordings, movies, computer games, and software are characterized by high fixed (first-copy) costs, but low costs for the production of additional copies. This cost structure is greatly exaggerated for digital information goods. King and Tenopir (2005) observe that the maturation and integration of communication technologies and the economics of the journal system, particularly pricing of traditional journal subscriptions and access to digital full-text databases through site licensing and packages have the potential either of destroying the scholarly journal system or substantially enhancing its considerable usefulness and value. They argue that "the new technologies should, if deployed with care, enhance the journal system (. . .), but contemporary pricing policies have been a greater threat to the journal system."

One of the challenges for information-good producers is that linear marginal-cost pricing for electronic information goods cannot result in efficient production or distribution because near-zero prices would not recover the initial fixed costs (Varian, 1995). On the other hand, a price above the marginal cost creates a deadweight loss and therefore is inefficient. However, the flexibility of digital technology permits a wide range of responses to the cost problem. For example, by selling individual items as well as bundles such as journals, CDs, and software packages, firms can potentially attract more consumers than with either pricing scheme.[1] Offering both schemes is much more feasible for digital goods available online than for information goods produced and distributed on physical media.

[1] This is assuming consumers' tastes are heterogeneous enough.

In this chapter I study how product configuration flexibility introduced by digital technology may affect competing producers of information goods and consumers. An important contribution of this work, for the analysis of both monopoly and competitive bundling, is that I study consumers with heterogeneous preferences.

Several authors have provided valuable insights about bundling under a simplifying assumption that consumers are characterized by a single variable and are ex ante homogeneous (e.g., Schmalensee, 1984; Zahray & Sirbu, 1990; Bakos & Brynjolfsson, 1999). However, the assumption of such preferences is quite restrictive: for example, it implies that a monopolist will maximize social welfare but will extract all available consumers' surplus (Bakos & Brynjolfsson, 1999). In this chapter, I study environments with high consumer diversity. Following Chuang and Sirbu (1998), I assume two-dimensional consumer preferences, which capture not only the consumers' reservation price for a set of goods, but also correlations between individual valuations within the set. The latter is particularly important in the context of mixed bundling.

For tractability, I restrict my attention to three types of product configurations, or bundling strategies: selling items individually (at the same per-item price), pure bundling (the whole collection of items is sold together), and mixed bundling (consumers are offered a choice between the two previous options). With N goods, there are 2^N-1 different bundles that could be offered, each at a potentially unique price, and consumer preferences would need to be specified on this (2^N-1)-dimensional space. The profit-maximization problem for setting prices is NP-hard and generally considered to be computationally intractable for modern computers when N is only moderately large.[2] In this study, the collection size of a single firm varies from 50 to 150. Complex pricing schemes would also create intractable consumer information-processing problems and impractical transactions costs.

There are several reasons why a firm may choose to bundle its goods. Previous authors provide at least five reasons for selling two or more products in a single package:

1. Cost savings in production and transactions associated with package selling in presence of economies of scale (e.g., Coase, 1960; Bakos & Brynjolfsson, 2000; Demsetz, 1968; Chuang & Sirbu, 1998);

2. Strategic bundling such as *tying* for leveraging market power (e.g., Carbajo, Meza, & Seidmann, 1990; Whinston, 1990);

[2] My empirical analysis also shows that a monopoly does not benefit from splitting its collections into two and applying independent bundling schemes to them. This holds for all of the preference distributions I have analyzed. This part of the analysis is not reported here, but is available upon request.

3. Complementarity (superadditivity) in consumption of bundling components (e.g., Adams & Yellen, 1976; Bakos & Brynjolfsson, 1999; Economides & Viard, 2004);

4. Second-degree price discrimination when marginal valuations are additive or subadditive (e.g., Adams & Yellen, 1976; McAfee, McMillan, & Whinston, 1989; Varian, 2000, 1989; Bakos & Brynjolfsson, 1999; Chuang & Sirbu, 1998); and

5. Reducing buyer diversity (via aggregation) (e.g., Schmalensee, 1984; Varian, 1989; Salinger, 1995; Bakos & Brynjolfsson, 1999, 2000).

The first point concerns savings on the supply side, whereas the last three represent bundling as a tool for extracting consumer surplus. The literature under the second category studies how a firm with monopoly power in one market can use its leverage provided by this power to forclose sales in, and thereby monopolize, a second market. Producers of information goods may choose to bundle for any of these reasons. In the paper-based production of scholarly articles, for example, packaging articles into journals and journals into subscriptions saves non-negligible reproduction and distribution costs. Software is often pre-installed on personal computers, which leverages the market power of the software producer. Software packages typically contain complementary programs. Producers of business news sub-bundle general and time-sensitive information – such as stock quotations or financial news – and delay the latter for those who are not willing to pay a premium. This is an example of price discrimination: the scheme induces casual readers and readers who use the information for business purposes to self-select into appropriate consumption groups. The aggregation effect often manifests itself in bundles of music recordings (CDs), articles (journals or magazines), TV shows (TV channels), etc. The consumers' willingness to pay for such bundles varies less than their willingness to pay for individual components. In this chapter, however, I am concerned only with two of all possible reasons for bundling: price discrimination and reducing buyer diversity. Also, I study the use of bundling schemes when firms *compete* in the selling of information goods. Thus, I address the interaction between market structure and bundled pricing strategies.[3]

In particular, I consider two firms, each producing differentiated collections of information goods. I assume that they have zero reproduction and distribution costs and face a population of consumers with heterogeneous tastes. I model heterogeneous tastes as a valuation function that depends on two parameters. One parameter is the consumer value of her most favored item on the market, and the other is the

[3]I restrict my attention to competition in the market for a single differentiated information good. Analysis of bundling on multiple linked markets for strategic purposes, which falls under the second category, is outside of the scope of this chapter.

percentage of the items on the market that she values positively. My preference model is a generalization of Chuang and Sirbu's monopoly model (Chuang & Sirbu, 1998) to two firms producing differentiated collections of items.[4] The firms can choose among three bundling schemes: they can offer individual items, each at the same price (pure unbundling), sell their collection as a single bundle (pure bundling), or offer both options at the same time (mixed bundling).

Within this framework, I address several questions. First, can simple mixed bundling effectively sort consumers under competition? How do the equilibrium profits compare to those when the firms are restricted to pure bundling or pure unbundling? How do the pure forms compare to each other? Under monopoly, mixed bundling must perform at least as well as either pure unbundling or pure bundling since it contains both pure strategies as sub-cases: maximizing over a larger set of options may not result in a lower profit. When two firms compete, however, the equilibrium payoffs in a game in which both firms have unrestricted strategy sets can be lower than the equilibrium payoffs in a game in which at least one firm's strategy set is restricted to a smaller set of options. In a price-competition game like the one studied here, a restricted set of pricing schemes may alleviate pricing wars and therefore result in higher profits for the firms.

In cases in which mixed bundling is more profitable in equilibrium than the pure schemes, it is important to know by how much. Pure bundling and pure unbundling are appealing due to their simplicity, which may be an important advantage in practical implementation. My second goal is therefore to quantify the gains from choosing the more complex pricing scheme of mixed bundling.

My third question is how much does competition curb the ability of bundling firms (in pure or mixed forms) to extract value from consumers? Although digital information goods have negligible marginal costs of reproduction and distribution, there may be significant first-copy costs. If competition significantly reduces operating profits, investment in new information products may be low.

Finally, what is the effect of information-good bundling under competition on social welfare, taking into account both consumers' surplus and profits? How do mixed-bundling equilibrium profits, social welfare, and consumer surplus compare to those when the firms are restricted to one of the pure schemes?

Due to the complexity induced by heterogeneous, multi-dimensional consumer preferences over multiple goods sold in various configurations, I have not been able to solve analytically for equilibrium of the pricing game. As in Chapter 1, I make

[4] My production model, however, is a special case of the Chuang-Sirbu model: Chuang and Sirbu assume a positive marginal cost and different levels of economies of scale, while I consider the extreme case of zero marginal cost and, consequently, no economies of scale.

use of the *empirical game-theoretic methodology* to analyze competitive bundling in a number of specific *environments*.

I find that depending on the consumer preference distribution, any of the three schemes can result in the highest equilibrium profits when two firms compete. While under monopoly mixed bundling is always weakly better due to its ability to sort consumers into different consumption groups, under duopoly it also leads to more aggressive price competition. However, mixed bundling is generally the most profitable scheme, and in the cases when a pure form yields higher profits in equilibrium, the difference from the mixed-bundling equilibrium profits is small. Under monopoly, a pure pricing scheme gives 80–100% of the mixed-bundling profits; under duopoly, at least 75–105%, depending on the preference distribution and the scheme employed by the other firm.

I also find that under competition, the profits of firms that bundle – in pure or mixed form – are lower by up to 21% relative to a monopoly employing the same (combination of) pricing schemes as the duopoly. The market efficiency is up to 16% higher, and the distribution of social welfare shifts toward consumers by up to 22%. These estimates are based on three different distributions of consumer preferences, each analyzed for a monopoly, a symmetric duopoly and a non-symmetric duopoly.

2.2 Related Work

In their seminal paper on bundling, Adams and Yellen (1976) consider a monopolist producing two goods. They analyze three pricing strategies: component selling (each good priced and sold individually), bundling (both goods sold together), and mixed bundling (consumers offered a choice between buying a bundle or individual components). They show by example that bundling can result in higher revenue than component selling and argue that the profitability of a bundle of different commodities can "stem from its ability to sort customers into groups with different reservation price characteristics, and hence to extract consumer surplus." However, bundling might result in inefficient consumption of components for which the marginal cost of producing that component exceeds its value to a consumer. In general, the profitability of bundling strategies depends on the prevailing level of marginal costs and on the distribution of customer valuations.

McAfee et al. (1989) extend the analysis to general demand functions, and provide conditions under which mixed bundling strictly dominates either of the two pure strategies. Salinger (1995) shows that if valuations for the two goods are not perfectly correlated, the demand for the bundle will be more price-elastic than the sum of the individual demands for each component. As a result, bundling is a way to smooth out idiosyncratic preferences enabling a monopolist to extract more consumers' surplus than is possible through pure components pricing.

Some more recent papers have extended the analysis to more than two goods. Hanson and Martin (1990) tackle the full 2^N bundle pricing problem with N=21. MacKie-Mason and Riveros (1998), and Hitt and Chen (2005) introduce a new alternative: generalized subscriptions. Consumers prepay for G tokens, then select G items from the entire collection after the items are created. In effect, individual-specific sub-bundles are created. Riveros (1999) comparatively evaluates generalized subscriptions against the three most studied strategies.

Brooks, Fay, Das, MacKie-Mason, Kephart, and Durfee (1999), Kephart, Das, Brooks, Durfee, Gazzale, and MacKie-Mason (2001), and Brooks, Gazzale, Das, Kephart, MacKie-Mason, and Durfee (2002) employ machine-learning techniques to explore the space of pricing policies. They are concerned with the trade-off between exploitation and exploration of pricing policies in static and dynamic environments. They show that simple pricing policies with few parameters to learn are more robust in both environment types. Kephart, Das, and MacKie-Mason (2000) acknowledge that information is an experience good. They allow consumers to learn their valuations and the seller to learn pricing parameters simultaneously. They find that dynamic market interactions when there is substantial uncertainty can lead to pathological outcomes if agents are designed with reasonable but not sufficiently adaptive strategies.

Bakos and Brynjolfsson (1998) illustrate that bundling with N goods is strictly preferred to component pricing when marginal cost is zero, consumer item preferences are homogeneous and identically distributed, and N is sufficiently large.

Chuang and Sirbu (1998) introduce two-dimensional heterogeneity of tastes and allow for the possibility of mixed bundling. I adopt these features of their model to study bundling under competition. As Chuang and Sirbu point out, by employing a single variable to model consumer heterogeneity, one can only capture consumers' *aggregate* valuations for a bundle. This is adequate in the pure-bundling context. In the mixed-bundling context, however, it is important to account for the correlation of values across different items as well. For example, mixed bundling may be more profitable relative to pure bundling when one consumer values positively only two of a hundred items in a collection (low correlation) and the other has the same total collection value more or less evenly distributed across fifty of the items (high correlation).

There are two important differences between the assumptions in my work and those in Chuang and Sirbu's study: Chuang and Sirbu assume a positive marginal cost and a single producer. At the center of their study is the question of how the economies of scale and marginal cost affect the relative profitability of pure unbundling, pure bundling, and mixed bundling, with the goal to demonstrate that scholarly journal publishers have incentives to unbundle their journals – i.e., to provide individual articles as well as journals – given how technology has been changing the costs. With a positive marginal cost, it is not surprising that they find that for a monopolist, mixed bundling is more profitable than either of the pure strategies. They also allow readers to value some articles at less than marginal production cost. Thus it is possible for component selling to strictly dominate bundling, since bundling results in the costly distribution of products that have below-cost value to some consumers. In this chapter, I set the marginal cost to zero to study the effects of price discrimination and buyer-diversity reduction of different bundling strategies under competition.

A number of authors have studied bundling in competitive settings. Fishburn, Odlyzko, and Siders (2000) consider a duopoly in which each firm produces an identical set of N information goods (perfect substitutes). By assumption, one firm bundles while the other offers component pricing. In nearly all of their numerical simulations, a price war ensues with both firms' prices falling towards zero (marginal cost).

Matutes and Regibeau (1992) and Farrell, Monroe, and Saloner (1998) consider duopolists that produce complements rather than substitutes. These authors study bundling of two complementary goods that may be purchased from separate firms. Nalebuff (2000) finds that a firm that sells a bundle of complementary products will have a substantial competitive advantage over rivals who sell the component products individually. Some authors examine bundling as a tying strategy when one firm is a monopolist over one product but faces potential competition for a second product (see, e.g., Carbajo et al., 1990; Whinston, 1990; Aron & Wildman, 1999). Nalebuff (1999) explores how bundling can be used as an entry deterrent.

Bakos and Brynjolfsson (1999) consider a duopoly in an N-good market and (stochastically) identical consumers. Goods are pair-wise substitutes: the demand for one good is independent of the demand for all but one of the remaining goods. This drastically restricts the possible strategic interactions. In this chapter, all item valuations are (weakly) subadditive, and thus all items depress one another's value.

There is an extensive body of literature on product differentiation under linear pricing. Goods are assumed to have a fixed and a positive marginal cost of production, and the firms choose their market niche as well as the price of their differentiated good. The linear-city model due to Hotelling (1929), the circular-city model due to Salop (1979), and a symmetric model of monopolistic competition due to Dixit and Stiglitz (1977) and Spence (1976) are examples of classic models of horizontal differentiation. In none of these studies, however, is the model of consumer tastes rich enough to study the questions I pose in this chapter. Hotelling (1929) and Salop (1979) characterize consumers by a single type variable, and consumer valuations of goods are assumed to be independently and identically distributed. As discussed above, one-dimensional heterogeneity of tastes is not adequate in the context of mixed bundling. In the Dixit-Stiglitz-Spence model, there exists a single, representative consumer with a constant-elasticity utility function that depends on the consumption levels of items produced by different sectors (firms). That is, there is no heterogeneity in tastes whatsoever. The rationale for product differentiation stems from the consumer's broad spectrum of interest rather than heterogeneity in the consumer population.

For this chapter, I build directly on the work of Fay and MacKie-Mason (1999) and Fay (2001), Chapter 2. Fay and MacKie-Mason show that under homogeneous consumer preferences, bundling achieves the first-best solution to a firm's profit-maximization problem and unbundling yields less profit than the first-best solution. Under heterogeneous preferences, they find that introducing competition from a second firm results in much lower prices than under monopoly, yet only a moderate profit reduction. However, the latter result may be attributed to the specific properties of the value function used in the analysis of environments with heterogeneous consumers. The value function used to model heterogeneity of consumer tastes had an undesirable property: the value of consuming only firm 2's bundle is positively correlated with the number of items in firm 1's bundle, even if the consumer does not have access to that bundle. The valuation function I use in this study does not have this property.[5] In addition, Fay and MacKie-Mason did not consider mixed bundling and were unable to find any mixed-strategy Bertrand equilibria of the pricing game in which both firms use pure bundling. I overcome these difficulties by using more sophisticated game-solving techniques (see Section 2.4).

2.3 Bundle-Pricing Game

2.3.1 Firms

Two firms each control a collection of items, sized $N_1 > 0$ and $N_2 > 0$, respectively (N_1 and N_2 are fixed and exogenous). An example of such a collection would be a set of different book titles, a collection of different news stories or scholarly articles, a collection of music recordings or movies or software programs. A book title, a news

[5]The observation about the undesirable property of the valuation function studied by Fay and MacKie-Mason (1999) is due to Scott A. Fay. Fay also proposed the valuation function presented in this chapter.

story, or a song would be an example of a single item. There may be substantial sunk costs to create the goods,[6] but the marginal cost of reproducing and distributing is zero. Although in practice marginal cost is rarely strictly zero, the simplifying zero-marginal-cost assumption allows one to explore what happens as marginal costs become vanishingly small, as is typically the case for digital information goods.

I restrict each firms' bundling choices to three *bundling strategies*: selling items individually at the same per-item price, to which I refer as *pure unbundling* (U), *pure bundling* (B) of the whole collection for a single bundle price, and *mixed bundling* (M) that combines the two options.[7] I exclude sub-bundling strategies for reasons of computational tractability. Even for this simplified problem, solving for a mixed-strategy equilibrium sometimes took days, depending on the consumer preference distribution. I analyze this market as a duopoly game in which the two firms simultaneously choose their bundling and pricing strategies. Thus, the pricing strategy of firm i, $i \in \{1, 2\}$, is a per-item price $p_i \geq 0$ under pure unbundling, a bundle price $P_i \geq 0$ under pure bundling, and a pair of prices (p_i, P_i) under mixed bundling.[8]

2.3.2 Consumers

In this section, I introduce a model of consumer preferences over information goods and define the consumer's choice problem. For the purpose of this study, different book titles are instances of the same information good (books), different software programs are instances of the information good "software", etc. Typically, consumers buy only one copy of a book, software, song, or a news story. I therefore make the following assumption.

Assumption 2.1. *Consumers have demand for at most one unit of a particular item.*

Consumers vary in the quantity of the information good in which they are interested. For example, some consumers may spend all day viewing videos on

[6]We can think of first-copy costs as sunk, which allows to treat such costs as exogenous variables in the model. In other words, my model does not encompass the decision to produce the goods.

[7]It is important to distinguish between a mixed-bundling strategy and a *mixed strategy*, where the latter is a term commonly used in game-theoretic literature. Here the mixed-bundling strategy (M) is defined as a pair of prices: a bundle price and a per-item price. This strategy can be *pure* in the sense that a firm can "play" this pair with probability one. For example, a firm can offer a choice between buying the whole collection at \$10 or any subset of items at \$1 per item. A mixed strategy is a probability distribution over a subset of pure strategies. For example, in half of its stores, the firm can offer the whole collection for \$7 and each individual item for \$2 instead of the previous combination. Then the firm's strategy can be viewed as a mixed strategy of two mixed-bundling strategies, (10,1) and (7,2), each chosen with probability $\frac{1}{2}$ at each store.

[8]Note that mixed bundling subsumes both pure schemes: setting p_i to infinity is equivalent to pure bundling, and setting P_i to infinity is equivalent to pure unbundling.

`youtube.com`, and some view only a couple. On the book market, consumers vary by the number of books they read per year. Consumers also differ in the intensity of their preferences. The preference intensity is high if the consumer has particularly high values over some items available on the market. For example, fans of the interactive web-based video series *lonelygirl15*[9] are an example of consumers with high-intensity preferences over videos. Academic achievement of biology professors depends on their knowledge of the latest biology research, and therefore they have high-intensity preferences over scholarly articles.

Following Chuang and Sirbu (1998), let us assume that the consumers' valuations of individual items in a collection are correlated. The correlation in a collection is defined by two parameters: intensity and breadth. The difference from the Chuang-Sirbu model here is that the consumer can choose from two collections rather than a single one. Let $w_i > 0$, $i \in \{1, 2\}$, be the value of the consumer's most favored item in collection i. This variable describes the intensity, or depth, of the consumer's tastes for collection i. Let $k_i > \frac{1}{N_i}$ represent the preference breadth of a given consumer for collection i. Roughly, k_i can be thought of as the fraction of the items that the consumer is willing to consume from collection i. In fact, this is an accurate interpretation of the parameter if $k_i \leq 1$. If k_i is greater than one, the concept of such a fraction is not well defined. The meaning of k in that case will become clear later. I impose another simplifying assumption on w and k.

Assumption 2.2. *For each consumer, her preference breadth and intensity are the same across different collections when collections are considered independently, i.e.,* $w_1 = w_2$ *and* $k_1 = k_2$.

I will therefore suppress the collection subscript and denote breadth simply by k and intensity by w. Equal breadth implies, for example, that a rock fan likes the same number of songs in two different music collections of the same size. In a collection of twice that size, she will like twice as many items. Equal intensity implies that the consumer is equally happy with her top choices from either collection. It is straightforward to extend the analysis to cases in which $w_1 \neq w_2$ and $k_1 \neq k_2$, but this is outside the scope of this work.

Consumers may place zero value on any number of items. Suppose each consumer has ranked all items by their individual valuations. I label the individual valuation of an item x by $v(x)$. Consider $n_1 < N_1$ highest-ranking items from collection 1 and $n_2 < N_2$ highest-ranking items from collection 2.

Assumption 2.3. *A consumer's preferences are represented by a quasilinear utility function* $U(n_0, n_1, n_2) = u(n_0) + V(n_1, n_2)$, *where* n_0 *is the numéraire and* $V(n_1, n_2)$

[9] `http://www.youtube.com/user/lonelygirl15`

is a subutility function *that depends on n_1 and n_2, where n_1 and n_2 are as defined above.*

I approximate quantities n_1 and n_2 with continuous variables. Consider first the following subutility function V':

$$\begin{aligned} V'(n_1, n_2) = \max_{m_1,m_2} & \, w\Big(m_1(1 - \frac{m_1}{2kN_1}) + m_2(1 - \frac{m_2}{2kN_2})\Big), \\ \text{s.t.} \quad & m_1 \geq 0, \quad m_2 \geq 0, \\ & m_1 \leq n_1, \quad m_2 \leq n_2. \end{aligned} \tag{2.1}$$

I will refer to the function that is maximized as f':

$$f'(x_1, x_2) = w\Big(x_1(1 - \frac{x_1}{2kN_1}) + x_2(1 - \frac{x_2}{2kN_2})\Big). \tag{2.2}$$

The maximization part in Equation (2.1) is to ensure that the property of free disposal is satisfied. If the consumption pair (n_1,n_2) is such that the constraints are binding – which is true for (n_1,n_2) where f' is increasing – function V' is simply f'. Otherwise, V' is set to the value achieved at the satiation point, i.e., at the point where f' is maximized.

I now show that V' is a straightforward generalization of the Chuang-Sirbu model to two collections. To obtain the item valuation function, we take the derivative of f' with respect to x_i, $i \in \{1, 2\}$:

$$\frac{\partial f'(x_i, x_j)}{\partial x_i} = w(1 - \frac{x_i}{kN_i}). \tag{2.3}$$

This is a downward-sloping straight line. Remember that the free-disposal property implies that item valuations cannot be negative. Thus, the valuation $W(n_i)$ of the consumer's n_ith highest-ranking item from collection i becomes

$$W(n_i) = \max\{0, w(1 - \frac{n_i}{kN_i})\}, \tag{2.4}$$

where $0 \leq n_i < N_i$.

This is exactly the Chuang-Sirbu valuation function for a single collection. Figure 2.1 displays $W(n_i)$ for $N_i = 100$, $w = 100$, and two different values of k: $k = 0.5$ and $k = 2$. It also clarifies the interpretation of k when it is greater than one: k defines the slope of the valuation function ($\frac{-w}{kN_i}$).

Another assumption concerns the consumer's behavior when she is indifferent between consuming or not consuming an item. Define marginal surplus of an item as marginal value minus marginal cost of obtaining the item.

Assumption 2.4. *Consumers do not consume items with zero marginal surplus.*

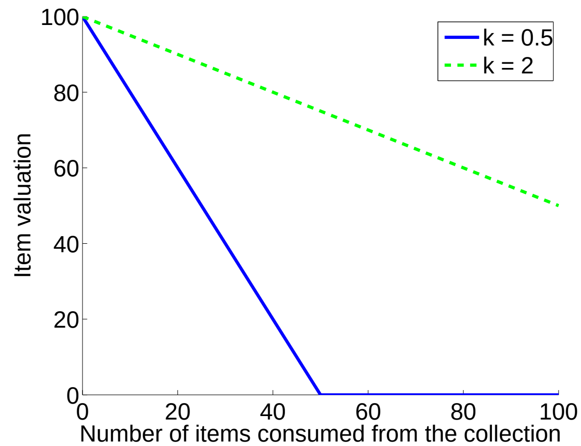

Figure 2.1: Valuation function $W(n_i)$ for $N_i = 100$, $w = 100$, and two different values of k: $k = 0.5$ and $k = 2$. On the horizontal axis, the items are in the decreasing order of preference. The slopes of the lines are $\frac{-w}{kN_i} = -2$ and -0.5, respectively. The 0.5-k consumer would consume $kN_i = 50$ out of 100 items. The cutoff number 50 is the intersection of the line with the horizontal axis. The rest of the items have zero value. The 2-k line intersects the horizontal axis at $n_i = 200 > N_i$. This consumer consumes all 100 items and has a strictly positive value of the least preferred item.

We can see from Equation (2.4) that the consumer's valuations of items from one collection are not related in any way to her valuations from the other collection. Therefore, with such preferences, we are still looking at the Chuang-Sirbu world, in which the interaction between each firm and consumers can be considered separately: they do not have to compete for consumers. In order to make the market more competitive, I introduce a parameter $\gamma > 0$ and modify the subutility function as follows:

$$
\begin{aligned}
V(n_1, n_2) = \max_{m_1, m_2} & \, w\Big(m_1(1 - \frac{m_1}{2kN_1}) + m_2(1 - \frac{m_2}{2kN_2})\Big) - \gamma(m_1 + m_2)^2, \\
\text{s.t.} \quad & m_1 \geq 0, \quad m_2 \geq 0, \\
& m_1 \leq n_1, \quad m_2 \leq n_2.
\end{aligned}
\tag{2.5}
$$

I use V as the consumer's subutility function to study competition. I will refer to the function being maximized as f:

$$
f(x_1, x_2) = w\Big(x_1(1 - \frac{x_1}{2kN_1}) + x_2(1 - \frac{x_2}{2kN_2})\Big) - \gamma(x_1 + x_2)^2. \tag{2.6}
$$

Note that the consumer value of some N items does not depend on which firm or how many firms the items belong to. This property is desirable when seeking to compare different market models. Suppose firm 1 owns a fraction α of the N items, and firm 2 owns the rest of them. Consider a consumer's n highest-ranking items. To simplify exposition, suppose that she values them all positively. Given Assumption 2.2, these n items will correspond to αn and $(1-\alpha)n$ highest-ranking items in the two collections, respectively. The consumers' subutility then equals

$$\begin{aligned}
&f(\alpha n,(1-\alpha)n) = \\
&w\Big(\alpha n(1-\frac{\alpha n}{2k\alpha N}) + (1-\alpha)n(1-\frac{(1-\alpha)n}{2k(1-\alpha)N})\Big) - \gamma(\alpha n + (1-\alpha)n)^2 = \\
&w\Big(n(1-\frac{n}{2kN})\Big) - \gamma n^2,
\end{aligned} \tag{2.7}$$

which is the consumer's subutility of n highest-ranking items from a single collection of size N. The more general case when some of the n items have zero value to the consumer can be proved similarly: those items have no effect on the utility.

With $\gamma > 0$, the consumer's valuations of positively valued items from different collections become *strongly subadditive*, as opposed to additive. Goods are subadditive if the combined set is worth less than the sum of its parts. This links the firms' markets in the following way: the price a consumer is willing to pay for an item depends on the number of items consumed from the other collection, which in turn depends on the other firms prices. In other words, firm i's prices affect the demand for firm j's goods.

To see that the parameter γ introduces subadditivity, it is easier to work with f. Note that the cross-derivative of f is negative for positive γ:

$$\frac{\partial^2 f(x_1,x_2)}{\partial x_1 \partial x_2} = -2\gamma. \tag{2.8}$$

This implies that for positively valued items, the marginal valuation of the n_ith preferred item from collection i decreases as the number of items n_j consumed from collection j increases, $i,j \in \{1,2\}$. Decreasing marginal valuation is equivalent to subadditivity (Lehmann et al., 2006).

It is important to mention that γ also introduces subadditivity between items from the *same* collection. To see that, let us calculate the marginal valuation function of the consumer's n_ith highest-ranking item from collection i for a given number of items n_j from collection j ($i,j \in \{1,2\}$):

$$MV(n_i) = \max\{0, w(1-\frac{n_i}{kN_i}) - 2\gamma(n_i+n_j)\}, \tag{2.9}$$

where $0 \le n_i < N_i$, and n_j is fixed at some number between 0 and N_j.

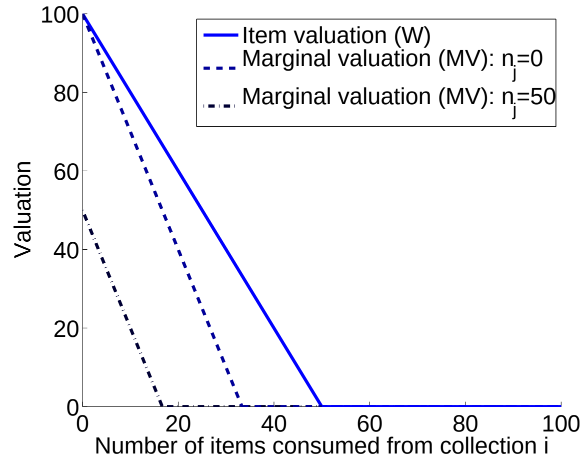

Figure 2.2: Valuation function $W(n_i)$ and marginal-valuation function $MV(n_i)$ for $N_i = 100$, $w = 100$, $\gamma = 0.5$, and two different values of n_j: $n_j = 0$ and $n_j = 50$. On the horizontal axis, the items are in the decreasing order of preference. The slope of the W line is $\frac{-w}{kN_i} = -2$. The slopes of the MV lines are $\frac{-w}{kN_i} - 2\gamma = -3$. MV for $n_j = 50$ is shifted down by $2\gamma n_j = 50$ units.

Figure 2.2 displays $W(n_i)$ and $MV(n_i)$ for $N_i = 100$, $w = 100$, $k = 0.5$, and $\gamma = 0.5$ at two different values of n_j: $n_j = 0$ and $n_j = 50$. Note how subadditivity between items from the same collection reduces the slope of the marginal-valuation line (by 2γ) and how subadditivity between items from different collections shifts the marginal-valuation line for collection i down by $2\gamma n_j$.

By assumption, the term $\frac{-w}{kN_i}$ is the slope of the item valuation function W, which describes the valuation of each item given no other items have been consumed. If $\gamma = 0$, the marginal-valuation function equals W. If $\gamma > 0$, the slope becomes steeper. That is, the marginal value of consuming an item is now below its valuation, and it is decreasing.[10] Therefore, the valuations of items from the same collections are subadditive if $\gamma > 0$. One consequence of this γ-effect is that the firms face a less elastic demand relative to the Chuang-Sirbu model, given the same k and w.

Let us now consider the consumer's choice problem when she faces two competing firms. Given the firms' pricing schemes and particular prices, consumers choose n_1 and n_2 to maximize their surplus (i.e., value minus the cost of acquiring the items). Remember that the pricing strategy of firm $i \in \{1, 2\}$, is a per-item price $p_i \geq 0$ under pure unbundling, a bundle price $P_i \geq 0$ under pure bundling, and a pair of prices

[10] Implicitly in this argument, I rely on Assumption 2.5, which I introduce at the end of this section. Given that the valuation function V and MV are defined for preference-ordered items, the conclusion does not directly follow if the items in the domains of MV and W have different orders.

$(p_i,\ P_i)$ under mixed bundling. In the most general case when the firms each offer a mixed bundle, each consumer has the following options: she can buy two bundles; buy only one bundle from one firm and zero or more individual items from the other; buy no bundles at all and zero or more individual articles from both firms. Let d_i, $i \in \{1, 2\}$, equal 1 if the consumer buys firm i's bundle, and equal 0 otherwise. Then the consumer's surplus (CS) of consuming n_1 and n_2 highest-ranking items from collections 1 and 2, respectively, is given by

$$CS(n_1, n_2) = V(n_1, n_2) - (1 - d_1)p_1n_1 - d_1P_1 - (1 - d_2)p_2n_2 - d_2P_2. \tag{2.10}$$

To find the consumer's optimal choice, we need to maximize CS subject to a number of constraints. Before stating the maximization problem, let me note that V is maximized whenever $f(n_1, n_2)$ is maximized, by definition of V (see Equations (2.5)–(2.6)). Other V-maximizing choices contain items with zero marginal values in addition to the f-maximizing set. According to Assumption 2.4, consumers do not consume items with zero marginal surplus. Therefore, they will never consume items with zero marginal value. Then the number of consumed items n_1 and n_2 will always be below or at the satiation point, and the consumer optimization problem can be written as follows:

$$\begin{gathered}
\max_{n_1,n_2} \Big(f(n_1, n_2) - (1 - d_1)p_1n_1 - d_1P_1 - (1 - d_2)p_2n_2 - d_2P_2 \Big), \\
\text{s.t.} \quad n_1 \geq 0, \quad n_2 \geq 0, \\
n_1 < N_1, \quad n_2 < N_2.
\end{gathered} \tag{2.11}$$

We need an additional assumption to make the consumer's choice well defined. Recall that the domain of the subutility function V is a rank-ordered set of items, with the items ranked according to their individual valuations $v(x)$. Suppose for a moment that we are back to one collection on the market. Given that the per-item price is the same for all items, consumers never have to consider item subsets other than those defined by the number of top-ranking items, when buying items individually. For example, if $v(x) > v(y) > v(z)$ for a particular consumer, then this consumer never has to consider the set $\{x, z\}$ to make the optimal choice, because it would never be optimal to exclude item y when choosing to consume z. Similarly, if a consumer buys a full collection as a bundle, the marginal cost of consuming each item is zero, and therefore the above argument applies. When there are two collections on the market, an additional assumption is required to ensure that a consumer will never want to consume a set of n items from collection i that is not the first n top-ranked items from i.

To see that, consider the following example. Suppose there are two collections (collection 1 and 2), each consisting of one economics article (E_i), one computer science article (CS_i), and one political science article (PS_i). Suppose also that the consumer is generally more interested in economics than in computer science and more in computer science than in political science. Let us extend the definition of $v(X)$ to be the consumer's valuation of *any* article subset X. For concreteness, suppose $v(E_1) = 30$, $v(E_2) = 25$, $v(CS_1) = 20$, $v(CS_2) = 15$, $v(PS_1) = 10$, $v(PS_2) = 5$. Moreover, the consumer prefers to read broadly, but does not have time to read more than one article in each category. More specifically, articles from the same fields have zero marginal values (subadditive preferences), while articles from different fields are additive: $v(\{E_1, E_2\}) = 30$, $v(\{CS_1, CS_2\}) = 20$, $v(\{PS_1, PS_2\}) = 10$; let i, j, h can be 1 or 2, then $v(\{E_i, CS_j\}) = v(E_i) + v(CS_j)$, and similarly for all other pairs of articles from different fields; and finally, $v(\{E_i, CS_j, PS_h\}) = v(E_i) + v(CS_j) + v(PS_h)$. Suppose now that the consumer is considering buying CS_1. Then the marginal value of CS_2 is zero, and the consumer may want to consider the subset $\{E_2, PS_2\}$ from collection 2. However, the subutility function V is not defined for this subset. It is only defined for the following subsets from collection 2: $\{E_2\}$, $\{E_2, CS_2\}$, and $\{E_2, CS_2, PS_2\}$. I therefore impose the following assumption on the marginal valuations.

Let $v(A)$ be the consumer's value of *any* set of items A. Let $mv(x|A)$ be the consumer's marginal value of item x in the set $\{x \cup A\}$.

Assumption 2.5. *Suppose $v(x) > v(y)$, where x and y are two single items. Then for any set A, $mv(x|A) > mv(y|A)$.*

This assumption is sufficient to ensure that no matter what items a consumer chooses to consume from one collection, the order of items from the other collection ranked by marginal valuations would be the same as the order by v. This in turn ensures that the consumer's choice problem is always well defined, even though V is undefined for some subsets of items.

2.4 Empirical Game Analysis

Recall that the purpose of this work is to analyze competition of bundling firms when consumers have heterogeneous tastes. Heterogeneity of the consumer population is defined by the distribution of two parameters, w and k introduced in Section 2.3.2. I have not been able to solve analytically for Nash equilibrium strategies in this game for the case of general distributions of w and k. As an alternative approach, I employ the *empirical game-theoretic* methodology discussed in Chapter 1. In particular, for each environment described in the following section, I use computer simulation to estimate the demand for a variety of the firms' strategy profiles. I explain the simulation procedure in Section 2.4.1. Given the demand, I calculate the firms' expected profits for each strategy profile. Then I solve for (mixed) Nash equilibria of the *restricted* competition game.[11]

To study relative performance in equilibrium of mixed bundling, pure bundling and pure unbundling, in addition to the unconstrained mixed-bundling game, I analyze all possible subgames in which at least one firm is restricted to either B or U. I label them as X_1X_2, where $X_1, X_2 \in \{\text{B, U, M}\}$ denote the bundling-strategy space to which firm 1 and firm 2, respectively, are restricted. There are eight such subgames: BB, BU, UB, BM, MU, UU, UM, MU; and the ninth possible game is the unrestricted MM.

2.4.1 Environments and Strategy Space

Particular environments are defined by specifying the market model (number of firms and collection sizes), and a preference model comprising the value of the substitution parameter γ and the probability distributions over parameters w and k in the population of consumers. I have analyzed a total of nine environments: for each of the three preference models described in Table 2.1, I analyzed three market models (Table 2.2).

The full strategy profile space is infinite ($\mathbb{R}^4_+$). To make the problem manageable, I limit sampling range to the following price intervals: $P_i \in (0, 1920]$ and $p_i \in (0, 120]$. The upper bounds on the price intervals were chosen based on a few rounds of preliminary simulations with $P_i \leq 2400$ and $p_i \leq 120$, for which all pricing strategies that survived iterated elimination of strongly dominated strategies (IESDS) were such that $P_i \leq 1800$ and $p_i \leq 115$. Similarly, zero price did not survive IESDS in any of the preliminary simulations, and therefore was not included in the intervals. For each of the nine environments, I evaluated the nine subgames (BB, BU, UB, BM, MU, UU, UM, MU, and MM) sampling prices within the same range and at the same constant intervals:

$$\begin{aligned} p_i &\in \{4, 8, 12, \ldots, 120\}, \\ P_i &\in \{40, 80, 120, \ldots, 1920\}. \end{aligned} \tag{2.12}$$

[11] To find Nash equilibria I used Gambit, a library of game theory software and tools for the construction and analysis of finite extensive and normal form games (McKelvey, McLennan, & Turocy, 2007). In particular, I used the `gambit-lcp` algorithm for solving two-player nonzero-sum games (documentation available at `http://gambit.sourceforge.net/doc/gambit-0.2007.01.30/gambit/`). This algorithm does not necessarily find all equilibria.

Thus, each firm's strategy space consists of 48 strategies in the BB game, 30 strategies in the UU game, and 1440 (48×30) strategies in the MM game. The interval sizes (40 for P_i and 4 for p_i) were primarily dictated by computational considerations. For example, it took about 2 days to do payoff estimation for each of the MM subgames for the preference model P3 (see Table 2.1 below), and about 5 more days for Gambit to find 189 equilibria for one of them. I have also analyzed the following larger "refined" strategy spaces:

$$\begin{aligned} p_i &\in \{4, 8, 12, \ldots, 120\}, \\ P_i &\in \{20, 40, 60, \ldots, 1920\}; \end{aligned} \tag{2.13}$$

$$\begin{aligned} p_i &\in \{2, 4, 6, \ldots, 120\}, \\ P_i &\in \{40, 80, 120, \ldots, 1920\}; \end{aligned} \tag{2.14}$$

$$\begin{aligned} p_i &\in \{2, 4, 6, \ldots, 120\}, \\ P_i &\in \{20, 40, 60, \ldots, 1920\}; \end{aligned} \tag{2.15}$$

$$\begin{aligned} p_i &\in \{3, 6, 9, \ldots, 120\}, \\ P_i &\in \{30, 60, 90, \ldots, 1920\}. \end{aligned} \tag{2.16}$$

For these strategy spaces, I was not able to solve all of the nine subgames. In particular, the MM game in the P3-SD environment (see Tables 2.1–2.2 below) was too large for Gambit to solve in less than a week. Also, it took prohibitively long to solve the MM subgame (for which all equilibria I could find are strictly mixed-strategy equilibria with relatively large supports) in any of the environments for (2.15). Since the MM subgame is central to the analysis of mixed bundling, I report results only for the strategy space (2.12), for which I was able to solve all the subgames of interest. These results are not significantly different from those obtained for the finer strategy sets: in all the equilibria I could find for (2.13)–(2.16), the equilibrium price ranges significantly overlap[12] and the equilibrium payoff of any player lies within –4% to +1% of the corresponding equilibrium payoff for the strategy space (2.12).

Table 2.1 summarizes the preference models I have analyzed. Model P3, in which both w and k have non-degenerate independent distributions, is the closest to the Chuang-Sirbu model.[13] In models P1 and P2, one of the parameters is constant. These models allow me to analyze the effect of competition for each type of the corresponding one-dimensional heterogeneity type.

[12] In all cases of multiple equilibria, the equilibrium price ranges always significantly overlap within the multiple-equilibrium set.

[13] They assume $\gamma = 0$, w is uniformly distributed between 0 and 1, and k is distributed exponentially with $\lambda = 13.8758$.

PID	w	k
P1	100	Exp(6.568)
P2	U[0,200]	0.152
P3	U[0,200]	Exp(6.568)

Table 2.1: List of preference models. In all models $\gamma = 0.5$. The choice of preference distributions and the specific distribution parameters is motivated by empirical studies by King and Griffiths (1995) (see Appendix B).

MID	N_1	N_2	Number of firms	Description
SD	100	100	2	Symmetric duopoly
ND	50	150	2	Non-symmetric duopoly
M	200	n/a	1	Monopoly

Table 2.2: List of market models.

For each preference model, I analyze three market structures given in Table 2.2. I label preference models and market models by their ID (PID and MID, respectively) and each environment as a pair PID-MID. Note that the total number of items on the market is 200 for all market models. To study the effect of competition, I compare the properties of the equilibrium outcomes of two duopoly models (one symmetric and one non-symmetric) to those of the optimal monopoly strategy.

The choice of preference distributions and the specific distribution parameters is motivated by empirical studies of distribution of scholarly articles performed by King and Griffiths (1995). Their results suggest that the parameter k follows an exponential distribution in the population of readers that the authors sampled (see Appendix B). Using the empirical data, I fit k to an exponential distribution with the parameters specified in preference model P1. See details in Appendix C. For the alternative heterogeneity model P2, I fix the breadth at the mean of the exponential distributions of model P1 ($k = 0.152$). I chose the range of w to be such that the mean intensity is equal to the fixed intensity value in model P1 ($\mathrm{E}(w) = 100$).

To estimate the firms' expected profits, I used the following procedure. For the preference model P2, in which $w \sim$U[0,200], I computed the demand of all consumers with integer $w \in \{1, \ldots, 200\}$. Below I refer to these types as w_g, where $g \in \{1, \ldots, 200\}$. For the model P1, in which $k \sim$Exp(6.568), I computed the demand for 284 consumers with k-types given by $k_{q+1} = k_q + 0.001\delta^{q-1}$ for $q = \{1, \ldots, 284\}$, where $k_0 = 0$ and $\delta = 1.03$. The 284th type is $k_{284} = 147.4186$. For the model P3, I computed the demand of 56800 (200×284) consumers with all combinations of w_g and k_q used in models P1 and P2. Given the demand, I computed the profit,

$\pi_i(w_g, k_q)$, that each firm i would earn from each estimated consumer type (w_g, k_q). The estimated expected profit of firm i, $i \in \{1, 2\}$ is given by

$$\mathrm{E}\pi_i = \sum_{\substack{g=\{1,\dots,200\}\\ q=\{1,\dots,284\}}} \mathrm{P}(w_{g-1} < W < w_g)\mathrm{P}(k_{q-1} < K < k_q)\pi_i(w_g, k_q), \qquad (2.17)$$

where $w_0 = 0$, $k_0 = 0$, and W and K are random variables following the distributions from Table 2.1.

Such procedure appears to produce rather accurate profit estimates. For the UU subgame under P3-preferences and three market models – SD, ND, and M – I estimated the equilibrium (optimal) prices and profits analytically.[14] For the M-model, I also found the optimal price and payoff for the case when the monopoly is restricted to pure bundling. For the five equilibrium (optimal) payoff estimates obtained in these four games,[15] the numerically estimated profits are all systematically lower, by at most 1.5%, than the theoretical estimates. The equilibrium (optimal) prices are within 6% of the theoretical estimates (lower by up to 4% and higher by up to 6%).

2.4.2 Empirical Results

In Table 2.3, I report for the duopoly market models the number of strategies that survived the iterated elimination of strongly dominated strategies (IESDS) in each of the nine subgames. Games with many survivors tend to have multiple strictly mixed-strategy equilibria. The largest number of multiple equilibria in a game is 189 and the second largest is 3. All multiple equilibria have very similar payoffs (see Table 2.4). Therefore, in the case of multiple equilibria, I report only one of them, giving priority to symmetric Pareto superior equilibria.

The largest support of a single mixed-strategy equilibrium includes 26 strategies for each firm, but most mixed-strategy equilibria do not have more than 5–10 strategies in the support. See Tables 2.5–2.7 for the supports of representative equilibria in each subgame. In the tables, a row and a column define a subgame. In each cell, the top number(s) is the equilibrium price (price pair) of the row firm, and the bottom number(s) is the equilibrium price (price pair) of the column firm. I denote a mixed-bundling price pair as *p*-*P*. For example, in the UM

[14] I did not impose the constraints that the quantity consumed by each consumer type from a collection is less than the collection size. Therefore, the analytically computed profits may be overestimated.

[15] In the non-symmetric duopoly, the two firms have different equilibrium (optimal) payoffs and prices.

Symmetric Duopoly				Non-Symmetric Duopoly			
P1-SD				P1-ND			
	Firm 2(1)				Firm 2 (the bigger)		
Firm 1(2)	U	B	M	Firm 1	U	B	M
U	1 1			U	1 1	1 1	2 2
B	1 1	1 1		B	1 1	1 1	2 4
M	1 1	5 4	26 26	M	1 1	6 5	8 11
P2-SD				P2-ND			
	Firm 2(1)				Firm 2 (the bigger)		
Firm 1(2)	U	B	M	Firm 1	U	B	M
U	1 1			U	1 1	1 1	2 2
B	1 1	6 6		B	1 1	2 6	2 9
M	1 1	35 6	47 47	M	1 1	13 6	14 21
P3-SD				P3-ND			
	Firm 2(1)				Firm 2 (the bigger)		
Firm 1(2)	U	B	M	Firm 1	U	B	M
U	1 1			U	1 1	1 1	1 1
B	1 1	7 7		B	1 1	2 6	1 1
M	1 1	38 7	116 116	M	1 1	7 5	26 37

Table 2.3: Number of survivors of iterated elimination of strongly dominated strategies (IESDS). A row and a column define a subgame. In each cell, the first number is the number of survivors for the row firm, and the second number is for the column firm. For example, in the MM subgame (row M and column M) in environment P1-ND, 8 strategies remained in the smaller (row) firm's strategy set, and 11 remained in the bigger (column) firm's strategy set. Each firm's original strategy space consists of 48 strategies in the BB game, 30 strategies in the UU game, and 1440 (48×30) strategies in the MM game.

Environment	Subgame	# of Multiple Eq.	Profit Diff. (%)
P1-SD	MM	3	0.02
P2-SD	BB	3	1.68
P3-SD	BB	3	0.57
P3-SD	MM	189	0.32

Table 2.4: Subgames in which Gambit found multiple equilibria. In the last column I report the (absolute) difference in the payoff of the unreported equilibria as a percentage of the reported equilibrium payoff.

subgame (row U and column M) for the symmetric duopoly under the P1-preference distribution (Table 2.5, P1-SD), firm 1's equilibrium (per-item) price is 40, and firm 2's equilibrium (mixed-bundling) price pair is 60-720, which means that it offers a choice between buying individual items for 60 each and buying the whole collection for 720. For large supports, I report ranges of prices. For example, "10 in $[60,80]\times[440,480]$" means that there are 10 price pairs in the support of this equilibrium mixed-bundling strategy, and the first component – the per-item price – varies from 60 to 80, and the second component – the bundle price – varies from 440 to 480.

The fact that many equilibria involving bundling are strictly mixed raises the question of their practical implementation and interpretation. In the context of information-good production, the support of mixed-strategy equilibria can be interpreted as a near-equilibrium price region. Note that in Table 2.6, which describes the non-symmetric duopoly, the sizes of the supports are all less than three, and the prices are rather close. Since firms are never perfectly symmetric, this suggests that in practice such price regions may be small. One example when a randomized choice from these regions may arise naturally is during a readjustment to some external shock, when both firms need to react immediately and simultaneously.

I first analyze the relative performance of different duopoly equilibrium pricing schemes, or the optimal prices in case of a monopoly. Table 2.8 presents equilibrium profits for the duopoly subgames in each of the environments. Table 2.9 presents total industry profits for duopolies as well as the monopoly. We can see that the relative payoffs of mixed bundling compared to pure unbundling and pure bundling depend on the type of consumer heterogeneity. The differences in profits are more clearly seen from Table 2.10. For the monopoly, we have the following result (see the last column in Table 2.10).

(i) Mixed bundling yields higher profits relative to pure bundling when breadth varies exponentially across consumers (preference models P1 and P3): pure bundling yields 80–83% of the mixed-bundling profits according to the empirical results.

(ii) Pure bundling is dominated by pure unbundling when breadth varies across consumers (models P1 and P3): the empirical analysis shows that pure bundling yields 80% the mixed-bundling profit vs. 90% produced by pure unbundling in P1; 83% vs. 99.2% in P3.

(iii) Pure bundling is as profitable as mixed bundling when the breadth is the same for all consumers, while the intensity of their preferences vary (model P2).

Symmetric Duopoly			
P1-SD			
	Firm 2(1)		
Firm 1(2)	U	B	M
U	44 44	40 480	40 60-720
B	480 40	360 360	400 440 56-520 64-400
M	60-720 40	56-520 64-400 400 440	10 in [60,80]×[440,480] 10 in [60,80]×[440,480]
P2-SD			
	Firm 2(1)		
Firm 1(2)	U	B	M
U	52 52	52 640	52 100-640
B	640 52	5 in [400×560] 5 in [400×560]	5 in [400,560] 5 in [72,96]×[400,560]
M	100-640 52	5 in [72,96]×[400,560] 5 in [400,560]	5 in [72,96]×[400,560] 5 in [72,96]×[400,560]
P3-SD			
	Firm 2(1)		
Firm 1(2)	U	B	M
U	52 52	48 640	48 60-960
B	640 48	5 in [320,520] 5 in [320,520]	5 in [440,600] 5 in [56,76]×[440,920]
M	60-960 48	5 in [56,76]×[440,920] 5 in [440,600]	26 in [56,88]×[440,920] 26 in [56,88]×[440,920]

Table 2.5: Strategies in the support of representative equilibria: Symmetric duopoly. A row and a column define a subgame. In each cell, the top number(s) is the equilibrium price (price pair) of the row firm, and the bottom number(s) is the equilibrium price (price pair) of the column firm. I denote a mixed-bundling price pair as p-P. For example, in the UM subgame (row U and column M) under the P1-preference distribution (P1-SD), firm 1's equilibrium (per-item) price is 40, and firm 2's equilibrium (mixed-bundling) price pair is 60-720, which means that it offers a choice between buying individual items for 60 each and buying the whole collection for 720. For large supports, I report ranges of prices. For example, "10 in [60,80]×[440,480]" means that there are 10 price pairs in the support of this equilibrium mixed-bundling strategy, and the first component – the per-item price – varies from 60 to 80, and the second component – the bundle price – varies from 440 to 480.

Non-Symmetric Duopoly			
P1-ND			
	Firm 2 (the bigger)		
Firm 1	U	B	M
U	44	36	36 40
	48	680	60-920 60-960
B	240	160	240 280
	44	600	60-840 60-880
M	56-360	56-240 64-200 76-200	60-240 64-240 68-280
	44	600 640 720	60-840 60-880 68-760
P2-ND			
	Firm 2 (the bigger)		
Firm 1	U	B	M
U	52	48	48 52
	56	920	84-960 88-920
B	280	240 280	240 280
	56	800 960	92-880 96-960
M	80-280	80-240 88-280	64-240 80-240 88-280
	56	800 960	92-800 92-880 96-960
P3-ND			
	Firm 2 (the bigger)		
Firm 1	U	B	M
U	48	44	44
	56	920	64-1560
B	320	200 240	320
	56	720 840	60-1520
M	60-440	76-240 80-240	64-360 68-320 72-360
	52	720 760	60-1440 68-1040 72-1040

Table 2.6: Strategies in the support of representative equilibria: Non-symmetric duopoly. A row and a column define a subgame. In each cell, the top number(s) is the equilibrium price (price pair) of the row firm, and the bottom number(s) is the equilibrium price (price pair) of the column firm. I denote a mixed-bundling price pair as p-P. For example, in the UM subgame (row U and column M) under the P1-preference distribution (P1-ND), firm 1's equilibrium strategy is a probability distribution over two (per-item) prices: 36 and 40; firm 2's equilibrium strategy is a probability distribution over two mixed-bundling price pairs: 60-920 and 60-960. The first price pair means that the firm offers a choice between buying individual items for 60 each and buying the whole collection for 920. Similarly for the second price pair.

Monopoly			
	P1-M	P2-M	P3-M
U	48	60	60
B	1040	1280	1320
M	60-1360	108-1280	68-1920

Table 2.7: Optimal strategies: Monopoly. The rows are the pricing schemes, and the columns are the preference models. For example, in row U, column P1-M, the monopoly is assumed to be restricted to pure unbundling and the consumers' preference distribution is described by the model P1. I denote a mixed-bundling price pair as p-P. For example, 60-1360 in row M, column P1-M, means that the monopoly offers a choice between buying individual items for 60 each and buying the whole collection for 1360.

(iv) Pure unbundling is almost as profitable as mixed bundling when both breadth and intensity vary (model P3): it yields 99.2% of the mixed-bundling profits according to the empirical results.

(v) When one of the preference parameters is constant across consumers (models P1 and P2), mixed bundling yields higher profits relative to pure unbundling: the latter yields 85–90% of the mixed-bundling profits according to the empirical analysis.

We already know from Chuang and Sirbu (1998) that pure unbundling can dominate pure bundling even in the presence of (weak) economies of scale and that mixed bundling strictly dominates the two when the marginal cost is positive. The authors conclude that "the choice of optimal bundling strategy lies in the balance between cost-savings from bundling and loss of surplus due to exclusion violation,"[16] where exclusion violation is the inefficiency that arises from consumption at submarginal-cost levels. The result above shows that another important factor is the distribution of consumer types. Moreover, as the transaction and distribution costs of digital information goods become negligibly small, the consumer side of the problem caries increasingly more weight.

The result underscores the intuition Chuang and Sirbu offered to motivate their two-dimensional preference model for studying mixed bundling: If the correlation between items (modeled as breadth here) within a collection is the same across consumers, as in model P2, then consumer reservation prices adequately capture the diversity in consumer tastes, and the monopolist can capture as much surplus through pure bundling as through mixed bundling. This is consistent with earlier

[16]Chuang and Sirbu (1998), p. 13.

studies that show that pure bundling allows a monopolist to extract all consumer surplus by reducing buyer diversity, when consumers have i.i.d. valuations over items (e.g., Bakos & Brynjolfsson, 1998). When correlation varies widely, however, as is often the case for scholarly journals, music CDs, TV channels, pure bundling is strictly dominated not only by mixed bundling, but can also be dominated by pure unbundling (models P1 and P3 here).

Under competition, we are interested in the equilibrium performance of different schedules. In Table 2.10, I present each firm's expected equilibrium profit as a percentage of their profit in the subgame in which they employ mixed bundling, while the other firm's scheme is fixed. We can see that the pattern is similar to that in the monopoly case. Particular numbers differ, but the relative profit sizes are largely the same. The only exception is the non-symmetric duopoly under P1-preferences: for the smaller firm, the relative profitability of U and B in equilibrium is reversed when the other (bigger) firm uses bundling in the pure or mixed form. This suggests that under competition, each pricing scheme works largely in the same way to extract consumer surplus as under monopoly.

In some environments, however, mixed bundling leads to more aggressive price competition, which overrides its benefits as a price-discrimination mechanism. In such environments, pure schemes yield higher equilibrium profits than the mixed-bundling scheme, although not by much. For example, in the case of a symmetric duopoly under P2-preferences, the equilibrium in the BB subgame is better for both firms than the equilibrium in the UUand MM games (see Table 2.8, environment P2-SD). Moreover, once in that equilibrium, neither firm would benefit from unilaterally changing its price schedule: the resulting equilibrium prices would produce lower profits for the deviating firm. In fact, the mixed bundling scheme is weakly "dominated" by pure bundling, in the sense that it always leads to weakly lower profits after prices reach an equilibrium given the schedules, although the potential loss is below 1%. Note that this holds only in the environment in which mixed bundling does not offer an advantage as a price-discriminating tool even under monopoly, i.e., when the preference breadth is constant in the population.

If breadth varies as well as depth, the equilibrium in the UU subgame for the symmetric duopoly (environment P3-SD) is more profitable for both firms than that in MM. In this case, however, the firms would want to unilaterally expand their scheme to mixed bundling. Interestingly, the other firm would not follow the move, because that would decrease both firms' profits by around 5% after the price equilibrium is re-established (see Table 2.10: in the UM game, firm 1 earns 104.9% of its profit in the MM game). This pattern is observed in the environment in which pure unbundling is almost as profitable as mixed bundling under monopoly, while

Symmetric Duopoly							Non-Symmetric Duopoly						
	P1-SD							P1-ND					
	Firm 2(1)							Big Firm					
F. 1(2)	U		B		M		Small	U		B		M	
U	244	244	204	214	216	269	U	125	367	94	322	99	406
B	214	204	204	204	212	246	B	107	335	103	308	114	388
M	269	216	246	212	247	247	M	137	344	119	316	132	388
	P2-SD							P2-ND					
	Firm 2(1)							Big Firm					
F. 1(2)	U		B		M		Small	U		B		M	
U	250	250	247	296	247	296	U	127	376	122	439	123	441
B	296	247	264	264	263	262	B	150.6	368	154	415	156	415
M	296	247	262	263	262	262	M	151.1	368	154	415	155	415
	P3-SD							P3-ND					
	Firm 2(1)							Big Firm					
F. 1(2)	U		B		M		Small	U		B		M	
U	217	217	194	189	205	222	U	112	324	94	277	104	329
B	189	194	160	160	181	194	B	100	305	93	255	101	311
M	222	205	194	181	196	196	M	115	311	99	261	113	313

Table 2.8: Duopoly expected equilibrium profits. A row and a column define a subgame. In each cell, the left number is the equilibrium expected profit of the row firm, and the right number is that of the column firm. For example, in the MM subgame (row M and column M) in environment P1-ND, the smaller (row) firm's equilibrium profit is 132, and the bigger (column) firm's equilibrium profit is 388.

Symmetric Duopoly				Non-Sym. Duopoly				Mon.	
P1-SD				P1-ND				P1-M	
	Firm 2(1)				Big Firm				
Firm 1(2)	U	B	M	Small	U	B	M		
U	488	419	485	U	493	416	506	U	495
B	419	408	458	B	443	411	502	B	437
M	485	458	493	M	481	435	520	M	548
P2-SD				P2-ND				P2-M	
	Firm 2(1)				Big Firm				
Firm 1(2)	U	B	M	Small	U	B	M		
U	500	543	543	U	503	561	564	U	505
B	543	528	526	B	519	569	571	B	592
M	543	526	525	M	519	569	570	M	592
P3-SD				P3-ND				P3-M	
	Firm 2(1)				Big Firm				
Firm 1(2)	U	B	M	Small	U	B	M		
U	435	383	427	U	436	370	434	U	440
B	383	320	375	B	405	349	413	B	367
M	427	375	390	M	426	360	426	M	444

Table 2.9: Total expected (equilibrium/optimal) profit. For the duopoly models, a row and a column define a subgame. For example, the number in a row U and a column M is the total expected equilibrium profit when the row firm is restricted to unbundling its items and the column firm can offer a mixed bundle (subgame UM).

Symmetric Duopoly				Non-Symmetric Duopoly								Mon.	
P1-SD				P1-ND								P1-M	
	Firm 2				Big Firm				Small Firm				
F. 1	U	B	M	Small	U	B	M	Big	U	B	M		
U	91	83.1	87	U	91	79	75	U	90	86	89	U	90
B	80	82.9	86	B	78	87	86	B	79	79	81	B	80
P2-SD				P2-ND								P2-M	
	Firm 2				Big Firm				Small Firm				
F. 1	U	B	M	Small	U	B	M	Big	U	B	M		
U	85	94	94	U	84	79	80	U	85	89	89	U	85
B	100.0	100.6	100.4	B	99.7	100.0	100.4	B	99.7	99.9	99.9	B	100.0
P3-SD				P3-ND								P3-M	
	Firm 2				Big Firm				Small Firm				
F. 1	U	B	M	Small	U	B	M	Big	U	B	M		
U	98.1	99.8	104.9	U	97	95	93	U	98.5	97.9	99.6	U	99.2
B	85	82	93	B	87	94	90	B	84	82	84	B	83

Table 2.10: Total expected (equilibrium/optimal) profit in percent of the mixed-bundling profit. For the duopoly models, the total expected equilibrium profit in the UX and BX subgames, where X can be U, B, or M, is given as a percentage of the equilibrium profits in the MX game.

pure bundling yields only 83% of the mixed-bundling profit. That is, when the population distribution is such that bundling fails to reduce buyer diversity while individual-article sales capture almost all surplus a monopoly could capture. Again, increased price competition from utilizing a mixed-bundling scheme hurts profits in an environment in which mixed bundling extracts little more surplus relative to one of the pure forms. Under P1-preferences, when mixed bundling has clear benefits for a single firm, the price competition that corresponds to $\gamma = 0.5$ is not strong enough to override those benefits.

The assumption of the two-dimensional heterogeneity of consumer tastes is key to these results: if the consumers are *ex ante* homogeneous, pure bundling achieves the first-best solution under monopoly (Bakos & Brynjolfsson, 1998) as well as under duopoly (Fay & MacKie-Mason, 1999, Proposition 1), while pure unbundling yields less profit than the first-best solution (Fay & MacKie-Mason, 1999, Proposition 2).

Comparing the duopoly total expected profits to those of a monopoly (see Table 2.11), we see that competition reduces the industry-wide profits by up to 21% depending on the preference model and the combination of pricing schemes the firms employ. Table 2.8 shows that the firms' shares of those profits are roughly proportional to their collection sizes. If we compare by pricing schemes, the effect of competition on profits of bundling firms is far greater than on firms relying on sales of individual items. In the worst case, the latter earn 1% below monopolistic profits. If firms employ pure or mixed bundling, competition can reduce profits by up to 13% if the firms use the same pricing scheme (column "Diagonal" in Table 2.11) and by up to 21% if the firms use different schemes (column "Overall" in Table 2.11). This is not surprising given that when firms bundle, a single purchase carries more weight.

Finally, in Table 2.12 I report market efficiency for all duopoly subgames as well as for the monopoly. Depending on the preference model and pricing scheme, competition increases efficiency by up to 16% (see the last column in Table 2.12). The distribution of welfare is also affected in a predictable way: under each pricing scheme, competing firms can capture a smaller share of the social welfare compared to a monopoly. As Table 2.13 shows, unbundling firms lose up to 5% of their welfare share, firms employing only pure bundling or only mixed bundling (BB and MM subgames) lose up to 15%, and the largest difference overall is 22%.

These findings contrast with the case of *ex ante* homogeneous consumers: since pure bundling achieves the first-best solution under both monopoly and duopoly, both market structures are efficient. The firms' profits also fall, however, but for the case of *ex ante* homogeneous consumers, this is a direct consequence of the change

Symmetric Duopoly				Non-Sym. Duopoly				Largest Diff.	
P1-SD				P1-ND				Diagonal	Overall
	Firm 2(1)				Big Firm				
F. 1(2)	U	B	M	Small	U	B	M		
U	-1	-15	-11	U	0	-16	-8	-1	
B	-15	-7	-16	B	-11	-6	-8	-7	
M	-11	-16	-10	M	-12	-21	-5	-10	-21
P2-SD				P2-ND					
	Firm 2(1)				Big Firm				
F. 1(2)	U	B	M	Small	U	B	M		
U	-1	-8	-8	U	0	-5	-5	-1	
B	-8	-11	-11	B	-12	-4	-4	-11	
M	-8	-11	-11	M	-12	-4	-4	-11	-12
P3-SD				P3-ND					
	Firm 2(1)				Big Firm				
F. 1(2)	U	B	M	Small	U	B	M		
U	-1	-13	-4	U	-1	-16	-2	-1	
B	-13	-13	-15	B	-8	-5	-7	-13	
M	-4	-15	-12	M	-4	-19	-4	-12	-19

Table 2.11: Difference between expected monopoly profits and total expected equilibrium profits in percent of monopoly profits. A row and a column define a duopoly subgame. For example, the profit of the subgame UB is at the intersection of row U and column B. The difference for each subgame X_1X_2 is calculated based on the assumption that the monopoly would chose the pricing scheme X_i, $i \in \{1,2\}$ with the highest profit. That is, the total expected equilibrium profit in a UB subgame, for example, is compared to the maximum of the monopoly's U- and B-profits; the total equilibrium profit in a UM subgame is compared to the monopoly's profit under the M-scheme. In the last two columns I report the largest differences across both duopolies for (a) the diagonal subgames UU, BB, and MM(note that these are all the SD-diagonals), and (b) over all possible subgames (note that these are all ND-values).

Symmetric Duop.				Non-Sym. Duop.				Mon.		Largest Diff.	
P1-SD				P1-ND				P1-M		Diag.	Overall
	Firm 2(1)				Big Firm						
F. 1(2)	U			Small	U	B	M	Mon.			
U	81	78	80	U	78	79	81	U	77	4	
B	78	85	83	B	78	83	81	B	72	13	
M	80	86	86	M	80	82	84	M	80	6	14
P2-SD				P2-ND				P2-M			
	Firm 2(1)				Big Firm						
F. 1(2)	U	B	M	Small	U	B	M	Mon.			
U	78	77	77	U	76	78	77	U	72	6	
B	77	87	87	B	77	80	80	B	75	11	
M	77	87	87	M	77	80	80	M	75	12	15
P3-SD				P3-ND				P3-M			
	Firm 2(1)				Big Firm						
F. 1(2)	U	B	M	Small	U	B	M	Mon.			
U	78	76	79	U	77	75	77	U	73	6	
B	76	86	80	B	75	79	76	B	70	16	
M	79	80	84	M	78	80	79	M	75	9	16

Table 2.12: Efficiency (in percent of maximum social welfare). For the duopoly models, a row and a column define a subgame. For example, in the UU subgame (row U and column U) in environment P1-SD, the efficiency is 81%. In the same subgame in P1-ND, it is 78%. In the last two columns, I report (a) the difference between the highest efficiency in duopoly subgames UU, BB, and MM and the monopoly efficiency for the schemes U, B, and M, respectively; and (b) the largest difference over all pricing schemes. For example, the difference for the U-scheme under preferences P1 is $\max(81, 78) - 77 = 4\%$.

in the distribution of the welfare due to competition.[17]

2.5 Conclusion

In this chapter, I have explored the interaction between competition and the bundling of electronically delivered information goods. One of the main contributions

[17] Fay and MacKie-Mason (1999) show that under duopoly, consumers retain $V_1 + V_2 - V_B$, where V_1 and V_2 are the consumer values for the first and second collections, respectively, and V_B is the value for both collections. Since Fay and MacKie-Mason assume that consumers' valuations of the collections are strictly subadditive, this quantity is positive (as opposed to zero under monopoly). The relative size of the consumer share, however, depends on the shape of the consumer valuation function.

Symmetric Duopoly				Non-Symmetric Duop.				Mon.		Largest Diff.	
P1-SD				P1-ND				P1-M		Diag.	Overall
	Firm 2(1)				Big Firm						
F. 1(2)	U	B	M	Small	U	B	M	Mon.			
U	61	54	61	U	64	53	63	U	65	-4	
B	54	48	55	B	57	50	63	B	61	-13	
M	61	55	58	M	61	53	63	M	69	-11	-21
P2-SD				P2-ND				P2-M			
	Firm 2(1)				Big Firm						
F. 1(2)	U	B	M	Small	U	B	M	Mon.			
U	53	59	59	U	55	60	61	U	58	-5	
B	59	51	50	B	56	59	59	B	66	-15	
M	59	50	50	M	56	59	59	M	66	-15	-15
P3-SD				P3-ND				P3-M			
	Firm 2(1)				Big Firm						
F. 1(2)	U	B	M	Small	U	B	M	Mon.			
U	53	48	52	U	54	47	54	U	58	-5	
B	48	36	45	B	52	42	52	B	50	-14	
M	52	45	45	M	52	43	52	M	57	-12	-22

Table 2.13: Firms' share of actual social welfare (in percent). For the duopoly models, a row and a column define a subgame. For example, in the UU subgame (row U and column U) in environment P1-SD, the total expected equilibrium profits account for 61% of the actual welfare. In the same subgame in P1-ND, it is 64%. In the last two columns I report (a) the difference between the lowest duopoly welfare share in subgames UU, BB, and MM and the monopoly welfare share for the schemes U, B, and M, respectively; and (b) the largest difference over all pricing schemes. For example, the difference for the U-scheme under preferences P1 is $\min(61, 64) - 65 = -4\%$.

of this work is to focus attention on bundling (monopoly and competitive) with heterogeneous consumers. In particular, consumers vary in the value of their most favored item in a collection as well as in the percentage of items they value positively in that collection. This approach allows me to capture not only the variation in the consumer reservation price for the collection, but also in the amount of correlation among item valuations within the collection, which is an important preference characteristic in the context of (mixed) bundling information goods.

I use an empirical-game methodology developed elsewhere in collaboration with other co-authors to solve the duopoly game. I find that when there are no marginal costs of production for existing information goods, the relative profitability of mixed bundling, pure bundling, and pure unbundling is defined by the preference distribution of the consumers. Generally, the relative performance of these schemes in a duopoly equilibrium – when firms compete in price after setting on the scheme – is the same as under monopoly:

- Mixed bundling *strictly* dominates the pure schemes in highly heterogeneous environments, working as a price-discrimination tool that provides incentives for different categories of consumers to self-select into buying individual items or bundles;

- Pure bundling can extract almost as much consumer surplus as mixed bundling through its aggregation effect (effect of reducing buyer diversity) when the number of items consumers value positively (preference breadth) in a collection is constant across the population and only the "intensity" of their preferences varies – that is, when the main source of heterogeneity lies in the reservation price for a collection;

- When the preference breadth varies, pure bundling loses its aggregation power, and pure unbundling can be more profitable, and even as profitable as mixed bundling.

When firms compete, the general mechanisms by which each pricing scheme works to extract consumer surplus remain the same, but price competition can reduce the price-discriminating power of mixed bundling to the extent that pure schemes can result in higher equilibrium profits than mixed bundling. This can happen under a symmetric duopoly when the distribution of consumer preferences is such that a monopolist is almost indifferent between mixed bundling and the pure scheme in question.

Comparing the effect of competition on social welfare and its distribution, I find that it has a negligible impact on the industry profits as well as consumer surplus if items are sold individually. If firms employ bundling schemes – in the pure or mixed form – the effect is noticeable. The market efficiency under duopoly is greater than under monopoly. According to my empirical analysis, it is greater by up to 16% depending on consumer preference distribution and the particular type of bundling scheme. The distribution of social welfare predictably shifts toward consumers under competition, by up to 22%. The largest drop in profits is 21% relative to a monopoly employing the same (combination of) pricing schemes as the duopoly.

CHAPTER 3

Learning Bayesian Nash Equilibrium: An Experimental Study

3.1 Introduction

Games of incomplete information are important in game theory, as many strategic interactions can be characterized by such games, including auctions, bargaining, contracting and competitive markets. Since Harsanyi (1968), Bayesian Nash equilibrium (BNE) has been proposed as the main solution concept for normal-form games of incomplete information. However, it relies on several strong assumptions. First, BNE requires strong assumptions on the rationality of the players. Unlike games of complete information, in which rational play only requires a best response given the payoff structure, in games of incomplete information, rational play requires a player to think about what she would do if she were another type. This raises the question as to whether BNE is a reasonable solution concept if the rationality assumption is relaxed. Of particular interest to practical applications is the question whether BNE can arise as a result of a learning process by which the players adjust their choices in time based on the history of observations such as their payoffs and/or others' actions and payoffs. Such process does not necessarily require the players to be able to compute the equilibrium or even have enough information about the game to compute it. The term *bounded rationality* is often used to describe the assumption on such agents' behavior. Theoretical research establishes that the answer is positive for agents with some forms of bounded rationality at least in the class of potential games (e.g., Monderer and Shapley (1996b), Monderer and Shapley (1996a), Facchini, van Magen, Borm, and Tijs (1997)) and games with strategic complementarities (e.g., Vives (2000), Milgrom and Roberts (1990), Blume (1993)). In these theoretical models, the agents are assumed to follow an adaptive rule such as the *best-response dynamic* or an even less restricted type of adjustment. The empirical question is under what conditions, if any, BNE can arise as a result of human learning, and in particular, whether the theoretical results carry over to practical applications.

Second, BNE assumes common knowledge of the distribution of Nature's move, i.e., the distribution of player types. This assumption is not satisfied in many real-world strategic interactions, including online markets. In online auctions, for example, information technology enables geographically dispersed bidders to participate in the same auction, thus making it easier for a relatively large group of bidders to obtain an object. In such an auction, the common prior assumption (on bidder types) is almost always violated. In the computer-science community, there has been a debate about the relevance of BNE for the design and analysis of mechanisms for the Internet precisely because of the common-prior assumption. Theoretical research establishes that quite restrictive assumptions are necessary to justify the concept of BNE without a common prior as a steady state of a learning process (Dekel, Fudenberg, and Levine (2004), Bergemann and Valimaki (2006)). Empirically, if it can be demonstrated that, at least in some classes of games, BNE can arise as a result of learning without a common prior, it will provide a foundation to investigate the empirical relevance of the solution concept.

To investigate the learnability of BNE, I design a laboratory experiment that separates learning into two levels. First, when the distribution of types is known to the players, how accurately does BNE predict the outcomes to which subjects converge as a result of learning other players' actions, if they converge at all? Second, how similar is the answer when the distribution of the types in unknown? Under both information conditions, I design a class of two-player Bayesian games from the family of duopoly games by varying two structural properties that have been shown to affect learning of players with bounded rationality: *strategic complementarity* and the existence of a *potential*. Theoretically, games with strategic complements and potential games both have robust convergence properties with respect to a broad class of learning processes. In this experimental study, I find that both properties affect the convergence dynamics. Under low information, convergence is more robust in supermodular and potential games, confirming theoretical findings. The Bayesian-information condition, however, proved to be a more complicated case. Consistently with earlier experimental studies of complete-information games in which the equilibrium is not Pareto optimal, we don't observe such a drastic improvement in the convergence level. In addition, I find evidence of the subjects' aversion to losses.

3.2 Literature Review

Of the large literature on learning in games,[1] I will focus on learning in games of incomplete information. In this section, I first review the theoretical literature. I then review the relevant experimental literature.

3.2.1 Theoretical Literature on Learning under Incomplete Information

I first present a Bayesian-game framework and notation. I then review theoretical results on the structural properties of games which lead to convergence to BNE. In what follows, I will only review results on learning a unique BNE in pure strategies, which is most relevant for the experimental design.

I organize my overview of related theory around the following questions: What factors make an equilibrium of a (Bayesian) game easy or hard to learn? How important for learning are the informational assumptions that underlie the BNE concept?

In games with a unique equilibrium, the idea of structural properties facilitating learning can be formalized by some notion of the "dynamic stability" of the equilibrium with respect to a class of learning processes. In games with multiple equilibria, learnability has more than one dimension. One important question concerns equilibrium selection and stability properties of the selected equilibrium. Another way to look at learnability is to study the dynamic stability of the equilibrium set in general. In Sections 3.2.2–3.2.5, I provide an overview of structural properties that theoretically facilitate learning in games. The object of most theoretical studies I overview is *normal-form* (or *strategic-form*) games with Nash equilibrium as the solution concept, but the definitions and, in some cases, analysis can be extended to *Bayesian games* as well. It is worth noting that the structural properties discussed in Sections 3.2.2–3.2.5 are each sufficient to ensure that there exists an equilibrium in pure strategies. In Section 3.2.6, I review some limitations of the assumption that players have correct common priors in models predicting learning outcomes.

Let $N \neq \emptyset$ be the set of players. I will assume that N is finite, although some of the theorems presented here hold for infinite N as well. Each player $n \in N$ has a strategy set X_n with typical element x_n. This element is also referred to as player n's *pure strategy*. The competitors' (pure) strategies are denoted by x_{-n} and a full strategy profile is denoted by $x = (x_n, x_{-n}) \in X$, where $X = \times_n X_n$. Elements of X

[1] See Fudenberg and Levine (1998) for a survey of the theoretical learning literature. For a survey of the experimental learning literature, see Camerer (2003), and more recently, Erev and Haruvy (2008).

are called *pure-strategy profiles* or simply *strategy profiles*. Player n's payoff function is $\pi_n(x_n, x_{-n})$. The object $\Gamma = \{N, (X_n, \pi_n, n \in N)\}$ is a game in *normal form*.

In general, players' choices do not have to be deterministic. Randomization of choices gives rise to what is called a mixed strategy. Given player n's finite pure strategy set X_n, a *mixed strategy* for player n, $\delta_n : X_n \to [0,1]$, assigns to each pure strategy $x_n \in X_n$ a probability $\delta_n(x_n) \geq 0$ that it will be played, where $\sum_{x_n \in X_n} \delta_n(x_n) = 1$. If X_n is infinite, a mixed strategy $\delta_n : X_n \to [0,1]$ is a probability distribution over X_n, which integrates to one: $\int_{x_n \in X_n} \delta_n(x_n) = 1$. I will denote the set of mixed strategies of player n by Δ_n and the set of mixed-strategy profiles, $\times_n \Delta_n$, by Δ.

It is sometimes important that the strategy spaces are ordered. This is the case, for example, in supermodular and submodular games, which I overview in Sections 3.2.3–3.2.4. Let each strategy set X_n be endowed with a partial order $\geq$, and the strategy profiles be endowed with the *product order*, that is, $x \geq x'$ means $x_n \geq x'_n$ for all $n \in N$. The object $\Gamma_o = \{N, (X_n, \pi_n, n \in N), \geq\}$ is a game in *ordered normal form*.

In Bayesian games, the players have privately observed characteristics, or *types*, which affect their and other players' payoffs. Formally, let a random variable $\theta_n \in \Theta_n$ chosen by nature and observed only by player n be n's type. The joint probability distribution of the θ_n's is given by $F(\theta_1, \ldots, \theta_{|N|})$, which is assumed to be common knowledge among the players. I will denote by $\hat{\mu}$ the common beliefs of the players, which is a probability measure on $\Theta = \times_n \Theta_n$. The measure $\hat{\mu}_n$ will represent the marginal on Θ_n. The payoff to player n is given by $\pi_n : X \times \Theta \to \mathbb{R}$. A *Bayesian game* is the object $\Gamma_b = \{N, (X_n, \pi_n, n \in N), \Theta, F(\cdot)\}$. A *pure strategy* for player n is a map $x_n : \Theta_n \to X_n$, which assigns an action to every possible type of the player.

A *mixed strategy* is a map $\delta_n : \Theta_n \to \Delta_n$.

3.2.2 Stability with Respect to the Cournot Tatonnement

One of the oldest notions of stability in games is stability with respect to the Cournot tatonnement (Cournot best-reply dynamic). Assume that the players follow some sort of adaptation process. "Myopic" players may react to the strategies played by the other players by deviating to best-reply strategies. Considering the underlying time axis to be a continuum, such deviations may be assumed to take place one at a time. The resulting dynamic system is the continuous Cournot tatonnement. Below I define it more formally. The definition applies to both normal-form and Bayesian games.

Definition 3.1. *Suppose that $x_n \in \mathbb{R}$. Let $g_n(x_{-n})$ be player n's best-reply function.*

The continuous Cournot tatonnement *is a first order system of differential equations*

$$\dot{x}_n(t) = g_n(x_{-n}(t)) - x_n(t), \quad n \in N, \tag{3.1}$$

where t is time and $\dot{x}_n(t)$ is the derivative of $x_n(t)$ with respect to time.

A solution of the system is any function $x(t) = (x_n(t))$ that satisfies Equation (3.1). Any constant-function solution $x_n(t) = x_n^*$, $n \in N$, of Equation (3.1) is an equilibrium.

Definition 3.2. *An equilibrium $x^* = (x_n^*)$, $n \in N$, is called* asymptotically stable *if every solution $x(t)$ which starts near x^* converges to x^* as $t \to \infty$.*

Definition 3.3. *An equilibrium $x^* = (x_n^*)$, $n \in N$, is called* globally asymptotically stable *if for any initial condition $x(0) = x_0$ (with the possible exception of some lower-dimensional set of x_0's), the solution of the initial value problem $\dot{x}_n = g_n(x_{-n}(t)) - x_n(t)$, $x(0) = x_0$ tends to x^* as $t \to \infty$.*

In other words, an equilibrium is globally asymptotically stable if just about every solution $x(t)$ tends to x^* as $t \to \infty$.

Equilibrium stability for linear systems in $|N|$ dimensions can be determined based on the eigenvalues of the coefficient matrix of the system. Stability of an equilibrium of a non-linear system of autonomous differential equations can be determined based on the eigenvalues of the Jacobian matrix of the system. See, for example, Simon and Blume (1994).

Vives (1990) provides an important stability result for a special case of strictly increasing continuously differentiable best replies. This is stated by the following theorem.[2]

Theorem 3.4. *Suppose that strategy spaces are compact intervals and that best replies are strictly increasing continuously differentiable functions $g_n(\cdot)$, $n \in N$ (that is, we have $\frac{\partial g_n}{\partial x_m} > 0, m \neq n$). Then the continuous Cournot tatonnement converges to an equilibrium point of the game for almost all starting points x_0. When the number of players $|N| = 2$ and best replies are either strictly increasing or strictly decreasing convergence everywhere, as opposed to almost everywhere, is obtained.*

Note that the convergence is to *an* equilibrium point. When the best replies are continuous monotone functions, this can be a stable internal equilibrium or a corner solution. Intuitively, the assumption of strictly increasing (as well as strictly decreasing for $|N| = 2$) functions ensures that the Cournot tatonnement will never oscillate in a cycle and will converge to either an internal fixed point or a boundary

[2]The result is directly implied by Hirsch (1985), Theorem 5.1 and Corollary 2.8.

of the strategy space. The assumption that the strategy spaces are compact intervals guarantees that any fixed point (internal or boundary) belongs to the strategy space. The definition of the Cournot best-reply dynamic then implies that its end point is an equilibrium.

In the following sections, I overview theoretical results on equilibrium existence and stability in n-player games with strategic complementarities and two-player games with strategic substitutes. These results can be thought of as extensions of this simple result to a broader class of players' payoff functions, actions spaces, and learning dynamics. The underlying structural property that prevents a learning dynamic from cycling in loops is that of strategic complementarity of the players' actions, which in the above theorem is formalized as strictly increasing best-reply functions. In two-player games, strategic substitutability is also sufficient for the result.

3.2.3 Games with Strategic Complements

Strategic complementarity of players' actions can be formalized as increasing best replies, increasing differences of the payoff functions, or *supermodular* payoff functions. In general, supermodularity is a stronger property than increasing differences, and they both imply increasing best replies (Vives, 1990).

I adopt Milgrom and Roberts' definition of supermodular games (Milgrom & Roberts, 1990). Supermodular games are games in which each player's strategy set is partially ordered, the marginal returns to increasing one's strategy rise with increases in the competitors' strategies and, if a player's strategies are multidimensional, the marginal returns to any one component of the player's strategy rise with increases in any other component.

Definition 3.5. *Suppose that a typical strategy for player n is $(x_{ni}; i = 1, \ldots, k_n) \in \mathbb{R}^{k_n}$ and that $\geq$ is the usual componentwise ordering. Then Γ_o is a* supermodular *game if, for each $n \in N$:*

(A1) X_n is an interval in $\mathbb{R}^{k_n}$, that is,

$$X_n = [\underline{x}_n, \overline{x}_n] = \{x | \underline{x}_n \leq x \leq \overline{x}_n\};$$

(A2) π_n is twice continuously differentiable on X_n;

(A3) $\frac{\partial^2 \pi_n}{\partial x_{ni} \partial x_{nj}} \geq 0$ for all n and all $1 \leq i < j < k_n$;

(A4) $\frac{\partial^2 \pi_n}{\partial x_{ni} \partial x_{mj}} \geq 0$ for all $n \neq m$, $1 \leq i < k_n$ and $1 \leq j < k_m$.

In other words, in a supermodular game the players' preferences can be represented as supermodular functions defined on the strategy space.[3]

Supermodularity of the payoff functions creates a favorable environment for learning. The results of Milgrom and Roberts (1990), for example, imply that in supermodular games with a unique equilibrium, a broad class of learning processes converge to the equilibrium. More generally, Milgrom and Roberts (1990) show that the smallest and largest elements of the sets of pure-strategy equilibria, correlated equilibria, and rationalizable strategies are the same in supermodular games (Theorem 5, Milgrom and Roberts (1990)). Also, for a class of models of dynamic adaptive choice behavior that encompasses best-reply dynamics, fictitious play, and Bayesian learning, the players' choices lie eventually within the same bounds (Theorem 8, Milgrom and Roberts (1990)). To put it simply, these results characterize supermodular environments as robust to the choice of the solution concept and learning behavior, which in turn leads to solid theoretical predictions about the game outcomes.[4]

However, as Milgrom and Roberts point out, the usefulness of their results depends partly on the equilibrium set. The outcome prediction is the strongest in games with a unique equilibrium. For games with multiple equilibria, the learned outcome is not even guaranteed to be an equilibrium: it only has to lie "between" some of them. Equilibrium choice also remains an issue. This point is illustrated in Example 3.2.1, in which I consider a supermodular game with corner solutions. In this game, Theorem 8, Milgrom and Roberts (1990), may not narrow down the set of possible learning outcomes to anything smaller than the original strategy space, while direct stability analysis of the Cournot tatonnement (see Section 3.2.2) predicts that the learning-outcome set consists of the corner solutions.

Example 3.2.1. *There are two players, the strategy spaces are closed intervals and*

[3] Formally, the definition of supermodularity for a twice continuously differentiable function is as follows. Let $I = [\underline{x}, \overline{x}]$ be an interval in $\mathbb{R}^n$, $n \geq 2$. Suppose that $f : \mathbb{R}^n \to \mathbb{R}$ is twice continuously differentiable on some open set containing I. The function f is *supermodular* on I if for all $x \in I$ and all $i \neq j$, $\frac{\partial^2 f}{\partial x_i \partial x_j} \geq 0$. There are no restrictions on $\frac{\partial^2 f}{\partial x_i^2}$. More generally, the supermodularity condition is $f(\min(x, x')) + f(\max(x, x')) \geq f(x) + f(x')$, $x, x' \in \mathbb{R}^n$. Thus, supermodularity holds trivially for $n = 1$.

[4] In general, the conclusion about possible outcomes of learning in games is rather negative. Shapley (1964), for example, established that fictitious play (and the Cournot best-reply dynamic as a special case) can lead to an infinite pattern of cycling behavior for two-player, finite strategy, general sum games. Fudenberg and Kreps (1988) show that learning may yield a larger set of strategies than is identified by Nash equilibrium. See Milgrom and Roberts (1990), p. 1268 for a discussion. Dekel et al. (2004) establishes that the set of *self-confirming* equilibria, a solution concept incorporating the players' beliefs and observation histories, and the set of (Bayes-)Nash equilibria coincide only for a narrow class of games. In contrast, structural properties of the payoff functions such as supermodularity in games leads to more positive results.

the payoffs are $\pi_1(x_1, x_2) = -x_1^2 + 6x_1x_2$ *and* $\pi_2(x_1, x_2) = -x_2^2 + 4x_1x_2$*, respectively. The game is supermodular by Definition 3.5, because the cross-derivatives of the payoff functions are both positive (6 and 4, respectively). The best responses are straight lines with slopes greater than 1 (3 and 2, respectively). Suppose the strategy spaces include the origin, which is the intersection point of the best replies, so that the internal equilibrium is feasible. It can be shown that the internal equilibrium is unstable with respect to the Cournot best-reply dynamic, which means that a best-reply learning process will eventually stop at a boundary of the strategy set (i.e., at one of the corner solutions, depending on the players' initial choices). Milgrom and Roberts' results, however, only tell us that the players' choices will eventually lie in the area between those corner solutions, a set that may be equal to the strategy space.*

In many interesting games, however, the shapes of the best replies are more complex, especially when there are more than two players: the strategy sets are multidimensional, and the payoff functions are not concave. In such cases, the theory of supermodular games becomes very useful. For example, supermodularity implies increasing best replies (i.e., strategic complementarity) even if the players' payoff functions are not quasiconcave, as is often the case in economic applications (Vives, 1990). In addition, the strategy space does not have to be a product of compact intervals in $\mathbb{R}$. It can be any *complete lattice*, a family of sets that includes non-compact and discrete sets:[5]

Definition 3.6. *The game* Γ_o *is a* supermodular *game if, for each* $n \in N$*:*

(A1') X_n *is a complete lattice;*

(A2') $\pi_n : X \to \mathbb{R} \cup \{-\infty\}$ *is order upper semi-continuous in* x_n *(for fixed* x_{-n}*) and order continuous in* x_{-n} *(for fixed* x_n*) and has a finite upper bound;*

(A3') π_n *is supermodular in* x_n *(for fixed* x_{-n}*);*

(A4') π_n *has increasing differences in* x_n *and* x_{-n}.

The Definition 3.6 is the original definition given by Milgrom and Roberts (1990), and the equivalence of the Definitions 3.5 and 3.6 for the class of smooth payoff functions defined on a product of compact intervals follows directly from Topkis's Characterization Theorem (Topkis, 1979). Games satisfying Definition 3.5 are also called *smooth supermodular games.* All the results presented in this section hold for the general class of supermodular games defined as in Definition 3.6.

[5] For theory of lattices see Birkhoff (1940). Milgrom and Roberts (1990) also provide a brief overview of lattice theory.

When detecting or formulating supermodular games, it is important to understand the role of the ordering imposed on the strategy space. Many economic models that are not supermodular games in their original formulation can be reformulated as supermodular ones. In Example 3.2.2 below, I demonstrate one way of doing it.

Example 3.2.2. *There are two players, the strategy spaces are closed intervals and the payoffs are $\pi_1(x_1, x_2) = -x_1^2 - 6x_1x_2$ and $\pi_2(x_1, x_2) = -x_2^2 - 4x_1x_2$, respectively. The cross-derivatives of these functions are both negative, and therefore the game does not satisfy the assumptions of Definition 3.5. Thus, the game is not supermodular. However, there may exist an alternative interpretation of strategy spaces, under which the game will become supermodular. For example, suppose that x_1 and x_2 denote how fast the players do some action. Then the natural order ranks faster actions higher. Redefine player 2's strategy as $x_2' = -x_2$, which is equivalent to ordering her strategies by how slow the actions are. The new payoffs become $\pi_1'(x_1, x_2') = -x_1^2 + 6x_1x_2'$ and $\pi_2'(x_1, x_2') = -(x_2')^2 + 4x_1x_2'$, which are supermodular, as established in Example 3.2.1.*

See Vives (1990) and Milgrom and Roberts (1990) for examples of well known economic models that can be reformulated as supermodular games.

I now summarize Milgrom and Roberts' key results in more detail. A *learning process* is an arbitrary history of the players' choices as they repeatedly play the game. In loose terms, a pure strategy x_n is *justifiable* if no combination of other strategies the player n has played is better that x_n, according to the payoff feedback n observed in the play history. A process is one of *adaptive dynamics* if the players' choices eventually lie in the interval defined by the set of such justifiable strategies.[6] The definition of an adaptive dynamic imposes a very weak restriction on the players' choices: they do not all have to be justifiable; being bounded by justifiable strategies is all that is required.

The first key theorem states that in supermodular games, the smallest and largest elements in the set of serially undominated strategies[7] are always well defined concepts, they always exist, and they are always pure-strategy equilibria (Theorem 5,

[6] A formal definition relies on a few game-theoretic concepts and a few bits of new notation. A pure strategy x_n for player n is said to be *strongly dominated by another pure strategy* $\hat{x}_n$ if it is the case that for all x_{-n}, $\pi_n(x_n, x_{-n}) < \pi_n(\hat{x}_n, x_{-n})$. Given a product set $\hat{X}$ of strategy profiles, define the set of n's undominated responses to $\hat{X}$ by $U_n(\hat{X}) = \{x_n \in X_n | (\forall x_n' \in X_n)(\exists \hat{x} \in \hat{X})\pi_n(x_n, \hat{x}_{-n}) \geq \pi_n(x_n', \hat{x}_{-n})\}$. Let $U(\hat{X}) = (U_n(\hat{X}); n \in N)$ be the list of undominated responses for each player, and let $\overline{U}(\hat{X})$ denote the interval $[inf(U(\hat{X}), sup(U(\hat{X})]$. Let $\{x(t)\}, x(t) \in X, t \in T$, denote a *learning process*. The time index t can be discrete or continuous. Given a process $\{x(t)\}$, let $P(T, t)$ denote the strategies played between times T and t: $P(T, t) = \{x(s) | T \leq s \leq t\}$. A process $\{x(t)\}$ is one of *adaptive dynamics* if $(\forall T)(\exists T')(\forall t \geq T)x(t) \in \overline{U}([inf(P(T, t)), sup(P(T, t))])$ (Condition (A6), Milgrom and Roberts (1990).

[7] A strategy is serially undominated if it survives iterated elimination of dominated strategies.

Milgrom and Roberts (1990)).[8]

One obvious corollary is that any supermodular game has a pure-strategy equilibrium. Another corollary says that if the supermodular game Γ_o has a unique pure-strategy equilibrium (i.e., if the smallest and largest equilibrium points coincide), then Γ_o is dominance solvable. These characteristics also tend to facilitate learning. They represent another piece of evidence suggesting that supermodular games are "learning-friendly".

Finally, Theorem 8, Milgrom and Roberts (1990), and its two corollaries, state that for any adaptive dynamic for a supermodular game Γ_o, the players' choices will eventually lie between the smallest and the largest pure-strategy equilibria of Γ_o.

Vives (1990) showed that the analysis of supermodular games can be extended to Bayesian games obtained from normal-form games by introducing multiple player types. The result consists of two parts. First, if the set of strategies for any given type is a compact subset of reals,[9] then any player's strategy space is a complete lattice (i.e., it satisfies Assumption 1 of Definition 3.6), provided the type space satisfies some regularity conditions.[10] The second step is to show that supermodularity is preserved under integration.

3.2.4 Two-Player Games with Strategic Substitutes

Games with strategic substitutability are characterized by decreasing best replies, decreasing differences of the payoff functions, or *submodular* payoff functions. For completeness, I provide a definition of a *smooth submodular game* below.

Definition 3.7. *Suppose that a typical strategy for player n is $(x_{ni}; i = 1, \ldots, k_n) \in \mathbb{R}^{k_n}$ and that $\geq$ is the usual componentwise ordering. Then Γ_o is a* supermodular *game if, for each $n \in N$:*

[8] Since the strategy space is assumed to be a *partially* ordered set, in general, not all (equilibrium or non-equilibrium) strategy profiles can be compared, and therefore the largest and smallest elements of a strategy subset are not always well defined concepts (i.e., they do not always exist). For example, consider a two-player game, where each player's strategy set is an interval with a natural order of real numbers: $X_n = [0, 1]$, $n = 1, 2$. The *product order*, frequently assumed as the partial order of X, is defined as follows: $x \geq x'$ if for all n $x_n \geq x'_n$. Thus, $(1, 1) \geq (0, 0)$ and $(1, 0) \geq (0, 0)$, but the relationship between (1,0) and (0,1) is not defined. Theorem 5 by Milgrom and Roberts (1990) says that the concept of the smallest and the largest rationalizable strategy is always well defined in supermodular games. Moreover, because all pure-strategy equilibria are always correlated equilibria, which in turn are always rationalizable strategies, the fact that the smallest and largest rationalizable strategies are pure-strategy equilibria implies that the three sets have identical bounds.

[9] More generally, Vives (1990) assumes that strategy spaces are compact lattice subsets of Euclidean spaces.

[10] The type sets are assumed to be non-empty complete separable metric spaces.

(B1) X_n *is an interval in* $\mathbb{R}^{k_n}$*, that is,*

$$X_n = [\underline{x}_n, \overline{x}_n] = \{x | \underline{x}_n \leq x \leq \overline{x}_n\};$$

(B2) π_n *is twice continuously differentiable on* X_n*;*

(B3) $\frac{\partial^2 \pi_n}{\partial x_{ni} \partial x_{nj}} \leq 0$ *for all* n *and all* $1 \leq i < j < k_n$*;*

(B4) $\frac{\partial^2 \pi_n}{\partial x_{ni} \partial x_{mj}} \leq 0$ *for all* $n \neq m$*,* $1 \leq i < k_n$ *and* $1 \leq j < k_m$*.*

Games with strategic substitutes do not have the robust dynamic stability properties as games with strategic complements except for the two-player case. Specifically, two-player submodular games can be reformulated as supermodular games by simply reversing the usual order of one player's strategy space. All theoretical results on supermodular games therefore apply to two-person submodular games. In practice, however, whether the game can be treated as supermodular may depend on players' perception of the game. In some games, there are multiple natural ways to order the strategy sets. For example, in a multiperiod arms race game the strategic choices can be ordered by the stock of arms held in each period or by the periodic level of investment in new armaments. As Milgrom and Roberts (1990) show,[11] the former game is supermodular, while the latter is not. It is reasonable to expect that over time the players can learn to see the game in its former formulation (or both) and therefore perceive it as supermodular. Vives (1990) showed that a variety of Cournot duopoly games are supermodular if one of the two players' strategy sets is given in the reverse of its usual order, but not the natural order. Here, to apply the theory of supermodular games, one has to rely on the firms to treat each others' productions asymmetrically and order them in a different way, but the idea that firms may evolve to perceive the game as supermodular or act as if it is supermodular is plausible. In an abstract game such as Example 3.2.2, however, the players should have some insight into game theory to do the trick in order to help their learning. Therefore, it is plausible that they will not perceive such a context-free submodular game as supermodular. I rely on this conjecture in the design of my experimental study (Section 3.4).

3.2.5 Potential Games

A second class of games with robust dynamic stability properties is the class of potential games.

[11] See pp. 1256, 1272–1274.

Definition 3.8. *A function $P : X \to \mathbb{R}$ is a* potential *for Γ, if for every $n \in N$ and for every $x_{-n} \in X_{-n}$*

$$\pi_n(x_n, x_{-n}) - \pi_n(x'_n, x_{-n}) = P(x_n, x_{-n}) - P(x'_n, x_{-n}),$$

for every $x, x' \in X_n$. Γ is called a potential game *if it admits a potential.*

It follows directly from the definition that the equilibrium set of a potential game coincides with that of a game in which all players' payoff functions are replaced with the potential function.[12] Thus, if players jointly maximize the potential function P, they end up in an equilibrium. The following is an example of a potential game and its potential function.

Example 3.2.3. *Consider the game in Example 3.2.1. The first player's payoff function is given by $\pi_1(x_1, x_2) = -x_1^2 + 6x_1x_2$. Redefine the second player's payoff function as follows: $\pi_2(x_1, x_2) = -2x_2^2 + 6x_1x_2$. This game is potential, and the potential function is given by $P = -x_1^2 - 2x_2^2 + 6x_1x_2$.*[13] *It is maximized at the origin.*

Unfortunately, I cannot offer any straightforward economic interpretation of the potential function, nor could Monderer and Shapley, who introduced potential games (Monderer & Shapley, 1996b).

Anderson, Goeree, and Holt (2001) generalize the definition of a potential to a broader class of games and define it, in loose terms, as a function of all players' decisions, which increases with unilateral changes that increase a player's payoffs, so that any Nash equilibrium is a stationary point of the potential function. Their intuition behind potential is that "if each player is moving in the direction of higher payoffs, each of the individual movements will raise the value of the potential, which ends up being maximized in equilibrium."

The class of learning processes for which positive results in potential games have been obtained is in fact broader than the best-reply dynamic. Monderer and Shapley (1996a) showed that every *fictitious play* process converges to equilibrium in finite potential games. Loosely speaking, fictitious play can be thought of as the best-reply dynamic in which the players play their best reply to the complete history of observed strategy profiles by other players, rather than to the last observation. Monderer and Shapley (1996b) showed that in any *bounded* potential game (with possibly infinite strategy sets), if players follow a "weaker" version of the best-reply

[12] The result in fact holds for a larger class of all finite *ordinal* potential games (see Lemma 2.1, Monderer and Shapley (1996b))

[13] Potential functions are unique up to a constant (Monderer & Shapley, 1996b).

dynamic, they will eventually come arbitrarily close to equilibrium. By a "weaker" version I mean that the players are assumed to "better-reply" rather than best-reply.

To summarize the results formally, I will need to introduce new notation. A *path* in X is a sequence $x = (x(t))_{t=1}^{\infty}$ of elements of X. A *belief path* associated with time t and a path $x(t)$ is a sequence $f^x(t) = \frac{1}{t}\sum_{i=1}^{t} x(t)$ in Δ.

Definition 3.9. *A path $x = (x(t))_{t=1}^{\infty}$ is a* fictitious play process *if for every $n \in N$*

$$x_n(t+1) = BR_n(f_{-n}^x(t)),$$

where BR_n is n's best reply.

Corollary 2.2, Monderer and Shapley (1996b), states that every finite potential game possesses a pure-strategy equilibrium.[14] Theorem 2.4, Monderer and Shapley (1996b), states that in every finite potential game, every fictitious play process converges in beliefs to an equilibrium.[15]

Although the proof of Theorem 2.4, Monderer and Shapley (1996b), is quite involved, the intuition behind these results can be illustrated for a simpler "better-reply" learning dynamic in a game with a finite strategy space X.[16]

Consider a sequence $\gamma = (x^0, x^1, \ldots), x^i \in X$ such that for every $k \geq 1$ there exists a unique player, say player n, such that $x^k = (x_{-n}^{k-1}, x_n)$ for some $x_n \neq x_n^{k-1} \in X_n$. A sequence $\gamma = (x^0, x^1, \ldots), x^i \in X$ is an *improvement path* with respect to Γ if for all $k \geq 1$ $\pi_n(x^k) > \pi_n(x^{k-1})$, where n is the unique deviator at step k. It follows from the definition of the potential game that $P(x^k) > P(x^{k-1})$ for all $k \geq 1$. Thus, for every improvement path $\gamma = (x^0, x^1, \ldots), x^i \in X$ we have $P(x^0) < P(x^1) < \ldots$. That is, the potential function is strictly increasing along any improvement path γ. As the strategy space X is a finite set, the sequence γ must be finite.[17] This demonstrates why the existence of potential implies that learning processes are unlikely to cycle in loops. If the players move in the direction of improving their expected payoff, making predictions based on the observed history of play, they will eventually come to a point where the potential is maximized, which by definition means that no player can further unilaterally improve her payoff. In other words, they will come to a pure-strategy equilibrium. An obvious corollary is that a pure-strategy equilibrium must exist.

[14]The result in fact holds for all finite *ordinal* potential games, which is a larger class of games than finite potential games (see Monderer and Shapley (1996b)).

[15]The result in fact holds for all finite *weighted* potential games, which is a larger class of games than finite potential games (see Monderer and Shapley (1996b)).

[16]Monderer and Shapley (1996b) use similar logic to extend the analysis to bounded potential games with possibly infinite strategy sets.

[17]This result is a slightly modified proof of Lemma 2.3, Monderer and Shapley (1996b), which states that it holds for a broader class of all finite *ordinal* potential games.

A similar result holds for any bounded potential game. The game Γ is a *bounded game* if the payoff functions π_n, $n \in N$, are bounded. This includes games with infinite strategy sets, not necessarily continuous. Monderer and Shapley extend the concept of an improvement path to that of an *ϵ-improvement path*, which requires that the unique deviator improves her payoff by at least an ϵ at each step, where ϵ is an arbitrarily small positive number. This concept is well defined for games with finite as well as infinite strategy sets. In any bounded potential game, if players follow a path with such properties, their choices will eventually lie arbitrarily close to an equilibrium (Lemmas 4.1 and 4.2, Monderer and Shapley (1996b)). Lemma 4.3, Monderer and Shapley (1996b), states that any continuous potential game with compact strategy sets possesses a pure-strategy equilibrium.

The following two theorems help detect potential games when the payoff functions are differentiable and the strategy sets are intervals of real numbers.

Theorem 3.10. *(Lemma 4.4, Monderer and Shapley (1996b)). Let Γ be a game in which the strategy sets are intervals of real numbers. Suppose the payoff functions $\pi_n : X_n \to \mathbb{R}$, $n \in N$, are continuously differentiable, and let $P : X \to \mathbb{R}$. The P is a potential for Γ if and only if P is continuously differentiable and*

$$\frac{\partial \pi_n}{\partial x_n} = \frac{\partial P}{\partial x_n} \quad \textit{for every } n \in N.$$

Theorem 3.11. *(Theorem 4.5, Monderer and Shapley (1996b)).*[18] *Let Γ be a game in which the strategy sets are intervals of real numbers. Suppose the payoff functions $\pi_n : X_n \to \mathbb{R}$, $n \in N$, are twice continuously differentiable. Then Γ is a potential game if and only if*

$$\frac{\partial^2 \pi_n}{\partial x_n \partial x_m} = \frac{\partial^2 \pi_m}{\partial x_n \partial x_m} \quad \textit{for every } n, m \in N. \tag{3.2}$$

Like the class of games with strategic complementarities, the class of potential games possess robust dynamic stability properties. Monderer and Shapley (1996a) showed that every *fictitious play* process converges to equilibrium in finite potential games. Myopic learning process based on one-sided better reply dynamics converges to the equilibrium set. Furthermore, Blume (1993) shows that a log-linear (noisy) learning process converges to the arg max set of the potential.

While Monderer and Shapley (1996b) analyze potential games of complete information, Facchini et al. (1997) extend Rosenthal's congestion model (Rosenthal, 1973) to an incomplete information setting and show that the related Bayesian games are potential games. Relevant to the experiment, Facchini et al. (1997) show that Bayesian potential game with inconsistent priors need not have a pure strategy BNE.

[18] The original theorem also provides a method of constructing the potential function based on the players' payoff functions.

3.2.6 Equilibrium Stability and the Common-Prior Assumption

Bayesian games can be thought of as games with a move by Nature, which draws a type for each player before the actual game. As we know from Section 3.2.1, the type distributions (which do not have to be the same for all players) are common knowledge, while the types are privately revealed to the players. Dekel et al. (2004) raise the question about the limitations of the common-knowledge assumption and address issues arising when players are learning about the distribution of Nature's move (i.e., the type distributions) as well as learning about opponents' strategies.

Dekel et al. (2004) do not formally model the dynamics of learning, but they appeal to the idea that a steady state of a learning process based on the players' beliefs about opponents' play should be a self-confirming equilibrium (SCE). In a sense, this is an equilibrium concept that incorporates some notion of stability with respect to learning processes that adapt based on the observed history of play. It turns out that the players' prior beliefs about the type distribution are critical for the learning outcome. Below I reproduce the notation, the definition of the self-confirming equilibrium, and the intuition behind the concept as provided by the authors.[19]

The key components of self-confirming equilibrium are: each player n's *beliefs* about Nature's move, each player's strategy, and each player's *conjecture* about the strategies used by the other players. Player n's beliefs, $\hat{\mu}_n$, are interpreted here as a point in the space $\Delta(\theta)$ of distributions over Nature's move, and the player's conjectures about opponents' play are assumed to be a $\hat{\delta}_{-n} \in \times_{-n}\Delta_{-n}$, that is, a (mixed-) strategy profile of n's opponents. The notation $\hat{\mu}_n(\cdot|\theta_n)$ refers to the conditional distribution corresponding to $\hat{\mu}_n$ and θ_n, while $\delta_n(x_n|\theta_n)$ denotes the probability that $\delta_n(\theta_n)$ assigns to x_n. Suppose also that $p(\theta_n)$ is the true distribution of the Nature's move and that after each play of the game, players receive private signals $y_n = y_n(x, \theta)$, representing information about Nature's move and the other players' moves. In other words, these are the players' private observations of others' play.

Definition 3.12. *A strategy profile δ is a* self-confirming *equilibrium (SCE) with conjectures $\hat{\delta}_{-n}$ and beliefs $\hat{\mu}_n$ if for each player n,*

(i) for all θ_n with $p(\theta_n) \neq 0$, $\hat{\mu}_n(\theta_n) = p(\theta_n)$,

and for any pair θ_n, $\hat{x}$ such that $\hat{\mu}_n(\theta_n) \cdot \delta_n(\hat{x}|\theta_n) > 0$ both the following conditions are satisfied

(ii) $\hat{x} \in \arg\max_{x_n} \sum_{x_{-n},\theta_{-n}} \pi_n(x_n, x_{-n}, \theta_n, \theta_{-n})\hat{\mu}_n(\theta_{-n}|\theta_n)\hat{\delta}_{-n}(x_{-n}|\theta_{-n})$,

[19] Dekel et al. (2004), pp. 286–287.

and for any $\overline{y}_n$ in the range of y_n

(iii)

$$\sum_{\{x_{-n},\theta_{-n}:y_n(\hat{x}_n,x_{-n},\theta_n,\theta_{-n})=\overline{y}_n\}} \hat{\mu}_n(\theta_{-n}|\theta_n)\hat{\delta}_{-n}(x_{-n}|\theta_{-n}) = \sum_{\{x_{-n},\theta_{-n}:y_n(\hat{x}_n,x_{-n},\theta_n,\theta_{-n})=\overline{y}_n\}} p(\theta_{-n}|\theta_n)\delta_{-n}(x_{-n}|\theta_{-n}).$$

Condition (i) is a consequence of the assumptions that players observe their own types, and the types are independently identically distributed over time. Condition (ii) says that any action played by a type of player n that has positive probability is a best reply to her conjecture about the other players' play and beliefs about Nature's move. Condition (iii) says that the distribution of signals (conditional on type) that the player expects to see equals the actual distribution.

The key result is in Proposition 2 (below), which gives rather restrictive conditions sufficient for the sets of SCE and Bayesian Nash equilibria to coincide, and a number of examples showing that the sets do not coincide when the assumptions of the proposition are violated. Together these results illuminate some limitations of the common-prior assumption about the type distributions in Bayesian games and offer an explanation of some experimental evidence that learning outcomes may differ from theoretical predictions. In the experiment, I do not attempt to test these theoretical results and design the treatments to satisfy the assumptions of Proposition 2. I reproduce Proposition 2 below.

Proposition 3.13. *(Proposition 2, Dekel et al. (2004).) If either*

1. *payoffs are generic (i.e., $\pi_n(x,\theta) \neq \pi_n(x',\theta')$ if $x \neq x'$ or $\theta \neq \theta'$) and observed, or*

2. *there are private values $\pi_n(x,\theta) \neq \pi_n(x',\theta_n)$ and observed actions,*

then the set of strategy profiles in self-confirming equilibria coincides with the set of Nash equilibrium profiles of the game with the correct (hence common) prior.[20]

3.3 Experimental Literature

While there has emerged a large experimental learning literature in the past decade, most of this literature is focused on learning Nash equilibrium under complete information. I refer the reader to Camerer (2003) and Erev and Haruvy (2008) for

[20]BNE in a Bayesian game is, using the terminology of Dekel et al. (2004), a Nash equilibrium of a simultaneous-move game preceded with Nature's move, in which the player's have the correct common prior over the distribution of Nature's move.

comprehensive surveys. In what follows, I summarize the experimental literature on learning in Bayesian games, games of strategic complementarity, and potential games.

In games of strategic complementarity, the experimental results differ in two classes of environments. First, when Nash equilibrium is Pareto optimal, learning leads to robust convergence to Nash equilibrium (Chen and Tang (1998) and Chen and Plott (1996)). Furthermore, this robust convergence result extends to near-supermodular games (e.g., Chen and Gazzale (2004)). When Nash equilibrium is not Pareto optimal, however, empirical results contradict theoretical predictions. At the heart of this contradiction is the assumption of rational behavior. Based on a series of human-subject experiments, Fehr and Tyran argue that under strategic complementarity, a small amount of individual irrationality may lead to large deviation from aggregate predictions of rational models (Fehr and Tyran (2005), Fehr and Tyran (2004b), Fehr and Tyran (2004a), and Fehr and Tyran (2001)). Under strategic substitutes, on the other hand, a minority of rational agents may suffice to generate outcomes consistent with the predictions of rational models. Fehr and Tyran's examples of individual irrationality include money illusion (nominal vs. real incomes) and base rate fallacy (the Monty Hall problem). Intuitively, the authors point out, the cost of a "mistake" is low under strategic complements, while the cost for others of not responding to irrational behavior is high. Therefore, in games with strategic complements, rational predictions are not achieved or achieved more slowly.

Similarly to this study, in the money illusion experiment, the equilibrium is not Pareto optimal. This characteristic apparently gives rise to irrational[21] behavior by some subjects, who favor Pareto-superior albeit non-equilibrium choices. Such behavior can be explained by either personal incentives to collude or by tensions between individual and social preferences.[22]

Potters and Suetens (2006) study collusion in supermodular and non-supermodular environments in which the equilibrium is not Pareto optimal. The payoff functions induced in Potters and Suetens (2006) are very similar to those induced in this experiment. The authors report results qualitatively similar to those by Fehr and Tyran: the level of convergence to the equilibrium in the supermodular environment was much lower than the theory of supermodular games suggests. Rather than learning the equilibrium, the subjects converge to collusive outcomes. While potential games with multiple equilibria, especially order statistics games, have been studied extensively in the experimental literature (Camerer (2003),

[21] By irrational I mean behavior that does not constitute best reply in the context of one-shot simultaneous-move games.

[22] Under social preferences, players tend to prefer *socially* beneficial outcomes.

Chapter 7), a few papers investigate learning in congestion games under different information conditions. Chen (2003) reports an experimental study of the serial and the average cost pricing mechanisms under different information conditions. Although the proportion of Nash-equilibrium play under both mechanisms is statistically indistinguishable under complete information, the serial mechanism performs robustly better than the average cost pricing mechanism under limited information, both in terms of the proportion of equilibrium play and system efficiency. When the environment is more complex, with four types of players, Chen, Razzolini, and Turocy (2007) show that the serial mechanism performs significantly better than the average-cost-pricing mechanism in all treatments in terms of efficiency, predictability measured as frequency of equilibrium play, and the speed of convergence to equilibrium.

3.4 Experimental Design

I design the experiment to explore conditions that affect people's learning of BNE in games of incomplete information. I vary game environments by supermodularity of the players' payoff functions, existence of potential, and by the amount of type-distribution information available to the players. In all treatments, the game has a unique BNE. I therefore do not address the question of equilibrium selection in Bayesian games in this study. In addition, in each treatment the equilibrium is globally asymptotically stable with respect to the Cournot best-reply dynamic (see Section 3.2.2). Also, the games all satisfy condition 2 of Proposition 2 by Dekel et al. (2004) (see Section 3.2.6), which, given the uniqueness of the BNE, implies that the equilibrium is also a *self-confirming equilibrium* in each treatment. Finally, the equilibrium in this experimental study is not Pareto optimal. As I discuss in Section 3.3, incentives to deviate in games where the equilibrium is not Pareto optimal may significantly affect the learning outcome. I was not able to ensure that the equilibrium is "equally Pareto non-optimal" across treatments, but I statistically control for the incentives to deviate in the analysis. The specific environments and the experimental procedures are discussed in the sections below.

3.4.1 Strategic Environments

I chose a generalized version of the Cournot duopoly as the economic environment for the experiment. The Cournot competition game can be supermodular, submodular, or neither; potential or not depending on the coefficients of the players' payoff functions. There are two players, the row (r) and the column (c), each of two types. The row player can be either red (r) or blue (b) with equal probability.

The column player can be either green (g) or purple (p) with equal probability. The players simultaneously choose a number from 0 to 21. Let x_n be player n's choice and x_m be player m's choice, where $n, m \in \{r, c\}$, $n \neq m$. The payoff function of player n of type θ_n, where $\theta_r \in \{r, b\}$ and $\theta_c \in \{g, p\}$, and given m's choice, is

$$\pi_n^{\theta_n} = a_n^{\theta_n} + b_n^{\theta_n} x_n - d_n x_n^2 + f_n x_m x_n, \tag{3.3}$$

where $a_n^{\theta_n} \geq 0$, $b_n^{\theta_n} > 0$, $d_n > 0$, and f_n is any real number.

The player's expected payoff function over player m's choices is therefore given by

$$E\pi_n^{\theta_n} = a_n^{\theta_n} + b_n^{\theta_n} x_n - d_n x_n^2 + f_n E(x_m) x_n. \tag{3.4}$$

The game is an example of a well known linear Cournot duopoly with uncertain constant marginal costs.[23] The difference in the strategic behavior of different types is defined by the parameter $b_n^{\theta_n}$. The constant $a_n^{\theta_n}$ does not affect the player's best-reply function, and I define it to be type specific only to ensure that subjects earn similar payoffs in expectation.

The cross-derivative of the expected payoff function is simply the parameter f_n:

$$\frac{\partial^2 E\pi_n^{\theta_n}}{\partial x_n \partial x_m} = f_n. \tag{3.5}$$

By Definitions 3.5, 3.7, and Theorem 3.11, the parameters f_n determine whether the game is supermodular, submodular, or neither; whether the game has a potential or not.[24] In Table 3.1, I lay out all the cases.

[23] See, for example, Fudenberg and Tirole (1991), pp. 215–216, or Vives (2000), pp. 225–226. In the latter, the expected payoff function is given by $E(\pi_i|c_i) = (a - c_i)x_i - x_n^2 - Ex_m(c_j)x_i$, where the parameter c_n is firm n's private information (type) representing the constant marginal cost of production, and a is the intercept of the linear demand curve. In Equation (3.3), the parameter defining the player's type, $b_n^{\theta_n}$, is the difference of the two parameters.

[24] Technically, Definitions 3.5 and 3.7 apply to normal-form rather than Bayesian games. However, it can be shown that each of the four normal-form games obtained by fixing the type profile is supermodular and that the type sets are non-empty complete separable metric spaces. Then the Bayesian game is supermodular according to Vives (1990). Similarly, Theorem 3.11 assumes that the game is a normal form. However, if $f_n = f_m = f$, it is possible to construct a function satisfying the original Definition 3.8 for expected payoffs. Let

$$P = \sum_{n \in N, \theta_n \in \Theta_n} (a_n^{\theta_n} + b_n^{\theta_n} x_n - d_n x_n^2) + f x_m x_n.$$

Let us fix the strategy and the type of player m and consider player n of type θ_n.

$$P(x_n, \cdot) - P(x_n', \cdot) = \sum_{n \in N, \theta_n \in \Theta_n} (b_n^{\theta_n}(x_n - x_n') - d_n(x_n^2 - (x_n')^2) + f(\cdot)(x_n - x_n') = \pi_n^{\theta_n}(x_n, \cdot) - \pi_n^{\theta_n}(x_n', \cdot).$$

Since the equality holds for every type of player m, it holds for the expected utility of player n over m's types:

	Not Potential	Potential
Submodular	$f_n < f_m < 0$	$f_n = f_m < 0$
Supermodular	$0 < f_n < f_m$	$0 < f_n = f_m$
Neither ("Mixed")	$f_n < 0 < f_m$	

Table 3.1: Strategic environments as defined by the parameters f_n.

The best-reply function of player n of type θ_n is given by

$$BR_n^{\theta_n} = \frac{b_n^{\theta_n}}{2d_n} + \frac{f_n}{2d_n}\overline{x}_m, \tag{3.6}$$

where $\theta_r \in \{r, b\}$, $\theta_c \in \{g, p\}$, and $\overline{x}_m$ is player m's expected choice over player m's types.

The (unique) internal BNE choice of player n of type θ_n is given by

$$x_n^{*\theta_n} = \frac{1}{2d_n}\left(b_n^{\theta_n} + \frac{f_n(f_m\overline{b}_n + 2d_n\overline{b}_m)}{4d_nd_m - f_nf_m}\right), \tag{3.7}$$

where $\theta_r \in \{r, b\}$, $\theta_c \in \{g, p\}$, and $\overline{b}_n$ is the expected $b_n^{\theta_n}$ over θ_n. For a two-type model with equally likely types, the expected type simplifies to

$$\begin{aligned} \overline{b}_r &= \frac{b_r^r + b_r^b}{2}, \\ \overline{b}_c &= \frac{b_c^g + b_c^p}{2}. \end{aligned} \tag{3.8}$$

3.4.2 Experimental Treatments

In Table 3.2, I summarize the main features of the experimental treatments. Two supermodular-potential treatments were dropped from the design during data collection. Based on the results obtained for the other six treatments, it became clear that the supermodular-potential treatments are unlikely to provide new information.

	Low Information		Bayesian Information	
	Not Potential	Potential	Not Potential	Potential
Submodular	SubNP$_l$	SubP$_l$	SubNP$_b$	SubP$_b$
Supermodular	SuperNP$_l$	—	SuperNP$_b$	—

Table 3.2: Features of experimental treatments

$$P(x_n, \cdot) - P(x_n', \cdot) = E\pi_n^{\theta_n}(x_n, \cdot) - E\pi_n^{\theta_n}(x_n', \cdot).$$

For each strategic environment in Table 3.2, I present in Table 3.3 the choice of specific payoff functions for the experiment. To minimize the possibility of a confounding factor, I use the same payoff functions across environments whenever possible. To define the environments, four different payoff functions for the row player and six for the column player were required. I label the corresponding payoff functions by $\pi_{n_k}^{\theta_n}$, where $n \in \{r, c\}$, k is the preference model subscript, $k \in \{1, 2, 3\}$, and θ_n is player n's type.

Environment	Color Type	Payoff Functions		Slope	Equilibrium
SubNP	red	$\pi_{r_1}^r$	$118x_r - 2.5x_r^2 - 3x_r x_c - 50$	-0.6	17
	blue	$\pi_{r_1}^b$	$88x_r - 2.5x_r^2 - 3x_r x_c + 90$	-0.6	11
	green	$\pi_{c_1}^g$	$214x_c - 5x_c^2 - 6x_r x_c - 250$	-0.6	13
	purple	$\pi_{c_1}^p$	$174x_c - 5x_c^2 - 6x_r x_c - 50$	-0.6	9
SubP	red	$\pi_{r_2}^r$	$236x_r - 5x_r^2 - 6x_r x_c - 750$	-0.6	17
	blue	$\pi_{r_2}^b$	$176x_r - 5x_r^2 - 6x_r x_c - 130$	-0.6	11
	green	$\pi_{c_1}^g$	$214x_c - 5x_c^2 - 6x_r x_c - 250$	-0.6	13
	purple	$\pi_{c_1}^p$	$174x_c - 5x_c^2 - 6x_r x_c - 50$	-0.6	9
SuperNP	red	$\pi_{r_3}^r$	$104x_r - 5x_r^2 + 6x_r x_c - 700$	0.6	17
	blue	$\pi_{r_3}^b$	$44x_r - 5x_r^2 + 6x_r x_c - 200$	0.6	11
	green	$\pi_{c_2}^g$	$23x_c - 2.5x_c^2 + 3x_r x_c + 200$	0.6	13
	purple	$\pi_{c_2}^p$	$3x_c - 2.5x_c^2 + 3x_r x_c + 200$	0.6	9

Table 3.3: Payoff functions. Notation: x_r and x_c are choices made by the row and column players, respectively; $\pi_{n_k}^{\theta_n}$ is the payoff function of player n of type θ_n, where $n = \{r, c\}$, $\theta_r = \{r, b\}$, $\theta_c = \{g, p\}$, and k is the preference model subscript, $k = \{1, 2, 3\}$. The BNE solutions for each type are given in the last column. The numbers in the column labeled "Slope" are the slopes or the best-reply functions.

The payoff functions in Table 3.3, together with the strategy spaces equal to the intervals $[0, 21]$ and the players' type distributions, define specific Bayesian games. In Appendix D, I present graphs of the best-reply functions for each game. I present a graphical representation of the players' preferences and strategic properties of the environments in Appendix E. The games have the following properties:

1. There is a unique internal BNE in each game.
2. The internal BNE solutions are the same across all games.
3. The internal BNE solutions are all integers.
4. The absolute value of the slope of the best-reply function is the same across environments and types. This condition is important for comparability of the

speed of convergence to the equilibrium in different environments. In general, the number of iterations it takes, for example, fictitious-play learners to get from an off-equilibrium play to an equilibrium depends on the steepness of the best-reply function.

5. There are no obvious "focal-point" equilibrium choices: The row and column players have different payoff functions; their equilibrium choices are also different from each other.

6. The internal equilibrium is a globally asymptotically stable steady state of the continuous Cournot tatonnement.[25]

7. There are no corner solutions.[26]

8. To mitigate fairness concerns, I chose the endowment coefficient, $a_n^{\theta_n}$, in such a way that on average all players would have similar earnings at the end of the experiment. The equilibrium earning ranges from 475 to 585 across players, with the average equal to 526.[27]

9. The strategy space intervals are such that all Nash equilibria of the normal-form games obtained by fixing a type profile are interior points (see Table 3.4). I discuss the rationale behind this property in Appendix D.

There are two information conditions for each set of strategic environments: Bayesian information and low information. Under the Bayesian-information condition, the players have a common correct prior about the distribution of their types, and that prior is common knowledge. In other words, the condition satisfies the theoretical type-information assumption of Bayesian games. Under the low-information condition, minimum possible information about the number of player

[25] It can be shown that the eigenvalues of the coefficient matrix of the simultaneous equations defining the continuous tatonnement Equation (3.1) for the Cournot environments are equal to $-1 \pm \sqrt{\frac{f_1}{2d_1}\frac{f_2}{2d_2}}$. The expression under the square root is the product of the slopes of the two best-reply functions. If their absolute values are less than one, the real part of both eigenvalues is negative, which implies global asymptotic stability of the steady state of the dynamic system. As Table 3.3 shows, the best replies have slopes equal to either -0.6 or 0.6. Therefore, the internal equilibrium in all the treatments is globally asymptotically stable with respect to the Cournot tatonnement.

[26] This follows from the global stability of the internal equilibrium with respect to the Cournot best-reply dynamic: at any point on the boundary of the strategy space, at least one of the players has an incentive to deviate toward the internal equilibrium.

[27] To predict the average earning, I computed the mean of the equilibrium payoff and the random-choice expected payoff. The latter was obtained by simulating the game play for 5000 periods with players programmed to use the random-choice strategy. I used the mean of the equilibrium and random-choice payoffs as my predicted average subject earnings. Actual type-average earnings varied around my prediction.

Treatment	$\{r, g\}$	$\{r, p\}$	$\{b, g\}$	$\{b, p\}$
SuperNP	20.56, 16.94	16.81, 10.69	11.19, 11.31	7.44, 5.06
SubNP	16.81, 11.31	20.56, 5.06	7.44, 16.94	11.19, 10.69
SubP	16.81, 11.31	20.56, 5.06	7.44, 16.94	11.19, 10.69
MixNP	17.09, 11.15	15.32, 8.21	12.68, 13.79	10.91, 10.85
SuperP	20.56, 16.94	16.81, 10.69	11.19, 11.31	7.44, 5.06

Table 3.4: Nash equilibria of the normal-form games obtained by fixing a type profile.

types, type distribution, and others' payoffs is revealed to the players. This condition is designed to approximate extremely limited information available to players interacting with each other through online systems. Under both conditions, the players observe the history of play, which includes their own actions, the actions of the other player, and their own payoffs. The choice to reveal both players' actions was made in order to satisfy condition 2 of Proposition 2 by Dekel et al. (2004) (see Section 3.2.6).

Before collecting data, I tested the experimental treatments using software agents programed to follow an established learning model capturing the principles of human learning behavior. I report the results in Appendix F. The simulated dynamics vary little across environments, consistent with theoretical predictions based on Cournot-tatonnement stability, yet the human behavior in the experiment below varies significantly.

3.4.3 Experimental Procedures

I programmed and implemented the experiment with z-Tree (Fischbacher, 2007).[28] I conducted three independent sessions for each treatment (a total of 18 sessions). The treatments are summarized in Table 3.2. Each session had eight subjects. A total of 144 subjects participated in the experiment.[29] All subjects were University of Michigan undergraduate and graduate students. I conducted all sessions at the School of Information Experimental Laboratory at the University of Michigan in February–April, 2008. The subjects were paid according to their performance in the experiment. Complete instructions are provided in Appendices G–I. Experimental data are available upon request.

[28] The official cite of z-Tree is located at `http://www.iew.unizh.ch/ztree/index.php`.

[29] One subject participated in two different sessions: first in a low-information session and a month later in a Bayesian-information session of the SuperNP environment. The subject was a row player in the former and a column player in the latter, i.e., the subject was never assigned the same payoff functions. The results do not change significantly when the subject is dropped from the analysis of the second session.

For each treatment, I conducted experimental sessions consisting of 100 rounds of the corresponding Bayesian game (the payoff functions are presented in Table 3.3). The instructions were read aloud to the subjects at the beginning of the experiment. Then the program randomly assigned a type (red, blue, green, or purple) to each subject according to the type distribution. The types remained the same for the first 50 rounds. At the beginning of round 51, red types became blue types and *vice versa*; green types became purple types, and *vice versa*. Thus, rows remained rows for the entire experiment, but each row player played as both row types, for 50 rounds each. Similarly for column players. This helped alleviate fairness concerns that some types would have higher payoffs than others. The subjects were randomly rematched in each round (rows were matched with columns). The type or identity of the match was never revealed.

For a baseline comparison, I conducted a total of nine sessions with SubNP, SubP, and SuperNP environments under the Bayesian-information condition. Before the first round, the players were informed about the structure of the game, the matching protocol, the type distribution and the payoffs of all types. The information about the type distribution was given in the instruction. The information about the payoffs was given in the form of payoff tables (see an example in Appendix H). In addition, the subjects could compute their payoff and the payoff of either type of the other player using a computerized tool called "What-If-Scenario Analyzer". The subjects learned how to use the tables and the analyzer in a computerized tutorial. At the beginning of round 51, the subjects were notified of their type change and reminded that the new type would remain for the rest of the experiment. I designed the Bayesian-information treatments to compare convergence to BNE in the tree one-shot games (SubNP, SubP, and SuperNP) when the assumption about the players' information exactly matches the theoretical assumption.

In addition to the instruction, the subjects in Bayesian-information sessions had a 15-minute computerized tutorial, in which they were trained on how to read the payoff tables and use the What-If-Scenario Analyzer. They also had a chance to practice using the Analyzer and compare its calculations to the entries of the payoff tables. The subjects' understanding of the instructions and the tutorial material was tested with a computerized quiz. The subjects were rewarded for correctly answering the questions. Most subjects answered all questions correctly. The lowest number of questions answered correctly was 12 out of 16. I explained or clarified the correct answers to all subjects who made a mistake or had questions during the quiz. There was no practice round in any of the sessions. The average time spent on the instructions, tutorial and quiz was 33 minutes, and the entire session took 102 minutes on average.

To measure convergence in settings in which the Bayesian-information assumption is relaxed (such as in online applications), I designed a low-information condition. Online applications are often distributed: participants are geographically dispersed, and their information is limited by what is provided through the application interface. Such information can be extremely limited, sometimes to the extent that the players may not even know their own type, if type depends on some exogenous parameters not revealed by the application;[30] the relationship between others' actions and own payoffs may also be obscured. The low-information condition in my experiment is designed to represent this extreme case. The only information the players had was their own action, the action of the other player, and the players' own payoffs. Neither the What-If-Scenario Analyzer nor the payoff tables were available, and no information about the subjects' own or others' types was revealed. Also, the subjects were not notified about their type change before round 51, but they were informed in the instruction that the environment may change and that they would not be notified when that happens. Since the instructions for the low-information case were straightforward, I did not test the subjects' understanding with a quiz. There was no practice round either. The average instruction time was 7 minutes, and the entire session took 52 minutes on average.

In all treatments, the strategy space was discretized as multiples of 0.01. In each round, the subjects made their choices simultaneously and independently of other players. Throughout the experiment, each subject's computer screen displayed the history of the subject's own choices, own earnings, and the choices of the subject's matches.

The exchange rate for all treatments was one dollar for 1,700 points.[31] The average earning per hour was $23, including a $5 show-up fee.

3.5 Hypotheses

Based on the theory, results from prior experimental work, and the design, I next identify my hypotheses. To do so, I first define and discuss two measures of convergence: the convergence level and the rate of improvement (ROI). In theory, convergence implies that all players play the stage-game equilibrium and no deviation is observed. This is not realistic, however, in an experimental setting. I therefore

[30]See, for example, an experimental study by Chen et al. (2007), in which the authors evaluate the performance of two different mechanisms for congestion allocation in distributed networks. The players' types are defined by the topology of the network and their location in that network. However, such information may not be available or costly to obtain in online networks.

[31]The exception is the first SubNP session, where the exchange rate was one dollar for 3,300 and the coefficients $a_n^{\theta_n}$ were all about two time higher than those in the following sessions.

define the level of convergence more generally as distance from equilibrium payoff ($\mathrm{PDist}_{\mathrm{n,t}}$).

Definition 3.14. *Let $\pi_n(t)$ be player n's actual payoff in round t and π_n^* be n's equilibrium payoff. Then the* distance from equilibrium payoff *in round t is given by*

$$\mathrm{PDist}_{\mathrm{n,t}} \equiv |\pi_{\mathrm{n}}(\mathrm{t}) - \pi_{\mathrm{n}}^*|. \tag{3.9}$$

A lower $\mathrm{PDist}_{\mathrm{n,t}}$ means a better degree of convergence. Alternative measures commonly used in experimental literature include distance between actual and equilibrium choices and proportion of ϵ-equilibrium play. The latter can be defined in terms of payoffs or choice units. Distance from equilibrium choice is a more direct measure of the convergence level than distance from equilibrium payoff. In some cases, however, it is more meaningful to consider an action to be close to the equilibrium if it yields near-equilibrium payoff, in particular when players' strategy spaces are not ordered. When the payoff function is single-peaked, as in this experiment, the two types of distances are roughly equivalent measures. The findings of this study do not appear to be sensitive to the choice of the distance measure, but the payoff-based measure gives a better regression-analysis fit.

Proportion of ϵ-equilibrium play punishes all non-ϵ-equilibrium choices equally, which is considered undesirable by some authors. Another issue is that the choice of ϵ may not be obvious, which is the case in this experiment. Analysis of the data using (1) proportion of payoff-based ϵ-equilibrium play with ϵ equal to 10% of average equilibrium payoff and (2) proportion of choice-based ϵ-equilibrium play with to 3 (choice units) supports the results presented here, but some regression coefficients and the tests have lower statistical significance and the standard errors tend to be higher. The analysis of the alternative measures is available upon request, but I do not report it here.

Ideally, the ROI should measure how quickly all players converge to equilibrium strategies. However, since convergence is never perfect in this experiment, I instead measure how quickly the players *approach* the equilibrium. In particular, I measure ROI as change in PDist per unit of time. I estimate this change using regression analysis in which I include a time variable (ln(Round)) and its interactions with treatment indicator variables (Supermod, Potential, and Info). The coefficients on these variables are interpreted as a change in PDist when the round counter changes by 1%.

I now formally state my null hypotheses. Recall that treatments with better convergence would have smaller average PDist. Since the Bayesian-information condition fully matches the theoretical informational assumption of Bayesian games, while the low-information condition is an extreme case violating this assumption, I hypothesize that convergence to BNE is better under the former:

Hypothesis 1 (Effect of Information on Convergence Level)**.** *In each environment, the distance from equilibrium payoff is smaller under the Bayesian-information condition than under the low-information condition, i.e.,*

(i) $\mathrm{PDist_b(SubNP)} < \mathrm{PDist_l(SubNP)}$,

(ii) $\mathrm{PDist_b(SubP)} < \mathrm{PDist_l(SubP)}$,

(iii) $\mathrm{PDist_b(SuperNP)} < \mathrm{PDist_l(SuperNP)}$,

where the subscripts b and l label the Bayesian-information and the low-information conditions, respectively.

Hypothesis 2 (Effect of Information on Convergence ROI)**.** *In each environment, convergence is faster under the Bayesian-information condition than under the low-information condition, i.e., the coefficient on* ln(Round) *is more negative in Bayesian information treatments.*

My hypotheses about the effect of potential relies on the theoretical studies predicting robust convergence in potential games:

Hypothesis 3 (Effect of Potential on Convergence Level)**.** *Under each information condition, the distance from equilibrium payoff is smaller in potential games than in the non-potential counterparts, i.e.,*

(i) $\mathrm{PDist_l(SubP)} < \mathrm{PDist_l(SubNP)}$,

(ii) $\mathrm{PDist_b(SubP)} < \mathrm{PDist_b(SubNP)}$,

where the subscripts l and b label the low-information and the Bayesian-information conditions, respectively.

Hypothesis 4 (Effect of Potential on Convergence ROI)**.** *Under each information condition, convergence is faster in potential games than the non-potential counterparts, i.e., the coefficient on* ln(Round) *is more negative in* SubP *than in* SubNP.

Based on the literature reviewed in Section 3.2, supermodularity is an important convergence factor. However, experimental evidence is mixed about the direction of the effect. I therefore formulate the following hypotheses about supermodularity:

Hypothesis 5 (Effects of Supermodularity on Convergence Level)**.** *Under each information condition, the distance from equilibrium payoff in supermodular games is different from the distance in the submodular counterparts, i.e.,*

(i) $\mathrm{PDist_l(SuperNP)} \neq \mathrm{PDist_l(SubNP)}$,

(ii) $\mathrm{PDist_b(SuperNP)} \neq \mathrm{PDist_b(SubNP)}$,

where the subscripts l and b label the low-information and the Bayesian-information conditions, respectively.

Hypothesis 6 (Effect of Supermodularity on Convergence ROI)**.** *Under each information condition, convergence is different in supermodular and non-supermodular counterparts, i.e., the coefficient on* ln(Round) *is different in supermodular and non-supermodular treatments.*

3.6 Results

In Figures 3.1–3.3, I present the time series of type-specific choice data in each treatment. Each figure consists of eight graphs, comparing participant behavior in the low-information and Bayesian-information conditions for each environment. In each graph, I present the average choices in each round (the dots), the error bars which are one standard deviation from the average, and the BNE prediction (the grey line). The figures indicate that both the information conditions and the strategic environment affect the learning dynamics.

As prior experimental work shows, incentives to collude may have a significant effect on convergence dynamics in games where the equilibrium is not Pareto optimal. This means that in this experiment, the distance from the equilibrium payoff may be (negatively) correlated with the distance from Pareto improvements over the equilibrium. To control for the effect in the analysis below, I define the following variables.

Definition 3.15. *Let x_n be player n's choice and x_m be player m's choice, where $n, m \in \{r, c\}$. Let θ_1 and θ_2 be m's two types. Let $\pi_n(\cdot)$ and $\pi_m^{\theta_T}(\cdot)$, $T \in \{1, 2\}$, be the players' respective payoffs and π_n^* and $\pi_m^{*\theta_T}$ be their equilibrium payoffs. The* joint weighted payoff maximum *for player n is given by*

$$\begin{aligned} \mathrm{JWPM_n} \equiv \max_{x_n, x_m} & \left(\pi_n(x_n, x_m) + \frac{\pi_m^{\theta_1}(x_n, x_m) + \pi_m^{\theta_2}(x_n, x_m)}{2} \right), \\ \text{s.t.} \quad & \pi_n(x_n, x_m) \geq \pi_n^*, \\ & \pi_m^{\theta_1}(x_n, x_m) \geq \pi_m^{*\theta_1}, \\ & \pi_m^{\theta_2}(x_n, x_m) \geq \pi_m^{*\theta_2}. \end{aligned}$$

Low Information Bayesian Information

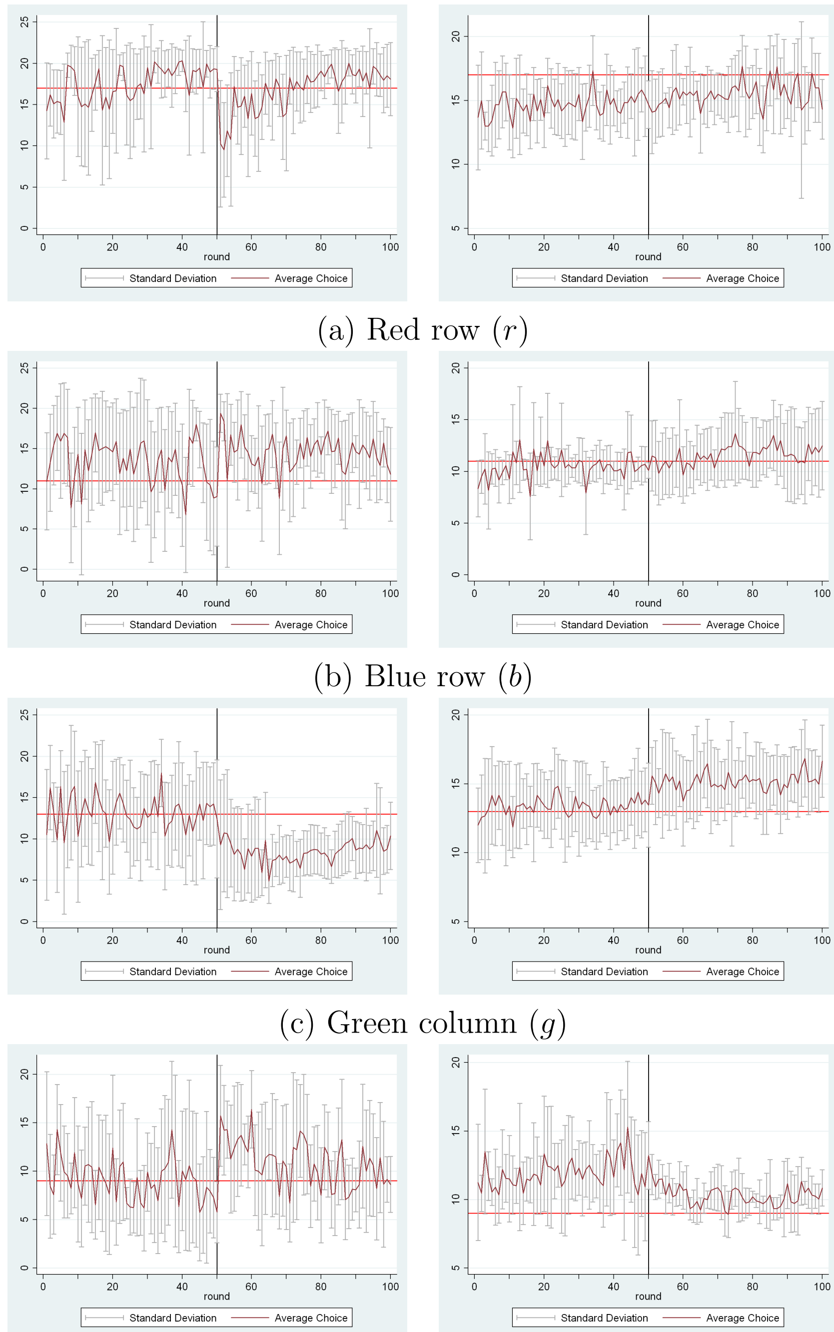

(a) Red row (r)

(b) Blue row (b)

(c) Green column (g)

(d) Purple column (g)

Figure 3.1: Convergence dynamics: Submodular & non-potential environment. The straight horizontal line is the equilibrium choice of the corresponding type. The straight vertical line at round 50 indicates the time of the type change. Thus, each type is represented by two subject pools: one in the first 50 rounds and the other in the last 50 rounds.

Low Information   Bayesian Information

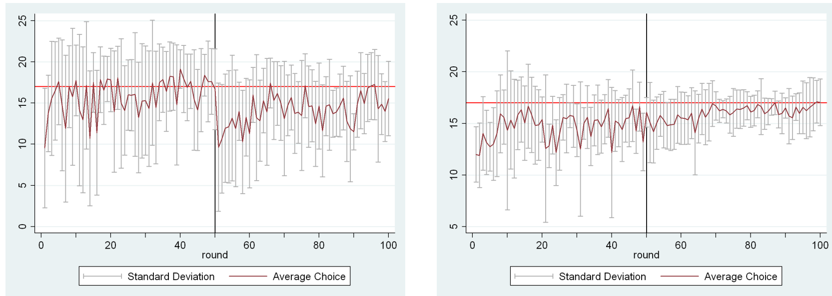

(a) Red row (r)

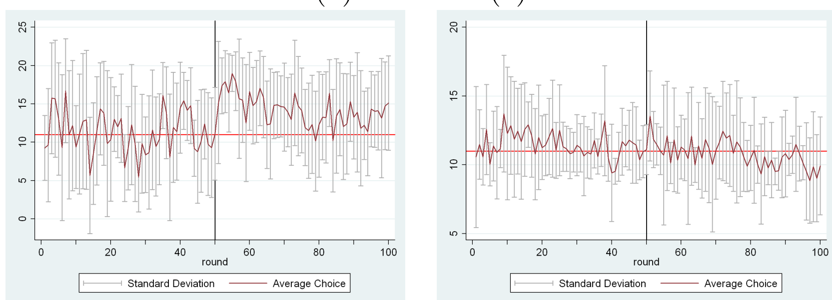

(b) Blue row (b)

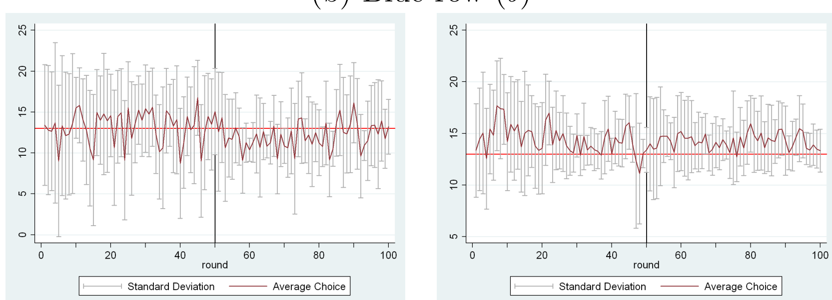

(c) Green column (g)

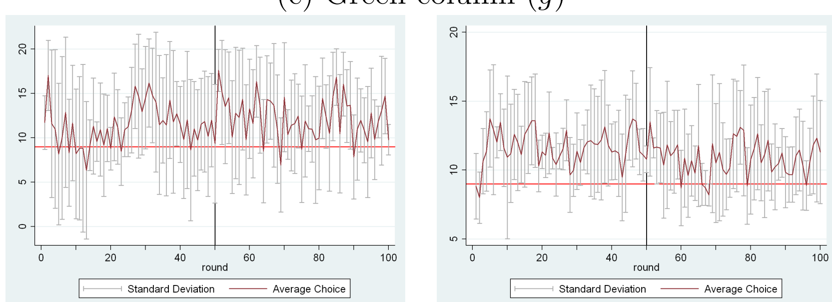

(d) Purple column (g)

Figure 3.2: Convergence dynamics: Submodular & potential environment. The straight horizontal line is the equilibrium choice of the corresponding type. The straight vertical line at round 50 indicates the time of the type change. Thus, each type is represented by two subject pools: one in the first 50 rounds and the other in the last 50 rounds.

Low Information　　Bayesian Information

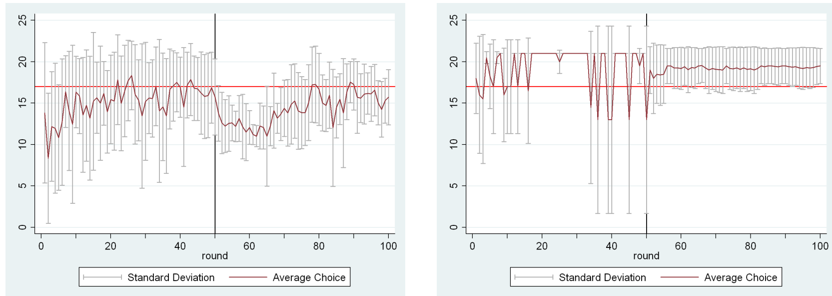

(a) Red row (r)

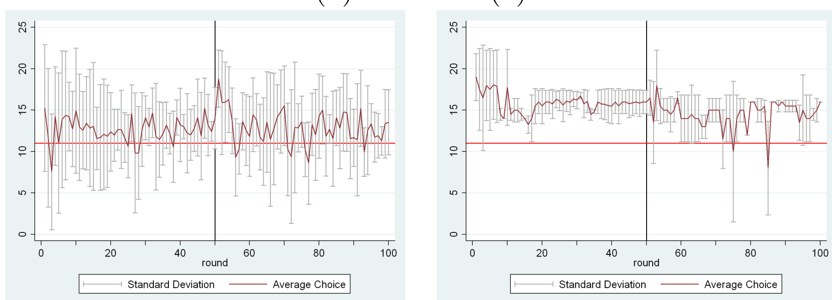

(b) Blue row (b)

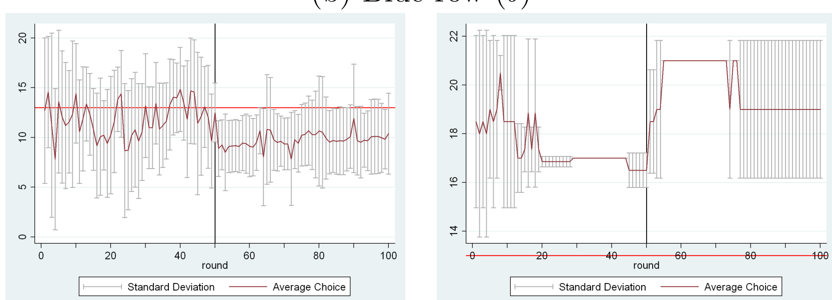

(c) Green column (g)

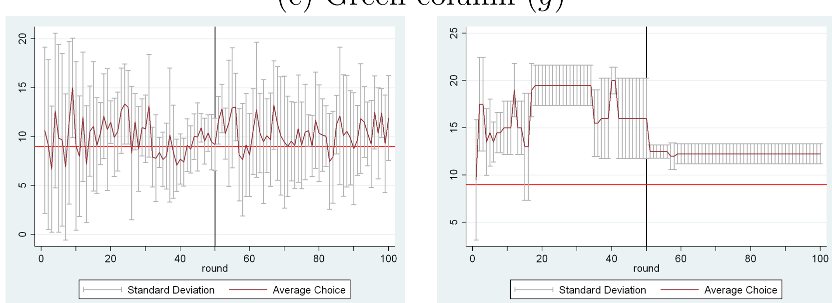

(d) Purple column (g)

Figure 3.3: Convergence dynamics: Supermodular & non-potential environment. The straight horizontal line is the equilibrium choice of the corresponding type. The straight vertical line at round 50 indicates the time of the type change. Thus, each type is represented by two subject pools: one in the first 50 rounds and the other in the last 50 rounds.

In Table 3.5, I report the values of JWPM, as well as JWPM choices for all types in all environments.

Definition 3.16. *Player n's* incentives to collude *in round t is the distance between n's actual payoff at round t-1, π_n(t-1), and n's* JWPM*:*

$$\mathrm{I2C_{n,t}} \equiv |\mathrm{JWPM_n} - \pi_\mathrm{n}(\mathrm{t}-1)|.$$

Another important variable measures average risk to the players of incurring big losses in an environment.

Definition 3.17. *Let $\pi^{\theta_n}_{\min}$ be the minimum possible payoff of player $n \in \{r, c\}$, of type θ_n, where $\theta_r \in \{r, b\}$, $\theta_c \in \{g, p\}$, in an environment* $\mathrm{E} \in \{\mathrm{SubNP}, \mathrm{SubP}, \mathrm{SuperNP}\}$*. The* average loss bound *in* E *is given by*

$$\mathrm{LossBound} = \frac{1}{4}(\pi^{\mathrm{r}}_{\min} + \pi^{\mathrm{b}}_{\min} + \pi^{\mathrm{g}}_{\min} + \pi^{\mathrm{p}}_{\min}).$$

I did not fully anticipate the effect of negative payoffs on subjects' behavior when designing the experiment, and therefore the measure is somewhat *ad hoc*. However, as I argue below based on statistical analysis of the data, it captures a characteristic of the environment whose effect overrides that of a potential. I report the value of LossBound for each environment in Table 3.5.

I first compare the convergence level achieved in the last ten rounds of each 50-round block in each treatment. Table 3.6 presents the average distance from BNE payoff in points and as percentage of the average payoff range over all types in all environments (2259 points). Using permutation tests of PDist between pairs of treatments, I test my original hypotheses.[32] As can be seen from Table 3.7 (top panel), only two of the seven relationships from the original hypotheses are confirmed (at the 5% level of statistical significance).

[32]Permutation tests are typically used for small samples to determine whether the observed difference between the sample means is large enough to reject the null hypothesis that the two groups have identical probability distribution. The data points in each sample are assumed to be independent. In my experiment, the average PDist in a session satisfies this assumption. Therefore, for the purpose of this test, I have a sample of size three for each treatment. If we pool the samples of two treatments of interest together and then make multiple random draws of three data points, we can obtain a distribution of the means of samples of size three. The idea of the test is that under the null hypothesis that the two groups have identical probability distribution, the actual means we observe (i.e., the overall averages in the corresponding treatment columns in Table 3.6) should be close to the mean of the distribution statistic. The test can be one or two-sided. For more details on the test, see, for example, Johnston and DiNardo (1996), Section 11.2.

Table 3.5: Values of JWPM, JWPM choices and LossBound. The payoffs are rounded to the nearest integer. The choices are rounded to two decimal places.

Environment	Color Type			
	Red	Green	Blue	Purple
JWPM				
SubNP	704	393	772	615
SubP	866	516	760	390
SuperNP	1925	1165	904	484
JWPM Choices				
SubNP	11.53	7.37	9.41	8.34
SubP	12.4	7.75	9.3	5.84
SuperNP	21	21	21	21
LossBound				
SubNP	-598			
SubP	-972			
SuperNP	-866			

Result 1 (Effect of Information on Convergence Level)**.**

*(i) In both submodular environments (*SubNP *and* SubP*), the distance from equilibrium payoff is smaller under the Bayesian-information condition than under the low-information condition, i.e.,*

$$\mathrm{PDist_b(SubNP)} < \mathrm{PDist_l(SubNP)},$$
$$\mathrm{PDist_b(SubP)} < \mathrm{PDist_l(SubP)},$$

where the subscripts b and l label the Bayesian-information and the low-information conditions, respectively.

*(ii) In the potential environment (*SubP*), the distance from equilibrium is smaller under the Bayesian-information condition than under the low-information condition (see the previous item).*

*(iii) In the supermodular environment (*SuperNP*), the difference in* PDist *under the Bayesian-information and low-information conditions is not statistically significant.*

Support: Table 3.6 presents the average distance from equilibrium payoff for each session. In Table 3.7, I also report the permutation tests for PDist. The tests confirm Parts (i) and (ii)(of Hypothesis 1 at $p = 0.05$; they reject Part (iii) of the hypothesis ($p = 0.9$ and $p = 0.15$ for two one-sided tests).

Table 3.6: Average distance from equilibrium payoff. The payoffs and percentages are rounded to the nearest integer.

Average Distance from Equilibrium Payoff (Points)						
All Rounds	Low Information			Bayesian Information		
Session	SubNP	SubP	SuperNP	SubNP	SubP	SuperNP
1	279	554	241	167	246	336
2	337	422	234	152	283	175
3	332	406	297	199	226	308
Overall	316	461	257	173	252	273
Last 10 Rounds	Low Information			Bayesian Information		
Session	SubNP	SubP	SuperNP	SubNP	SubP	SuperNP
1	256	579	193	217	195	309
2	317	346	175	166	268	187
3	354	279	236	189	245	289
Overall	309	401	201	191	236	262
Average Distance from Equilibrium Payoff (Percent of Average Payoff Range)						
All Rounds	Low Information			Bayesian Information		
Session	SubNP	SubP	SuperNP	SubNP	SubP	SuperNP
1	12	25	11	7	11	15
2	15	19	10	7	13	8
3	15	18	13	9	10	14
Overall	14	20	11	8	11	12
Last 10 Rounds	Low Information			Bayesian Information		
Session	SubNP	SubP	SuperNP	SubNP	SubP	SuperNP
1	11	26	9	10	9	14
2	14	15	8	7	12	8
3	16	12	10	8	11	13
Overall	14	18	9	8	10	12

Table 3.7: Monte Carlo permutation tests: Average distance from equilibrium payoff in the last 10 rounds of 50-round blocks. The results are based on 1,000,000 samples.

Effect	Original Hypotheses	p-value
Information by treatment	$\mathrm{PDist_b(SubNP)} < \mathrm{PDist_l(SubNP)}$	0.05**
	$\mathrm{PDist_b(SubP)} < \mathrm{PDist_l(SubP)}$	0.05**
	$\mathrm{PDist_b(SuperNP)} < \mathrm{PDist_l(SuperNP)}$	0.9
Supermodularity: LIC	$\mathrm{PDist_l(SuperNP)} \neq \mathrm{PDist_l(SubNP)}$	1.0
Supermodularity: BIC	$\mathrm{PDist_b(SuperNP)} \neq \mathrm{PDist_b(SubNP)}$	0.15
Potential: LIC	$\mathrm{PDist_l(SubP)} < \mathrm{PDist_l(SubNP)}$	0.8
Potential: BIC	$\mathrm{PDist_b(SubP)} < \mathrm{PDist_b(SubNP)}$	0.95

*Note: Significant at: * 10-percent level; ** 5-percent level; *** 1-percent level.*

Result 1 provides the first empirical evidence on the role the Bayesian-information assumption plays in learning BNE. It shows that the information condition affects learning in the intuitive direction predicted in the original hypothesis, with the exception of the supermodular environment, in which we observe no significant difference. Result 2 offers more insight into the interaction effect between supermodularity and the information condition.

I further investigate the effect of the treatment variables – information, supermodularity, and potential – on the convergence level and ROI – using regression analysis (Table 3.8). Unlike in permutation tests, one can control for confounding variables in a regression model. In my regression analysis, I control for I2C and LossBound, which I defined above. The regression analysis also provides insight into the interactions between dependent variables and informs a new set of convergence-level hypotheses.

I now briefly explain the regression models presented in Table 3.8. I use the subject ID as the panel variable. The Wooldridge test for autocorrelation in panel-data models[33] rejects the null hypothesis that there is no first-order autocorrelation at the 0.001% level (p-value = 0.0000 and the $F_{1,143}$ statistic equals 107.719). I therefore estimate a GLM model with the first-order[34] autoregression correlation structure within group using the GEE method (Yee Liang & Zeger, 1986) implemented in Stata 10.[35] Without the autoregression correlation assumption, the model is asymptotically equivalent to the weighted-GLS estimator (GLS with random effects) and to the full maximum-likelihood random-effects estimator. In

[33] Wooldridge (2002), pp. 282–283.

[34] Using the second-order autoregression correlation structure does not significantly affect the coefficients and standard errors, although the latter tend to be larger with AR(2).

[35] `http://www.stata.com/stata10/`. The regression command is `xtgee` with the option `corr(ar1)`.

balanced data, the model produces the same results as the maximum-likelihood random-effects estimator.[36] To estimate the standard errors, I use a version of the Huber/White/sandwich estimator that allows for correlation within subjects and produces valid standard errors even if the correlations within group are not as hypothesized by the specified correlation structure.[37] I use ln(Round) as the time variable. Usually, the variation in choices decreases with time, which suggests a logarithmic transformation of the round variable. This appears to be the case in this experiment.

Specification (1) reflects the original hypotheses. It includes all of the treatment variables and their interactions, as well as the I2C control variable. Specifications (2)–(3) offer an alternative set of regressions, which help explain a puzzling relationship that specification (1) reveals. In specification (1), the estimated coefficient on the potential indicator variable is positive. That is, we observe the opposite of the effect stated in Hypothesis 3, which is based on theoretical results suggesting that potential games have robust convergence properties. It is even more puzzling that the coefficient on ln(Round)×Potential is nonetheless negative, implying that subjects approach equilibrium faster in potential games than in non-potential, as originally hypothesized. In an attempt to find an explanation, I re-examined the design for possible confounding variables. The most plausible confound appeared to be the difference in the minimum possible payoffs, which are negative for all types in all treatments.[38] There is an extensive literature on loss aversion originated with a seminal paper by Tversky and Kahneman (1979), in which the authors argue that people tend to strongly prefer avoiding losses to acquiring gains. Two observations indicated that such a tendency may have had an effect on PDist. First, the SubP and SubNP environments are almost identical, with identical best-response functions and indifference-curve maps. The only difference is the row player's payoff functions, whose non-constant coefficients in SubP are twice those in SubNP. As a result, the

[36] StataCorp (2007), p. 126.

[37] The Stata option for the `xtgee` command is `vce(robust)`. Other maintained assumptions are (1) that the model correctly specifies the mean, and (2) that errors between groups are independent (StataCorp (2007), p. 121). The latter assumption is not satisfied in my experimental data. Ideally, errors should be grouped at the session rather than individual level. However, such an option is not available for the `xtgee` command in Stata 10. Therefore, the reported standard errors may be underestimated. However, due to the high significance of the regression coefficients (Table 3.8), I believe that the results reported here are valid.

[38] I introduced the possibility of earning negative payoffs in order to make the task look more challenging for the subjects and encourage higher effort on their part. However, I was not able to ensure that the range of possible payoffs on [0,21] was the same across different types and at the same time keep other important characteristics constant across the environments. I tried to minimize the risk of a confound by adjusting $a_n^{\theta_n}$ in the payoff functions in such a way that all types could incur significant losses and the types with bigger possible losses had somewhat higher expected payoffs. However, as my analysis shows, this measure proved to be insufficient.

payoff variation is higher for the row player in SubP, and the minimum possible payoff is more negative. Second, from Figure 3.2 we can see that the choice dynamics in SubP has a very distinct pattern, especially pronounced under Bayesian information: the subjects do not experiment as much as in other environments. Such behavior may be indicative of more cautiousness on the subjects' part, which in turn could be due to exposure to larger losses.

To test the conjecture, I defined the LossBound variable. Unfortunately, this variable is perfectly collinear with a linear combination of the Potential and Supermodular indicator variables and the constant. Thus they cannot each be estimated simultaneously. In specification (2), I estimate a model with LossBound, and with Potential dropped. The coefficient on LossBound is significant at the 1% level, while the constant term becomes small (-5.184) and statistically insignificant (the standard error is 73). I therefore suppress the constant to be able to estimate the coefficient on Potential, as well as the rest of the original treatment variables, while controlling for LossBound. The resulting model is specification (3). Note that (a) the coefficient on LossBound is the same as in specification (2): a one-point increase in the average loss bound results in a 0.6 decrease in the distance from equilibrium payoff; it is significant at the 1% level, as in specification (2); and (b) the sizes of the coefficients on both Potential and Supermodular are quite different from those in specification (1), which suggests that LossBound is in fact a confounding variable.[39] I summarize the above in the following observation.

Observation 1 (Effect of Average LossBound on Convergence Level)**.** *Distance from equilibrium payoff decreases as the average loss bound increases.*

Support: Specifications (2)–(3) show that a one-point increase in the average loss bound results in a 0.6 decrease in the distance from equilibrium payoff. The coefficients on LossBound are statistically significant at 1%.[40]

In light of the above, it is clear that the simple permutation tests in Table 3.7 may not accurately reflect the effects of potential and supermodularity on the average distance from equilibrium payoff. I therefore rely on specification (3) in the rest of the analysis. Two of the coefficients in specification (3) (Potential and ln(Round)×Supermodular) are statistically insignificant. However, I keep the variables in the model because they are variables of interest. If these variables are removed, the rest of the coefficients are not significantly affected, and all the results presented below remain valid. In Table 3.9, I present regression coefficients by

[39] The rest of the treatment coefficients remain unchanged.

[40] The interaction effect between LossBound and Info could not be identified due to collinearity problem.

treatment and report their statistical significance based on Wald tests (Table 3.10). Tables 3.8 and 3.9 serve as the main support of the following results and observations.

Two notes before I present the next result. First, since in my design there is no treatment that is both a potential and supermodular game, we can ignore coefficients on Potential and its interactions when analyzing supermodular treatments and *vice versa*.

Second, ROI may be negative in this experiment, i.e., convergence does not always improve with time. More specifically, we have the following result.

Result 2 (Effect of Information on Initial Choices and ROI)**.**

(i) [Initial Choices] The initial distance from equilibrium payoff is smaller under Bayesian information than under low information.

(ii) [ROI] In each environment, the Bayesian-information condition reduces convergence ROI relative the low-information treatments. More specifically:

(a) [ROI, LIC] Under the low-information condition, the distance from equilibrium payoff reduces with time in each environment.

(b) [ROI, BIC] Under the Bayesian-information condition, there are two cases:

*i. In the potential environment (*SubP*), the distance from equilibrium payoff also reduces with time, but at a slower rate than under low information.*

*ii. In the non-potential environments (*SuperNP *and* SubNP*), the distance from equilibrium payoff does not change or even increases with time.*

Support: First consider initial choices, whose distance from equilibrium payoff is measured by the intercepts. Specification (3) in Table 3.8 shows that Bayesian information reduces the initial PDist by 237.2 points in SubNP, by 299.67 ($-237.2-62.47$) points in SubP, and by 64.2 ($-237.2+173$) points in SuperNP. These differences constitute, respectively, 10.5%, 13.3%, and 2.8% of the average payoff range, which is equal to 2259 points. I next examine the coefficients on ln(Round), which measure convergence ROI, under different information conditions. Under low information, the distance from equilibrium payoff decreases by 0.27 ($\frac{-27.32}{100\%}$) points with every one-percent increase in the round counter. Under Bayesian information, consider first the potential environment (SubP). The coefficient on ln(Round) shifts up by 7.36 ($38.28-30.92$), but it remains negative: $-27.32+7.36=-19.96$. Thus, under Bayesian information in SubP, the distance from equilibrium payoff decreases by 0.2 ($\frac{-19.96}{100\%}$) points with every one-percent increase in the round counter. In

Dependent variable: Distance from equilibrium payoff (PDist)			
	(1)	(2)	(3)
ln(Round)	-27.32***	-27.32***	-27.32***
	(5.753)	(5.753)	(5.753)
Info	-237.2***	-237.2***	-237.2***
	(27.59)	(27.59)	(27.59)
ln(Round)×Info	38.28***	38.28***	38.28***
	(5.963)	(5.963)	(5.963)
I2C	0.114***	0.114***	0.114***
	(0.0118)	(0.0118)	(0.0118)
Supermodular	-70.17**	-230.7***	-228.4***
	(31.11)	(30.32)	(38.90)
Supermodular×Info	173.0***	173.0***	173.0***
	(38.66)	(38.66)	(38.66)
Potential	224.7***		3.244
	(36.52)		(45.68)
Potential×Info	-62.47**	-62.47**	-62.47**
	(29.43)	(29.43)	(29.43)
ln(Round)×Potential	-30.92***	-30.92***	-30.92***
	(7.869)	(7.869)	(7.869)
ln(Round)×Supermodular	-8.455	-8.455	-8.455
	(5.929)	(5.929)	(5.929)
LossBound		-0.600***	-0.592***
		(0.0976)	(0.0424)
Constant	353.8***	-5.184	
	(25.35)	(73.00)	

Notes: Semi-robust standard errors are in parentheses.
*Significant at: * 10% level; ** 5% level; *** 1% level.*

Table 3.8: GLM model with an autoregression correlation structure within group. The panel variable is the subject ID. The standard errors are reported as semi-robust, because the estimator requires that the model correctly specifies the mean. The errors are adjusted for clustering at the individual level. The ln(Round) variable is the logarithm of the round count. The count was reset before round 51, when subjects were assigned a new type. The Info indicator variable equals one for the Bayesian-information condition and zero otherwise. The Supermodular indicator variable equals one in SuperNP and zero in both submodular environments (SubNP and SubP). The Potential indicator variable equals one in SubP and zero in both non-potential environments (SuperNP and SubNP). Excluded is the submodular non-potential indicator variable, which equals one in SubNP and zero in SuperNP and SubP.

Coefficients by treatment relative to baseline (SubNP)					
		Info=0		Info=1	
		P=0	P=1	P=0	P=1
Intercept	S=0	baseline	3.244	-237.2	-296.426
	S=1	-228.4	—	-292.6	—
ln(Round)	S=0	baseline	-30.92	38.28	7.36
	S=1	-8.455	—	29.825	—
Actual coefficients by treatment					
		Info=0		Info=1	
		P=0	P=1	P=0	P=1
Intercept	S=0	0	3.244	-237.2***	-296.426***
	S=1	-228.4***	—	-292.6***	—
ln(Round)	S=0	-27.32***	-58.24***	10.96**	-19.96***
	S=1	-35.775***	—	2.505	—

*Significant at: * 10% level; ** 5% level; *** 1% level.*

Table 3.9: Coefficients by treatment for specification (3) in Table 3.8: relative to the baseline treatment SubNP (top) and actual (bottom). Info=0 and Info=1 refer to the low- and Bayesian-information conditions, respectively. P=0 and P=1 refer to the non-potential and potential treatments, respectively. S=0 and S=1 refer to the submodular and supermodular treatments, respectively. Some of the actual treatment coefficients are directly estimated according to specification (3) and therefore are simply copied from Table 3.8. Others are obtained by adding two or more of the estimated coefficients. In such cases, statistical significance is determined based on Wald tests (see Table 3.10).

Cell in Table 3.9		Wald Test for Linear Hypotheses	Chi-2(1)	p-value
Intercept	Info=1 P=1 S=0	BayesInfo + Potential + BayesInfo×Potential = 0	39.78	0.0000
Intercept	Info=1 P=0 S=1	BayesInfo + Supermodular + BayesInfo×Supermodular = 0	32.08	0.0000
ln(Round)	Info=0 P=1 S=0	ln(Round) + ln(Round)×Potential = 0	49.34	0.0000
ln(Round)	Info=1 P=0 S=0	ln(Round) + ln(Round)×Info = 0	5.87	0.0154
ln(Round)	Info=1 P=1 S=0	ln(Round) + ln(Round)×Info + ln(Round)×Potential = 0	10.75	0.0010
ln(Round)	Info=0 P=0 S=1	ln(Round) + ln(Round)×Supermod = 0	46.38	0.0000
ln(Round)	Info=1 P=0 S=1	ln(Round) + ln(Round)×Info + ln(Round)×Supermodular = 0	0.25	0.6140

Table 3.10: Statistical significance of the regression coefficients reported in Table 3.9. Statistical significance is determined using Wald tests for linear hypotheses for specification (3) in Table 3.8.

the non-potential environments (SubNP and SuperNP), however, the coefficient on ln(Round) is higher by 38.28 under Bayesian information and equals $-27.32+38.28 = 10.96$. In other words, the distance from equilibrium payoff increases by 0.11 ($\frac{10.96}{100\%}$) points under Bayesian information if the game is not potential.

Part (ii) of Result 2 rejects Hypothesis 2. Two observations may explain this. First, the sizes of the intercept shifts (-64.2 and especially -237.2 and -299.67) suggest that under Bayesian information, almost all of the convergence that could have happened, happened in the initial rounds. In other words, there was not as much room left for the subjects to further converge relative to the low-information condition, under which they had to start with random choices. Since in a linear model the convergence ROI can only be captured as a slope of a straight line, the effect can manifest itself as a smaller slope (in absolute terms) in Bayesian-information treatments. This can explain point i of Part (iib) of the result, but not point ii.

As for the latter, my conjecture is that this is due to Pareto inferiority of the equilibrium (the second observation). Here information about the payoff structure has two effects. One is that in initial rounds, it helps subjects to identify near-equilibrium areas, which is supported by Part (i) of Result 2. However, it also helps to identify Pareto superior outcomes that benefit both players, which may cause them to move away from equilibrium with time rather than even closer to it. Whether that happens or not may depend on the size of the incentives to collude and also on how "risky" the collusive areas are, as well as on the structural characteristics of the payoff functions (i.e., supermodularity and existence of potential). In this experiment – if we maintain my conjecture – the second effect is only observed in SubNP and SuperNP, but not in SubP, which suggests a particular role of potential. Result 5 will shed more light on why the second effect is present in supermodular games with Pareto inferior equilibrium.

To sum up, Result 2 suggests that under low information, the players – quite naturally – start with random choices, gradually approaching the equilibrium neighborhood. Under Bayesian information, subjects generally start playing the outcomes in the neighborhood of the equilibrium. In potential games, they move even closer to it with time. In non-potential games, however, if the equilibrium of the game is Pareto inferior, and especially if both players can gain from deviating, subjects may move away from it toward more profitable profiles.

Next, I present two results on potential games.

Result 3 (Effect of Potential on Convergence ROI)**.**

*(i) Under low information, where convergence to equilibrium improves with time in each environment, convergence is faster in the potential environment (*SubP*)*

*than in the non-potential counterparts (*SubNP *and* SuperNP*).*

*(ii) Under Bayesian information, the potential environment (*SubP*) is the only environment where convergence improves with time.*

Support: First note that the coefficient on the interaction effect ln(Round)×Supermodular is statistically insignificant, which implies that ROI in SubNP and SuperNP is the same (i.e., the difference is not statistically significant). Second, the coefficient on the interaction between Potential and ln(Round) is negative and significant at the 1% level. More specifically, a one-percent increase in the round counter increases the distance from equilibrium payoff by 0.31 ($\frac{30.92}{100\%}$) points faster in SubP than in SubNP and SuperNP. This makes the coefficient on ln(Round) in the potential environment more negative under low information ($-27.32-30.92 = -58.24$) and changes it from positive ($-27.32 + 38.28 = 10.96$) to negative ($-27.32 + 38.28 - 30.92 = -19.96$) under Bayesian information.

Result 3 confirms Hypothesis 4. In fact, the established result is stronger than originally hypothesized: it is the only one of the two structural characteristics that has a robust effect on the rate of improvement under both information conditions.

Result 4 (Effect of Potential on Convergence Level)**.** *Under each information condition, the distance from equilibrium payoff is smaller in potential games than in the non-potential counterparts, i.e.,*

(i) $\mathrm{PDist_l(SubP)} < \mathrm{PDist_l(SubNP)}$*,*

(ii) $\mathrm{PDist_b(SubP)} < \mathrm{PDist_b(SubNP)}$*,*

where the subscripts l and b label the low-information and the Bayesian-information conditions, respectively.

Support: In Result 3, I have shown that PDist reduces faster in SubP than in SubNP. It only remains to show that the intercept of the regression line in SubP does not exceed (is the same or more negative) that in SubNP to have sufficient evidence supporting Result 4. This is in fact so: under low information, there is no statistical difference between the baseline treatment SubNP and the potential treatment SubP: the coefficient on Potential is statistically insignificant. Under Bayesian information, the intercept is 62.47 points lower in SubP.

Result 4 supports Hypothesis 3. This and the above results are the first empirical evidence of the role of potential in learning BNE.

Finally, I discuss supermodular games.

Observation 2 (Effect of Supermodularity on Convergence ROI)**.** *The data provide no evidence that supermodularity affects convergence ROI.*

Support: The interaction effect between supermodularity and ln(Round) is statistically insignificant (Table 3.8).

To relate this observation to the original hypotheses, it means that there is no evidence supporting Hypothesis 6.

Result 5 (Effect of Supermodularity on Convergence Level)**.**

1. *Under each information condition, the distance from equilibrium payoff in supermodular games is smaller in supermodular games than in the non-supermodular counterparts, i.e.,*

   (i) $\text{PDist}_\text{l}(\text{SuperNP}) < \text{PDist}_\text{l}(\text{SubNP})$,

   (ii) $\text{PDist}_\text{b}(\text{SuperNP}) < \text{PDist}_\text{b}(\text{SubNP})$,

   where the subscripts l and b label the low-information and the Bayesian-information conditions, respectively.

2. *There is a strong interaction between supermodularity and the information condition: the effect of information on the average distance from equilibrium payoff is almost cancelled out in the supermodular environment.*

Support: As we already know from Result 3, convergence ROI in SubNP and SuperNP is the same (more precisely, the difference in the coefficients on ln(Round) is not statistically significant in these two environments). It only remains to show that the intercept of the regression line in SuperNP does not exceed (is the same or more negative) that in SubNP to have sufficient evidence supporting Result 5. This is in fact so: under low information, the intercept in SuperNP is 228.4 points (10.1% of the average payoff range) lower than that SubNP. Since supermodularity does not significantly affect the convergence ROI, this implies that distance from equilibrium payoff is on average 228.4 points smaller in SuperNP than in SubNP – under low information. Under Bayesian information, the intercept shifts down only by 55.4 points ($-228.4 + 173 = -55.4$), which is 2.5% of the average payoff range. Again, since the difference between the coefficients on ln(Round) in SuperNP and in SubNP is not statistically significant, I conclude that under Bayesian information, supermodularity decreases PDist on average by 55.4 points. The smaller effect of supermodularity in the Bayesian-information treatments is due to the strong interaction effect between supermodularity and information: the coefficient on ln(Round)$\times$Supermodular equals 173. Another way of stating the result is that the effect of information on the average PDist in SuperNP is only $-237.2+173=-62.2$ compared to -237.2 in SubNP.

Part (1) of Result 5 supports Hypothesis 5 and is consistent with theoretical predictions of the effect of supermodularity on convergence. Part (2) of Result 5 is the first empirical evidence of the striking interaction between strategically important information and the effect of supermodularity in Bayesian games. It suggests that information plays a key role as to whether supermodularity ensures the predictions of models based on the rationality assumptions – such as convergence to an equilibrium theoretically demonstrated by Milgrom and Roberts (1990) and experimentally confirmed by Chen and Tang (1998) and Chen and Plott (1996), among others – or reinforces behavioral biases, such as those described by Fehr and Tyran (2005). Under low information, when subjects do not have the payoff-structure information in front of them, the best they can do is to use unsophisticated heuristics such as those satisfying the definition of Milgrom and Roberts' adaptive dynamic. Therefore, consistently with Milgrom and Roberts (1990), supermodularity drastically improves convergence to the equilibrium. Under Bayesian information, however, the subjects know the payoff structure of all players and are more prone to "irrational" behavior, for example attempts at implicit cooperation in a one-shot game, or behavioral biases, such as loss aversion. In supermodular environments, such biases become reinforced, as argued by Fehr and Tyran (2005), and therefore reduce the effect predicted by the theory that does not model them explicitly.[41]

I now make two additional observations, which are not central research questions in this work, but are nevertheless relevant to the discussion of supermodular and potential games.

Observation 3. *Under each information condition, incentives to collude increase distance from equilibrium payoff.*

Support: Specifications (1)–(3) in Table 3.8 show that the distance from equilibrium payoff increases by 0.114 points per one-point increase in I2C. The coefficient on I2C is statistically significant at 1%. Although not reported here, the coefficient on the interaction between I2C and Info is statistically insignificant.

This observation is consistent with earlier findings of experimental studies of complete-information games.

[41] Since I statistically control for I2C and LossBound, the difference in the supermodularity coefficients under low and Bayesian information is unlikely to be fully due to the difference in potential gains from collusion. The large gap between the coefficients on Supermodular persists under different measures of incentives to collude: I have explored five alternative measures, and observed the phenomenon in all of the alternative models.

Observation 4. *If a non-supermodular game can be reformulated as a supermodular game, that does not ensure supermodular convergence properties in practice.*

Support: As the results above show, the convergence dynamics in the supermodular environment are significantly different from those in the submodular environments. This is true despite, as I point out in Section 3.2.1, the fact that both can be reformulated as supermodular games.

This observation is important for practical applications of the theory of supermodular games and mechanism design. It suggests that the learnability of the induced supermodular game is sensitive to the framing of players' strategies.

3.7 Conclusion

In this chapter, I report experimental results on convergence to Bayesian Nash equilibrium (BNE) in a variety of environments. The environments differ in the payoff structure of the players: they can be (1) a supermodular (and not potential) game, (2) a potential (and not supermodular) game, or (3) neither a potential nor a supermodular game. Each of these environments is studied under two information conditions. The first satisfies the theoretical assumption in Bayesian games. The other simulates online settings by giving the subjects very limited information about the game. Online markets are a challenging area for application of mechanism design theory due to extremely limited information available to the agents. The experimental results suggest that Bayesian information plays an important role for convergence to BNE. Generally, under Bayesian information, subjects' payoffs at early stages are about 12% closer to the equilibrium payoff, where percent is taken with respect to the average payoff range over all types and environments, and the average convergence level is significantly higher than under low information. The exception is the supermodular environment. In the supermodular environment, information improves the average convergence level only by about 3% of the average payoff range. This is consistent with findings of earlier experiments with complete-information games with Pareto inferior Nash equilibrium and with findings by Fehr and Tyran, who argue that supermodularity of the players' payoff functions reinforces mistakes and behavioral biases (I control for incentives to collude, and therefore the observed effect cannot be fully attributed to collusion). Under low information, however, my findings for supermodular games are more consistent with theoretical predictions and findings in experiments with complete-information games with Pareto optimal Nash equilibrium: the players' payoffs are 10% closer to the equilibrium payoff in the supermodular environment, where percent is taken with respect to the average payoff range over all types and environments. I also find that under both information conditions, the rate of convergence improvement is higher in potential games than in the non-potential counterparts. Finally, the subjects' behavior proved to be very sensitive to the size of potential losses. This underlines the importance for convergence to BNE of people's tendency to treat losses differently from gains.

APPENDICES

APPENDIX A

Bidding Strategies for Simultaneous Ascending Auctions: Appendix

In this appendix, we provide details on how we constructed the perceived-price function for each bidding strategy. Most of our bidding strategies are well-defined for any type distribution. Two exceptions are the demand-reduction strategy (Section A.2) and the own-effect price predictor (Section A.3) defined only for homogeneous-good environments.

A.1 Prediction-Based Perceived-Price Strategies

In this section, we report how we constructed initial predictions for each prediction-based strategy. Some of our initial predictions are based on Monte Carlo sampling. To obtain a prediction for a particular distribution of agents' preferences, we simulate a large number of game instances with agents drawn from that preference distribution and count the number of times each final price has occurred in each of the auctions. We refer to the total number of simulated game instances as the sample size. To obtain a distribution prediction of final prices, we divide all positive counts by the sample size. The result is always an m-vector of marginal price distributions. To obtain a point price prediction, which is simply an m-element vector of numbers, we compute the means of the marginal distributions. One exception is the $PP(F(\boldsymbol{\pi}^{EDPE}))$ point prediction, which is based on an expectation over the estimated demand rather than over final prices. See Section 1.5.2 for details.

We assume that prices are bounded by some constant V. For technical reasons, we require that initial distribution predictions are such that for each good, all integer prices in $\{0, \ldots, V\}$ have a positive probability of occurring. To ensure that this requirement is satisfied for initial distribution predictions obtained from empirical samples, we add a one to the counts of all prices from 0 to V before dividing them by the sample size. This increases the sample size by $V+1$. In our empirical analysis,

Strategy & notation	Strategy parameter	Example	Section
Straightforward bidder, SB	N/A	SB	1.3.1
Sunk-aware agent, SA(k)	Sunk-awareness parameter, k	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$	1.3.2
Point price predictor, $PP(\boldsymbol{\pi}^x)$	Predictions about average final prices of the goods, $\boldsymbol{\pi}^x$, where x labels method of generating prices	$PP(\boldsymbol{\pi}^{Zero})$ $PP(\boldsymbol{\pi}^{\infty})$ $PP(\boldsymbol{\pi}^{EPE})$ $PP(\boldsymbol{\pi}^{EDPE})$ $PP(\boldsymbol{\pi}^{SB})$ $PP(\boldsymbol{\pi}^{SC})$	A.1
Point price predictor with participation only, $PP(\boldsymbol{\pi}^x)$ w/ P.O.	Predictions about average final prices of the goods, $\boldsymbol{\pi}^x$, where x labels method of generating prices	$PP(\boldsymbol{\pi}^{\infty})$ w/ P.O. $PP(\boldsymbol{\pi}^{Zero})$ w/ P.O. $PP(\boldsymbol{\pi}^{EPE})$ w/ P.O. $PP(\boldsymbol{\pi}^{EDPE})$ w/ P.O. $PP(\boldsymbol{\pi}^{SB})$ w/ P.O. $PP(\boldsymbol{\pi}^{SC})$ w/ P.O.	A.1
Price distribution predictor, $PP(F^x)$ $PP(G(\mu(x), \sigma(y)))$ $PP(F(\boldsymbol{\pi}^x))$	Predictions about marginal final-price distributions, F, where F is labeled by method of generating prices	$PP(F^{Zero})$ $PP(F^{U})$ $PP(F^{SB})$ $PP(F^{CE})$ $PP(G(\mu(CE), \sigma(CE)))$ $PP(G(\mu(SB), \sigma(SB)))$ $PP(F(\boldsymbol{\pi}^{EPE}))$ $PP(F(\boldsymbol{\pi}^{EDPE}))$ $PP(F(\boldsymbol{\pi}^{SB}))$ $PP(F(\boldsymbol{\pi}^{SC}))$	A.1
Demand-reduction agent, $DR(\kappa)$	Demand-reduction parameter, κ	$DR(\kappa)$ $\kappa = \{1, 2, \ldots 30, 32, 34, 36, 38, 40, 44, 48, 50, 52, 56, 60, 70, 80, 90, 100, 110, 120\}$	A.2
Own-effect price predictor, $OEPP(\boldsymbol{\pi}^x)$	Predictions about own effect on final prices, $\boldsymbol{\pi}^x$, where x labels method of generating prices	$OEPP(\boldsymbol{\pi}^{SB})$	A.3

Table A.1: Strategy Index. P.O. refers to *participation-only* prediction (see MacKie-Mason et al. (2004)).

Environment	V	Num. Profiles	Min	Max	Average
Complementary $E(3,3)$	50	104	0.4	10.4	2.919
Complementary $E(3,5)$	50	462	0.6	0.65	0.643
Complementary $E(3,8)$	50	3023	0.37	10.44	0.453
Complementary $E(5,3)$	50	84	0.6	0.6	0.6
Complementary $E(5,5)$	50	462	0.6	0.6	0.6
Complementary $E(5,8)$	50	3023	0.5	9.61	0.601
Complementary $E(7,3)$	50	84	0.6	0.6	0.6
Complementary $E(7,6)$	50	924	0.6	0.6	0.6
Complementary $U(3,3)$	50	104	0.6	7.6	2.071
Complementary $U(3,5)$	50	462	0.4	6.4	0.72
Complementary $U(3,8)$	50	3023	0.41	6.44	0.48
Complementary $U(5,3)$	50	104	0.6	5.6	1.783
Complementary $U(5,5)$	50	4457	0.2	200.46	7.01
Complementary $U(5,8)$	50	3023	0.6	6.6	0.65
Complementary $U(7,3)$	50	104	0.6	6.61	2.049
Complementary $U(7,6)$	50	944	0.6	6.6	0.759
Complementary $U(7,8)$	50	3023	0.6	4.6	0.629
Substitutable $U(5,5)$	127	16995	0.04	34.485	0.986

Table A.2: Minimum, maximum, and average number of game instances (in millions of samples) generated per strategy profile in complementary and substitutable environments. The number of profiles is a total over all restricted games analysed for the environment.

V is at most 0.3% of the size of the empirical sample, and therefore the effect on the shape of the probability distribution is negligible.

Formally, let G be the number game instances in an empirical sample. If we use this sample to generate an initial distribution price prediction, our marginal-probability estimate that a good will have final price $p \in \{0, \ldots, V\}$ is given by

$$\Pr(p) = \frac{G_p + 1}{G + (V+1)}, \tag{A.1}$$

where G_p is the number of times the final price equals p in the original sample.

In Table A.2 we summarize for each environment the number of profiles, the minimum, maximum, and average number of game instances generated per strategy profile.

In the following sections we describe a few initial price predictions. We denote a specific point price-prediction strategy by $PP(\boldsymbol{\pi}^x)$, where x labels a particular initial point prediction. We denote the strategy based on a particular distribution predictor by $PP(F^x)$, where x labels various initial predictions about final price distributions.

If the initial prediction is based on Monte Carlo sampling, we write x_u and x_e to distinguish between predictions obtained using draws from uniform and exponential preference distributions.[1]

A.1.1 Zero Prediction

Zero *point* prediction is simply an m-element vector of zeros. We denote the strategy by $PP(\boldsymbol{\pi}^{Zero})$. The bidding behavior of $PP(\boldsymbol{\pi}^{Zero})$ is identical to that of SB (defined in Section 1.3.1).

To construct a zero *distribution* prediction, we create for each good an artificial empirical sample in which the zero price occurs $G = 1,000,000$ times and the (integer) prices in $\{1, \ldots, V\}$ never occur. We then compute the marginal PDFs according to Equation (A.1). We denote the strategy by $PP(F^{Zero})$. Note that as soon as the ask price of a good exceeds zero, the updated belief for the good becomes uniform (see Chapter 1 for the update rule).

A.1.2 Infinite Point Prediction

By infinite prediction we mean a price prediction that is higher than the maximum price any agent is ever willing to pay given the agents' preference distribution. We implement it as an m-element price vector with each price equal to 1,000. We denote the strategy by $PP(\boldsymbol{\pi}^{\infty})$. $PP(\boldsymbol{\pi}^{\infty})$ serves as a useful performance benchmark for point price predictors. The agent bids if and only if it has single-unit demand, in which case its bidding is identical to that of SB.[2]

A.1.3 Uniform Distribution Prediction

For each good, we create an artificial empirical sample in which all prices in $\{0, \ldots, V\}$ occur 20,000 times. We then compute the marginal PDFs according to Equation (A.1). We denote the uniform-distribution predictor by $PP(F^U)$.

[1] In the results reported in Chapter 1, all agents have predictions derived for the preference distribution of their environment. We therefore suppress the subscripts of predictions to simplify the notation. Thus, if we consider a uniform environment, x refers to initial predictions based on samples from uniformly distributed types. If we consider an exponential environment, x refers to initial predictions based on samples from exponentially distributed types. One exception is the 53-strategy game that we constructed for the 5×5 uniform environment. In this restricted game, some agents have predictions based on the exponential distribution. The 53 strategies are described in Section A.4.1.

[2] Remember that any perceived-price predictor reverts to SB if the agent has single-unit demand (see Chapter 1).

(Complementary) Environment	SB-based distribution mean (=SB-based point prediction)						
$E(3,3)$	15.834	6.743	2.168				
$E(3,5)$	21.573	11.497	4.578				
$E(3,8)$	26.957	15.575	7.751				
$E(5,3)$	13.841	5.874	2.437	1.034	0.375		
$E(5,5)$	19.807	10.267	5.133	2.357	0.861		
$E(5,7)$	23.920	13.139	7.377	3.702	1.416		
$E(5,8)$	25.576	14.218	8.282	4.331	1.699		
$E(7,3)$	13.140	5.472	2.353	1.142	0.546	0.251	0.094
$E(7,6)$	21.461	11.247	6.051	3.340	1.803	0.898	0.340
$E(7,8)$	25.067	13.645	7.924	4.673	2.671	1.401	0.549
$E(7,9)$	26.567	14.595	8.687	5.263	3.089	1.655	0.661
$U(3,3)$	14.740	8.353	2.855				
$U(3,5)$	19.564	12.733	5.548				
$U(3,8)$	23.941	16.026	8.653				
$U(5,3)$	11.149	7.595	4.777	2.613	1.007		
$U(5,5)$	14.842	10.702	7.550	4.631	1.898		
$U(5,7)$	17.410	12.450	9.156	6.121	2.742		
$U(5,8)$	18.522	13.105	9.714	6.696	3.132		
$U(7,3)$	9.181	6.708	4.800	3.369	2.191	1.277	0.558
$U(7,6)$	13.454	10.019	7.735	5.978	4.319	2.697	0.200
$U(7,8)$	15.511	11.288	8.725	6.904	5.232	3.458	1.602
$U(7,9)$	16.438	11.825	9.103	7.244	5.581	3.781	1.792
$Ex1(2,2)$	15.000	14.501					

Table A.3: SB-based price predictions in complementary environments. $Ex1(2,2)$ refers to the fixed preferences described in Example 1.3.1.

A.1.4 SB-Based (Baseline) Prediction

We simulate one million ($G = 1,000,000$) game instances in which all players follow the SB strategy (described in Section 1.3.1). The SB-based *point* prediction is an m-element vector of the average final prices. The SB-based *distribution* prediction an m-vector of marginal price distributions, each computed according to Equation (A.1). We denote the SB-based point prediction by $PP(\boldsymbol{\pi}^{SB_x})$, where x labels the preference distribution of the SB-players: uniform (U) or exponential (E). Similarly, we denote the SB-based distribution predictor by $PP(F^{SB_x})$. We also refer to the SB-based prediction as baseline. In Table A.3, we report the means of SB-based distribution predictions for a number of complementary environments. The means also represent SB-based point predictions (by definition).

A.1.5 Competitive Equilibrium Prediction

We repeatedly sample agents' preferences from the preference distribution (exponential or uniform) and apply Walrasian tatonnement to obtain a crude Monte Carlo estimate of the expected price equilibrium. The process is described in Section 1.5.2. To construct a competitive-equilibrium distribution prediction, we apply Equation (A.1) to the sample of price-equilibrium estimates. We denote the strategy by $PP(F^{EPE_x})$, where x labels the preference distribution. We also construct two types of point predictions, $PP(\boldsymbol{\pi}^{EPE_x})$ and $PP(\boldsymbol{\pi}^{EDPE_x})$. They differ in the order in which averaging and the tatonnement are applied. See Section 1.5.2 for details.

We have implemented this prediction method only for the $U(5,5)$ and $E(5,5)$ complementary environments. We found the prices to which tatonnement converges to be sensitive to the choice of initial prices and other parameters of the tatonnement algorithm. In Section A.4.1 of this appendix, we provide all the competitive-equilibrium predictions we have obtained. All these predictions are based on 25,000 draws from the corresponding preference distribution.

A.1.6 Self-Confirming Prediction

In this section, we present SC predictions that we derived for a number of complementary environments. In substitutable environments, we analyzed a modification of SC prediction that we believe has a higher potential in such environments (see Section A.3).

In Tables A.5 and A.6 below, we report approximate SC point predictions and the means of approximate SC distribution predictions. In environments in which the iterative process did not reach a fixed point, the reported price vector is the mean of the last-iteration final-price distribution. In all cases but a few exceptions, the maximum number of iterations is 100 and the number of game instances per iteration is 1 million ($G = 1,000,000$). The exceptions are given in Table A.4. Environments $Ex1(2,2)$ and $Ex2(m,n)$ refer to the specific preferences described in Examples 1.3.1 and 1.5.1 respectively.

A.1.7 Gaussian Distribution Prediction

Let $\boldsymbol{\mu}$ be a vector of expected final prices and $\boldsymbol{\sigma}$ be a vector of standard errors. We can approximate the final-price (marginal) distribution of good $i \in \{1,\ldots,m\}$ with a Gaussian-shaped distribution defined on $\{0,\ldots,V\}$ and centered around μ_i. We draw $G = 1,000,000$ samples from $N(\mu_i, \sigma_i)$ for each $i \in \{1,\ldots,m\}$, round the

Environment	Prediction	Iterations	Games per iteration
$U(5,5)$a	Point	70	500,000
$U(5,5)$	Point	40	500,000
$U(5,5)$	Distribution	50	1 million
$E(7,9)$	Point	87	1 million
$Ex1(2,2)$	Point	100	10,000
$Ex1(2,2)$	Distribution	100	10,000

Table A.4: Deriving approximate self-confirming (SC) price predictions in complementary environments: Environments in which the maximum number of iterations or the number of games per iteration deviates from the standard setting (100 and 1 million respectively). Point predictions in rows $U(5,5)$a and $U(5,5)$ differ in the initial prediction (see Table A.5 for details). In environment $E(7,9)$, the maximum number of iterations was originally set to 100, but the simulation was aborted by the system after iteration 87.

prices and discard those outside the $[0,\ V]$ interval.[3] Then we compute the final-price probabilities according to Equation (A.1). We denote the Gaussian prediction by $G(x,y)$, where x and y label the expected-price and the standard-error vectors respectively. For example, the Gaussian prediction of $PP(G(\mu(EPE_u),\sigma(EPE_u)))$ is based on the means and standard errors of the F^{EPE_u} marginals; the prediction of $PP(G(\boldsymbol{\pi}^{EDPE_u},\sigma(EPE_u)))$ is based on vector $\boldsymbol{\pi}^{EDPE_u}$ the standard errors of F^{EPE_u}.

A.1.8 Degenerate Distribution Predictions

Let $\boldsymbol{\pi}$ be an m-element price vector. If the prices are all integers, we create for each good $i \in \{1,\dots,m\}$ an artificial empirical sample in which these prices each occur $G = 1,000,000$ times and the rest of the integer prices in $\{0,\dots,V\}$ never occur. We then compute the marginal PDFs according to Equation A.1. If the price of good i in a particular price vector $\boldsymbol{\pi}$ that we want to use is non-integer, we either split a one-million sample between the closest integer prices so that the distribution mean equals the original price $\boldsymbol{\pi}^i$ or simply round the price before generating the distribution. In the former two cases, we denote the degenerate-distribution predictor by $PP(F(\boldsymbol{\pi}))$. For example, the mean prediction of $PP(F(\boldsymbol{\pi}^{SB_u}))$ equals exactly $\boldsymbol{\pi}^{SB_u}$. In the latter case, we mark the price vector with a prime. For example, $PP(F(\boldsymbol{\pi}^{EPE'_u}))$ is based on the rounded $\boldsymbol{\pi}^{EPE_u}$.

[3] One obvious shortcoming of this approach is that the mean of the resulting distribution is different from μ_i, unless $\mu_i = \frac{V}{2}$.

(Complementary) Environment	Final-price distribution mean in the last iteration (= approx. SC point prediction if a fixed point reached)						
$E(3,3)$	14.321	5.436	1.651				
$E(3,5)$	20.097	8.624	2.969				
$E(3,8)$	25.344	11.519	4.436				
$E(5,3)$	12.329	4.789	2.034	0.920	0.355		
$E(5,5)$	18.071	7.883	3.915	1.872	0.670		
$E(5,7)$	22.083	9.943	5.354	2.745	0.980		
$E(5,8)$	23.795	10.840	5.923	3.117	1.107		
$E(7,3)$	11.626	4.439	1.970	1.005	0.571	0.298	0.104
$E(7,6)$	19.742	8.646	4.591	2.670	1.521	0.807	0.323
$E(7,8)$	23.354	10.480	5.910	3.606	2.167	1.181	0.466
$E(7,9)$	24.818	11.223	6.431	4.047	2.487	1.376	0.544
$U(3,3)$	12.988	6.738	2.140				
$U(3,5)$	17.742	9.785	3.421				
$U(3,8)$	21.786	11.729	4.367				
$U(5,3)$	9.533	6.235	3.758	2.012	0.778		
$U(5,5)$a	12.930	8.576	5.334	2.988	1.177		
$U(5,5)$	13.038	8.668	5.424	3.035	1.183		
$U(5,7)$	14.995	9.560	5.922	3.344	1.324		
$U(5,8)$	15.477	9.437	5.635	3.133	1.220		
$U(7,3)$	7.765	5.535	3.851	2.641	1.694	0.989	0.444
$U(7,6)$	11.795	8.272	5.800	4.025	2.662	1.567	0.697
$U(7,8)$	12.352	7.917	5.052	3.297	2.108	1.232	0.550
$U(7,9)$	11.876	6.932	3.928	2.339	1.430	0.847	0.384
$Ex1(2,2)$	14.749	14.251					
$Ex2(2,2)$	13.919	13.241					
$Ex2(2,5)$	15.708	15.046					
$Ex2(5,2)$	6.699	6.452	6.080	5.792	5.596		
$Ex2(5,5)$	8.060	7.599	7.433	7.303	7.213		

Table A.5: Deriving approximate self-confirming (SC) point price predictions in complementary environments: Distribution means in the last iteration. Price oscillation persisted only in $Ex1(2,2)$. In the rest of the environments, the price vectors satisfy the definition of the approximate SC point price prediction. We used the zero initial prediction $PP(\boldsymbol{\pi}^{Zero})$ (see Section A.1.1) to derive all the predictions in this table except that in row $U(5,5)$. In row $U(5,5)$, the initial prediction is the average of the last 50 (of 70) iterations used to generate the prediction in row $U(5,5)$a.

(Complementary) Envir.	KS_{marg} distance in last iter.	Last iter.	Final-price distribution mean in the last iter. (= mean of approximate SC distribution prediction if a fixed point is reached)						
$E(3,3)$	-0.00147	100	12.21	3.80	1.20				
$E(3,5)$	-0.00184	100	17.86	6.55	2.17				
$E(3,8)$	0.000899	100	23.40	9.34	3.40				
$E(5,3)$	-0.00126	100	10.52	3.45	1.52	0.76	0.33		
$E(5,5)$	0.00179	100	16.14	6.08	2.95	1.47	0.58		
$E(5,7)$	-0.00158	100	20.33	8.09	4.19	2.15	0.82		
$E(5,8)$	-0.0011	100	22.12	8.96	4.75	2.48	0.95		
$E(7,3)$	0.0014	100	9.92	3.20	1.48	0.82	0.50	0.296	0.132
$E(7,6)$	0.001252	100	17.92	6.84	3.59	2.09	1.24	0.717	0.307
$E(7,8)$	-0.00235	100	21.70	8.67	4.75	2.86	1.75	1.017	0.427
$E(7,9)$	0.000955	100	23.30	9.48	5.28	3.22	1.99	1.154	0.476
$U(3,3)$	0.001499	100	10.40	4.58	1.52				
$U(3,5)$	0.001845	100	15.24	7.29	2.52				
$U(3,8)$	0.005117	100	19.76	9.71	3.67				
$U(5,3)$	0.001061	100	7.32	4.42	2.69	1.54	0.69		
$U(5,5)$	0.007133	6	10.77	6.55	4.08	2.34	1.03		
$U(5,7)$	-0.00242	100	13.21	7.95	5.07	2.96	1.29		
$U(5,8)$	0.003864	100	14.22	8.39	5.40	3.19	1.40		
$U(7,3)$	-0.00234	100	5.96	4.02	2.83	2.00	1.37	0.871	0.433
$U(7,6)$	0.002896	100	9.74	6.46	4.59	3.26	2.23	1.400	0.690
$U(7,8)$	-0.00213	100	11.54	7.44	5.28	3.78	2.60	1.643	0.811
$U(7,9)$	-0.00211	100	12.29	7.78	5.53	3.97	2.74	1.728	0.853
$Ex1(2,2)$	-0.5055	100	8.00	7.50					
$Ex2(2,2)$	-0.20994	100	10.83	10.19					
$Ex2(2,5)$	0.172697	100	12.21	11.59					
$Ex2(5,2)$	0	2	1	0	0	0	0		
$Ex2(5,5)$	0	2	1	0	0	0	0		

Table A.6: Deriving approximate self-confirming (SC) distribution price predictions in complementary environments: Distribution means in the last iteration. We used uniform initial predictions F^U to derive all the distribution predictions in this table except that in row $U(5,5)$, for which we used the SB-based baseline prediction F^{SB_u}. The KS_{marg} convergence criterion in $U(5,5)$ is 0.01, i.e., we programmed the simulation to stop if the KS_{marg} distance between the current- and previous-iteration final-price CDFs for all goods was below 0.01. The criterion was satisfied at iteration 6 in $U(5,5)$. For the rest of the environments in this table, the KS_{marg} criterion is 0.00001. This threshold was never reached in any other U- or E-environments, and the simulation stopped at the threshold on the number of iterations, which equals 100 in all the environments. However, KS_{marg} reached 0.01 within the first 11 iterations in all the U-environments and within the first 7 iterations in all the E-environments.

A.2 Demand-Reduction Strategy

The demand-reduction strategy family, $DR(\kappa)$, was introduced for homogeneous-good environments (see Section 1.7.1). In a homogeneous-good environment, each auction sells one unit of a homogeneous indivisible good, and the bidders' marginal value for one more unit of the good is weakly decreasing. We implemented such preferences by randomly drawing marginal values v_k for the kth unit from $U[0, v_{k-1}]$, where v_0 is a uniform upper bound on the marginal value of a unit, which equals 127 in our empirical-game analysis.

$DR(\kappa)$'s perceived price of the unit with the lth lowest *myopically perceived price* (see Section 1.3.1) is given by

$$\rho_l(\boldsymbol{B}) \equiv \begin{cases} \beta_l + \kappa(l-1) & \text{if winning the unit} \\ \beta_l + 1 + \kappa(l-1) & \text{otherwise,} \end{cases} \tag{A.2}$$

where $\boldsymbol{\beta}$ is the vector of current bid prices. Thus, the parameter $\kappa \in 0, \ldots, 127$ defines the degree of the agent's demand reduction. An agent with a larger κ bids on fewer units. When $\kappa = 0$, the agent's bidding behavior is equivalent to that of SB (see Section 1.3.1). In the other extreme case when $\kappa = 127$, the agent never bids at all. For more details on the strategy, see Section 1.7.1.

A.3 Own-Effect Price-Prediction Strategy

The own-effect price predictor, $OEPP(\boldsymbol{\pi}^x)$, is designed for homogeneous-good environments (see Section A.2; for analysis, see Section 1.7.2) and is a modification of the price predictor (see Section A.1). Its prediction is an $m \times m$ matrix of *predicted own-effect prices*. Each element of the matrix, which we denote $\pi_{iq}(\boldsymbol{B})$, is a predicted final price of unit i given that the agent tries to win q units and its information state at the current round is $\boldsymbol{B}$. In our analysis, the *initial* price prediction is equal across auctions: $\pi_{iq}(\varnothing) = \pi_{jq}(\varnothing)$ for all i and j, for all purchase sizes q. In other words, the initial-prediction matrix consists of m identical rows. We label it by $\boldsymbol{\pi^x}$, where the subscript x labels initial predictions. Of particular interest is what we call *self-confirming own-effect prices*, which we describe in the following section.

A.3.1 Self-Confirming Own-Effect Prices

Self-confirming own-effect prices satisfy the condition that if one of the agents (the "explorer") bids to win q units ignoring its preferences and the other agents "exploit" their own-effect price predictions, that prediction on average *is correct* for all q. See Section 1.7.3 for a formal definition and the derivation procedure. We

1	2	3	4	5
19.0538	39.994	60.934	80.4613	101.15
19.7819	39.6585	60.9936	80.3963	101.755
19.3541	40.2281	60.9409	80.4295	102.034
19.2679	39.8432	60.8178	80.2026	101.359
18.5337	39.5263	60.5431	80.4075	101.501
19.5411	40.2917	61.0213	80.5834	101.373
18.9921	40.0143	61.2818	80.377	101.792
18.2261	39.8009	61.2354	80.4773	101.684
18.73	39.5478	61.2418	80.572	101.787
18.4487	39.7332	61.4828	80.5034	101.204
18.99294	39.8638	61.04925	80.44103	101.5639

Table A.7: Approximate self-confirming own-effect predictions and their average. We used the average as $\boldsymbol{\pi}^{SC}$ in our empirical-game analysis in Section 1.7.4. The columns are the possible target purchase sizes ($q \in \{1, \ldots, m\}$, where $m = 5$ auctions). V=127 in our analysis of homogeneous-good environments.

denote the self-confirming own-effect price matrix by $\boldsymbol{\pi}^{SC}$ and the self-confirming own-effect price-prediction strategy by $OEPP(\boldsymbol{\pi}^{SC})$.

We have implemented $OEPP(\boldsymbol{\pi}^{SC})$ for a homogeneous-good environment with 5 units and 5 agents (see Section 1.7). We set the initial predictions to a matrix of zeros, and simulated 10,000 games for every purchase size of the explorer. The explorer's purchase size was changed 500 times, i.e., each purchase size was updated 100 times. We created 10 approximate self-confirming own-effect predictions, which we report in Table A.7. Since any own-effect initial-prediction matrix consists of m identical rows, we present one row per prediction. In our empirical-game analysis, we used the average of these 10 predictions as $\boldsymbol{\pi}^{SC}$.

A.4 Restricted SAA Games

A.4.1 53-Strategy Game for 5×5 Uniform Complementary Environment

In Table A.8, we describe 53 strategies that we constructed for the 5×5 uniform environment ($U(5,5)$) analyzed in Chapter 1. In Table A.9, we report all initial point predictions for the price-predicting strategies derived for this environment.

A.4.2 Strategies for Alternative Complementary Environments

For 11 alternative complementary environments ($E(3,3)$, $E(3,8)$, $E(5,8)$, $U(3,3)$, $U(3,5)$, $U(3,8)$, $U(5,3)$, $U(5,8)$, $U(7,3)$, $U(7,6)$, $U(7,8)$), we have evaluated

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	20	SA(k) $k = 0, 0.05, 0.1, \ldots 0.95$
Point Price Predictor	13	$PP(\boldsymbol{\pi}^{\infty})$, $PP(\boldsymbol{\pi}^{SB_u})$, $PP(\boldsymbol{\pi}^{SB_u})$ w/ P.O., $PP(\boldsymbol{\pi}^{EPE_u})$, $PP(\boldsymbol{\pi}^{EPE_u^*})$, $PP(\boldsymbol{\pi}^{EPE_u^{**}})$, $PP(\boldsymbol{\pi}^{EPE_e})$, $PP(\boldsymbol{\pi}^{EPE_e^*})$, $PP(\boldsymbol{\pi}^{EDPE_u})$, $PP(\boldsymbol{\pi}^{EDPE_u^*})$, $PP(\boldsymbol{\pi}^{EDPE_e})$, $PP(\boldsymbol{\pi}^{EDPE_e^*})$, $PP(\boldsymbol{\pi}^{SC_u})$
Price Distribution Predictor	19	$PP(F^{Zero})$,$PP(F^{U})$ $PP(F^{SB_u})$,$PP(F^{SC_u})$, $PP(F^{EPE_u})$, $PP(G(\mu(EPE_u), \sigma(EPE_u)))$, $PP(G(\mu(SB_u), \sigma(SB_u)))$, $PP(G(\boldsymbol{\pi}^{EDPE_u}, \sigma(EPE_u)))$, $PP(G(\boldsymbol{\pi}^{EDPE_u}, \sigma(SB_u)))$, $PP(G(\boldsymbol{\pi}^{SC_u}, \sigma(EPE_u)))$, $PP(G(\boldsymbol{\pi}^{SC_u}, \sigma(SB_u)))$, $PP(F(\boldsymbol{\pi}^{EPE_u}))$, $PP(F(\boldsymbol{\pi}^{EPE'_u}))$, $PP(F(\boldsymbol{\pi}^{EDPE_u}))$, $PP(F(\boldsymbol{\pi}^{EDPE'_u}))$, $PP(F(\boldsymbol{\pi}^{SB_u}))$, $PP(F(\boldsymbol{\pi}^{SB'_u}))$ $PP(F(\boldsymbol{\pi}^{SC_u}))$, $PP(F(\boldsymbol{\pi}^{SC'_u}))$

Table A.8: 53 strategies for the complementary 5×5 uniform-environment game (complementary $U(5,5)$). We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family. P.O. refers to *participation-only* prediction (see MacKie-Mason et al. (2004)). Different point predictions obtained using the same (non-deterministic) algorithm are marked by asterisks.

Predictions/Good	1	2	3	4	5	Appendix Section
	Initial Point Predictions					
$\boldsymbol{\pi}^{\infty}$	1000	1000	1000	1000	1000	A.1.2
$\boldsymbol{\pi}^{SB_u}$	14.8	10.7	7.6	4.6	1.9	A.1.4
$\boldsymbol{\pi}^{SC_u}$	13.0	8.7	5.4	3.0	1.2	A.1.6
$\boldsymbol{\pi}^{EPE_u}$	16.6	10.8	6.5	3.1	0.7	A.1.5
$\boldsymbol{\pi}^{EPE_u^*}$	16.5	10.7	6.4	3.1	0.8	A.1.5
$\boldsymbol{\pi}^{EPE_u^{**}}$	26.0	14.2	6.9	2.5	0.3	A.1.5
$\boldsymbol{\pi}^{EPE_e}$	6.0	4.1	1.8	0.6	0.1	A.1.5
$\boldsymbol{\pi}^{EPE_e^*}$	30.5	11.9	6.0	2.7	0.4	A.1.5
$\boldsymbol{\pi}^{EDPE_u}$	20.0	12.0	8.0	2.0	0.0	A.1.5
$\boldsymbol{\pi}^{EDPE_u^*}$	20.8	11.4	8.2	1.8	0.0	A.1.5
$\boldsymbol{\pi}^{EDPE_e}$	25.0	10.0	5.1	0.9	0.0	A.1.5
$\boldsymbol{\pi}^{EDPE_e^*}$	24.5	10.5	5.5	1.5	0.0	A.1.5

Table A.9: Initial predictions (rounded to one decimal place) for the 5×5 uniform complementary environment (complementary $U(5,5)$). Column 1 is the notation, and column 2 is the point-prediction vectors. The goods are numbered from 1 through 5. The monotonicity of the prices is due to the specifics of the scheduling-game preferences (see MacKie-Mason et al. (2004) and Reeves et al. (2005)). For information about price-predicting strategies, see Section A.1. The subsections most relevant to particular predictions are given in column 3.

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	20	SA(k) $k = \{0, 0.05, 0.1, \ldots 0.95\}$
Point Price Predictors	3	$PP(\boldsymbol{\pi}^{\infty})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{SC})$
Price Distribution Predictors	6	$PP(F^{U})$, $PP(F^{SB})$

Table A.10: 26 deviators for 11 alternative environments. We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family.

27 profiles: one with all $PP(F^{SC})$, and for each of 26 other strategies s, one profile with $(n-1)$ $PP(F^{SC})$ and one s, where s is a strategy from Table A.10. Also, for 16 alternative complementary environments reported in Table A.2 (all environments other than complementary $U(5,5)$, which include the 11 above-mentioned environments and 5 additional E models), we evaluated complete 7-cliques. In Table A.11, we describe the 7-clique restricted games for each of the alternative environments. We suppress the preference-distribution labels in the tables, because they all match the corresponding environment (e.g., $PP(F^{SB})$ for $E(5,3)$ refers to $PP(F^{SB_e})$, and $PP(F^{SB})$ for $U(5,3)$ refers to $PP(F^{SB_u})$).

A.4.3 51-Strategy Game for 5×5 Uniform Homogeneous Environment

To analyze the 5×5 homogeneous-good environment in Section 1.7.4, we constructed a 51-strategy restricted game described in Table A.12.

Environment	Strategies
$E(3,3)$	$PP(F^{SC})$, $PP(F^{SB})$, SA(0.6), SA(0.7), SA(0.75), SA(0.8), SA(0.85)
$E(3,5)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.85)
$E(3,8)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.85)
$E(5,3)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SB
$E(5,5)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.35)
$E(5,8)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.35)
$E(7,3)$	$PP(F^{SC})$, $PP(F^{SB})$, SA(0.55), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$E(7,6)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SB
$U(3,3)$	$PP(F^{SC})$, $PP(F^{SB})$, SA(0.65), SA(0.7), SA(0.75), SA(0.8), SA(0.85)
$U(3,5)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, SA(0.7), SA(0.75), SA(0.85)
$U(3,8)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.9)
$U(5,3)$	$PP(F^{SC})$, $PP(F^{SB})$, SA(0.6), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$U(5,8)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SA(0.9)
$U(7,3)$	$PP(F^{SC})$, $PP(F^{SB})$, SA(0.55), SA(0.65), SA(0.7), SA(0.75), SA(0.8)
$U(7,6)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, SA(0.75), SA(0.8), SA(0.9)
$U(7,8)$	$PP(F^{SC})$, $PP(F^{SB})$, $PP(F^{U})$, $PP(\boldsymbol{\pi}^{SC})$, $PP(\boldsymbol{\pi}^{SB})$, $PP(\boldsymbol{\pi}^{\infty})$, SB

Table A.11: 7-clique restricted games for 16 alternative environments. $E(m,n)$ and $U(m,n)$ refer to the environments with m goods, n agents, and exponential and uniform preference distribution respectively (see Section A.1 for information about the environments).

Strategy Family	Number of Representatives in Game	Strategies
Straightforward Bidder	1	SB
Sunk-aware Agent	1	SA(k) $k = 0.5$
Price Distribution Predictor	1	$PP(F^{SB_u})$
Demand-Reduction Agent	47	$DR(\kappa)$ $\kappa = \{1, 2, \ldots 30, 32, 34,$ $36, 38, 40, 44, 48, 50, 52, 56,$ $60, 70, 80, 90, 100, 110, 120\}$
Own-Effect Price Predictor	1	$OEPP(\boldsymbol{\pi}^{SC})$

Table A.12: 51-strategy restricted game for the 5×5 homogeneous-good uniform environment (also referred to as substitutable $U(5, 5)$). We report the strategies in column 3 and the strategy family to which they belong in column 1. Column 2 is the total number of strategies from each strategy family.

APPENDIX B

Empirical Distribution of Number of Articles Read in a Journal (King and Griffiths, 1995)

Figure B.1 displays a histogram of the empirical distribution of the number of scholarly articles read by a sample of readers (King & Griffiths, 1995) from a collection of about 100 articles. The horizontal axis is the number of articles. The vertical axis is the percentage of the readers who read the corresponding number of articles. The shape of the histogram approximates the density of an exponential distribution.

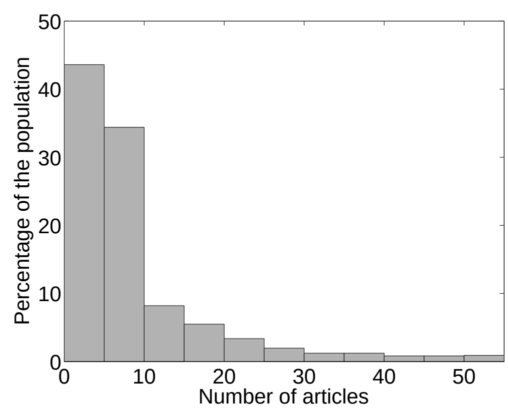

Figure B.1: Distribution of Number of Articles Read in a Journal (King and Griffiths, 1995)

k	Number of articles read	$P(K < k)$
0.05	5	.436
0.1	10	0.78
0.15	15	0.8621
0.2	20	0.9171
0.25	25	0.9508
0.3	30	0.9795
0.4	40	0.9828
0.5	50	0.991
≥ 0.5	≥ 50	1

Table B.1: Data points for Equation (C.2) in Chapter 2. The last column describes the empirical cumulative distribution of the number of scholarly articles read by a sample of readers (King & Griffiths, 1995) from a collection of about 100 articles. *Source: Chuang and Sirbu (1998), p.155.*

APPENDIX C

Parameters of Exponential Preference Distribution Estimated Based on King and Griffith's Empirical Data (1995)

In this appendix, I demonstrate how the parameters for the distribution of k were estimated. The cumulative distribution function of the exponential distribution is given by

$$F(k) = 1 - e^{-\lambda k}, \tag{C.1}$$

where λ is the distribution parameter.

Given γ and assuming $w = 100$ and $N = 100$, we can estimate the mean of the distribution, $\theta = \frac{1}{\lambda}$. Substituting $\theta = \frac{1}{\lambda}$ in (C.1), rearranging the terms and taking logarithm of the both sides of the equation, we obtain

$$k = -\theta \text{Log}(1 - F(k)). \tag{C.2}$$

Let k_i be the preference breadth such that a consumer with access to a collection of a hundred articles would choose to read i of them. Assume also that intensity is constant and equals one hundred ($w = 100$). Then King and Griffiths' data provide eight data points for Equation (C.2) (see Appendix B). The estimated values of k depend on the assumed value of the substitution effect γ. In Table C.1 I report a subset of the estimated values of the mean of the distribution θ, as well as the parameter λ, for a range of values of γ.

γ	$\theta = \frac{1}{\lambda}$	λ
0.	0.0927717	10.7791
0.01	0.0934613	10.6996
0.02	0.0941623	10.62
0.03	0.0948751	10.5402
0.04	0.0956	10.4602
0.05	0.0963374	10.3802
0.1	0.100223	9.97771
0.2	0.109162	9.16066
0.3	0.120132	8.32414
0.4	0.134001	7.46265
0.5	0.152255	6.56794
0.55	0.163819	6.10428
0.6	0.177723	5.62675
0.65	0.194853	5.13208
0.7	0.216652	4.61569
0.75	0.245647	4.07088
0.8	0.286764	3.48719
0.85	0.351227	2.84716
0.9	0.472042	2.11845
0.91	0.510953	1.95713
0.92	0.558992	1.78894
0.93	0.620017	1.61286
0.94	0.700454	1.42765
0.95	0.811852	1.23175

Table C.1: Mean ($\theta = \frac{1}{\lambda}$) and parameter λ of the exponential preference distribution estimated using Equation (C.2) in Chapter 2. The parameter is estimated for a range of γ (column 1) based on empirical data due to King and Griffiths (1995). The collection size is $N = 100$ items. The intensity is assumed to be fixed at $w = 100$. *Data source: Chuang and Sirbu (1998), p. 155.*

APPENDIX D

Learning Bayesian Nash Equilibrium: Best-Reply Functions

The figures in this appendix display the players' best-reply functions. Since the players' payoff functions are quadratic, the best replies are straight lines. Note that the best-reply functions do not depend on the other player's type: in each environment, same-color lines look the same. Also, the two submodular environments, SubNP and SubP, have the same best-reply graphs (Figure D.1). This is because the column player's preferences in the two environments are exactly the same, and the non-constant coefficients of the row player's payoff functions in SubP equal those in SubNP multiplied by the same constant (2). See Table 3.3, functions π_{c_1} for the column players' preferences and functions π_{r_2} and π_{r_1} for the row players' preferences.

Each panel displays the best-reply functions of two players—the row and the column—given their color types. For example, the top left panel on each figure displays the best replies of the red row and the green column players. The row player's best reply (solid lines) is a function of the column player's choices (the horizontal axis). The column player's best reply (dashed lines) is a function of the row player's choices (the vertical axis). Consider, for example, panel (a) on Figure D.1 (submodular environments, {red row, green column}). The best reply of the red row player (red solid line) to column's choice equal to 10 is 17.5. To find the best reply of the green column player (green dashed line) to 10, find 10 on the vertical axis and read the value of the best reply on the horizontal axis: it is around 15. Note that in the submodular environments, the slopes of the best-reply functions are negative. This property, called strategic substitutability, is implied by the condition of submodularity. On Figure D.2 displaying the supermodular (and non-potential) environment, the best-reply slopes are positive. This property is called strategic complementarity and is implied by supermodularity.

Unconstrained best replies to some feasible choices can be outside the player's strategy space. Consider again panel (a) on Figure D.1. The boundaries of the strategy space are outlined by the gray rectangle. If the green column player chooses 0, the red row player's unconstrained best reply is around 24, which is outside the rectangle area. In the experiment, however, the subjects can observe payoffs only for feasible choice profiles, which means that 21 would be the row player's best reply in this case. In fact, 21 is the row player's best reply to a range of column player's choices from 0 to around 4, which means that the red row player's best reply in the submodular environments is a kinked rather than a straight line. I summarize all cases of best-reply functions with a kink in Table D.1. In all cases, the binding strategy-space limit is 21.

Environment	Player type	Kink location
SubNP & SubP	red row	column player's choice is 4.33
SubNP & SubP	green column	row player's choice is 0.67
SuperNP	red row	column player's choice is 17.67

Table D.1: Best replies.

A game obtained by assigning each player of the original Bayesian game a particular type is a complete-information simultaneous-move game. The (Nash) equilibria of such games are in general different from the equilibrium of the original Bayesian game, but they may convey some useful information about the Bayes-Nash equilibrium. Therefore, I chose the strategy space and the players' payoff function parameters in such a way that these equilibria are feasible strategy profiles (graphically represented as intersections of the best-reply functions on Figures D.1–D.2). This is particularly important under Bayesian-information condition, where subjects are given a set of payoff tables displaying the players' payoffs for each combination of the row and column types. Since I have no knowledge about how subjects may use the strategic information they can read on such payoff tables, I impose the condition of feasibility of the Nash equilibria of the games displayed on each table in order to reduce the probability of an unknown behavioral confound.

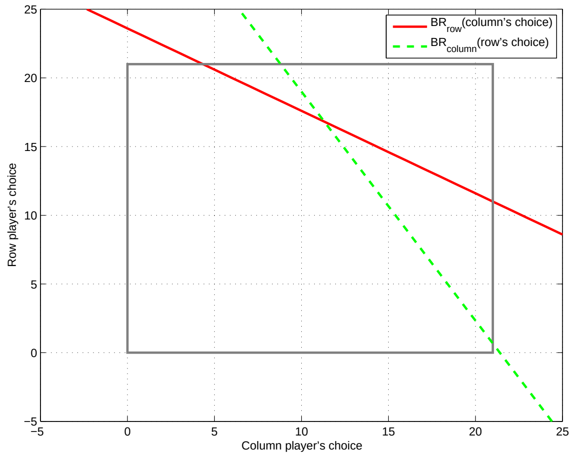

(a) {red row, green column}

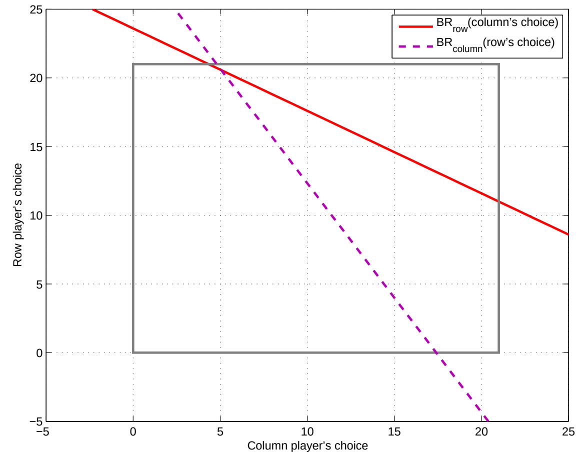

(b) {red row, purple column}

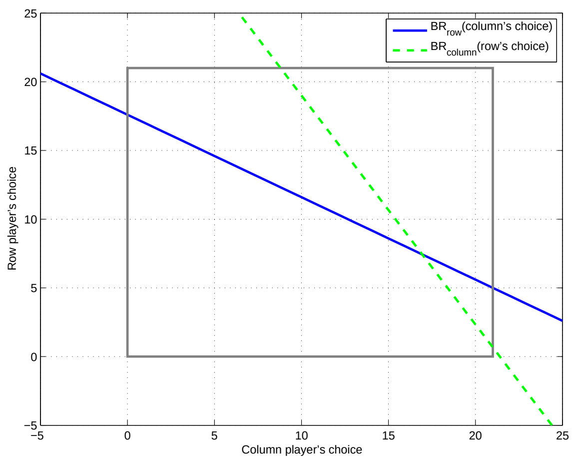

(c) {blue row, green column}

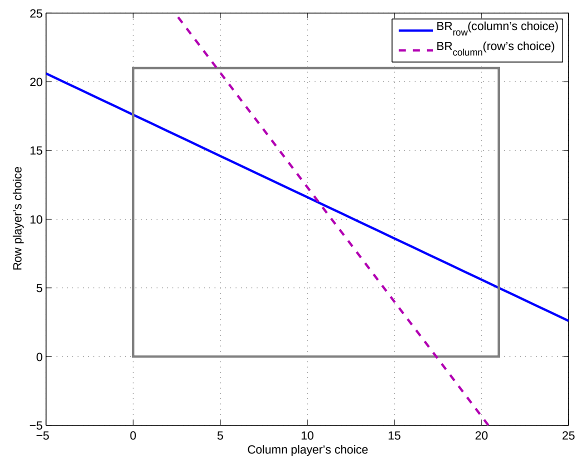

(d) {blue row, purple column}

Figure D.1: Players' best-reply functions: Submodular environments (SubNP and SubP). The panels are labeled by the players' type profile: {row's type, column's type}.

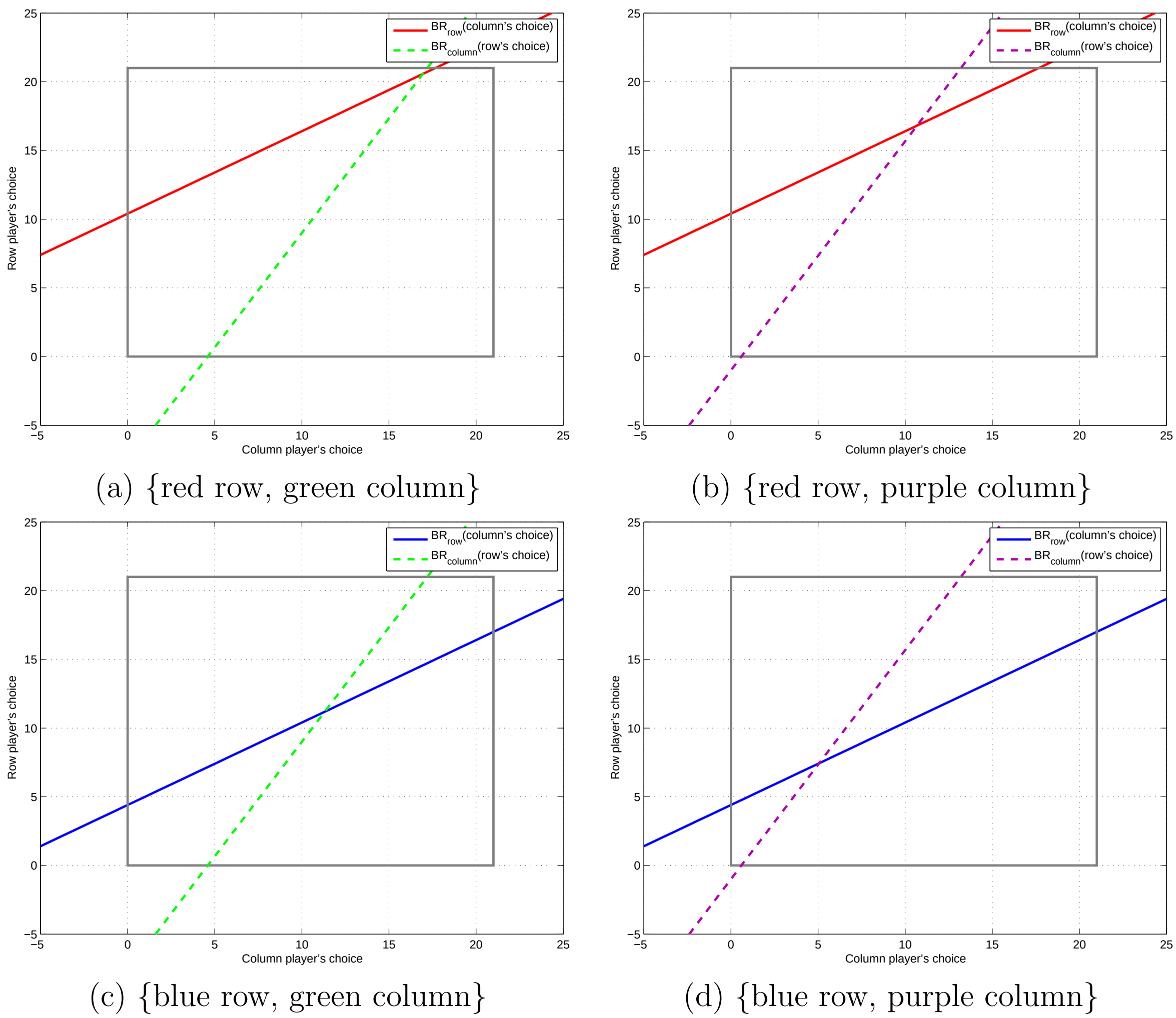

(a) {red row, green column}

(b) {red row, purple column}

(c) {blue row, green column}

(d) {blue row, purple column}

Figure D.2: Players' best-reply functions: Supermodular & non-potential environment (SuperNP). The panels are labeled by the players' type profile: {row's type, column's type}.

APPENDIX E

Learning Bayesian Nash Equilibrium: Payoff Gradient Fields

In Apeendix D, I use best-reply functions to summarize the strategic properties of the experimental environments. The plots I present in this appendix are another visual representation of the players' preferences. They are gradient fields of the players' payoff functions, which highlight each player's incentive to deviate at any given strategy profile.

Figures E.1–E.2 are superimposed gradient plots of the players' payoff functions. Figure E.1 displays submodular environments (potential and non-potential), and figure E.2 displays the supermodular (and non-potential) environment. Each panel displays the payoff functions of two players—the row and the column—given their color types. For example, the top left panel on each figure shows the payoffs of the red row and the green column players.

The figures can be thought of as a continuous-strategy normal-form representation of a stage game: the row player's strategies are on the vertical axis, increasing from top (0) to bottom (21); the column player's strategies are on the horizontal axis, increasing from left (0) to right (21). The curved lines are indifference curves, which gradually grow thicker as the corresponding payoff function grows. The curves effectively convey the shape of the payoff functions. They are displayed on all payoff tables available to the subjects under Bayesian-information condition (see an example in Appendix H). The vertical arrows on the plots show the row player's payoff gradient along her strategy space. Similarly, the horizontal arrows show the column player's gradient. The longer the arrows, the steeper the function at the point. Thus, the dots are near-equilibrium area of the normal-form game induced by fixing the type profile (but generally not of the Bayesian game).

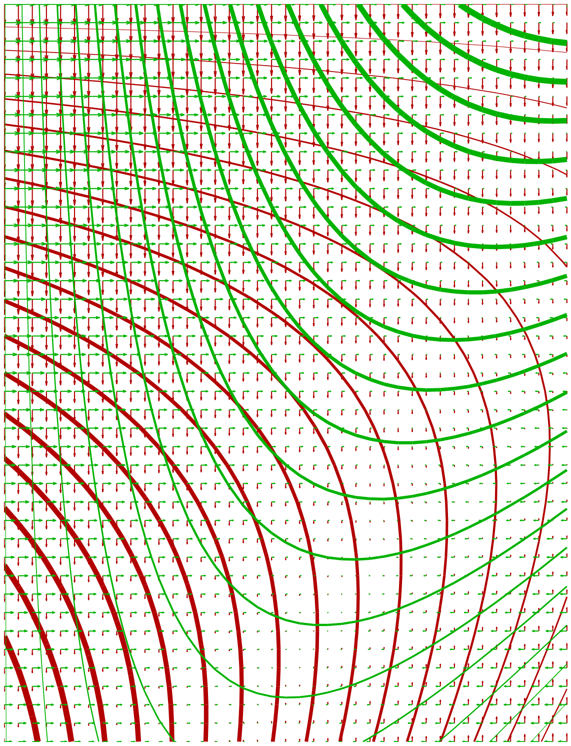

{red row, green column}

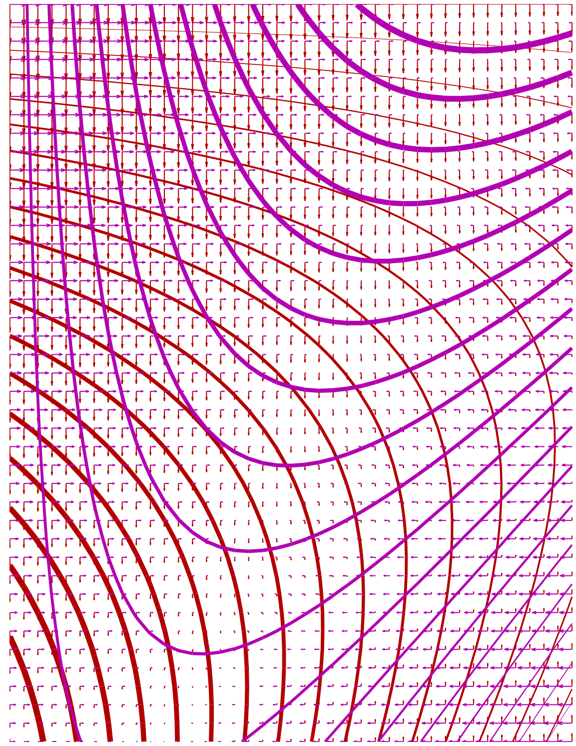

{red row, purple column}

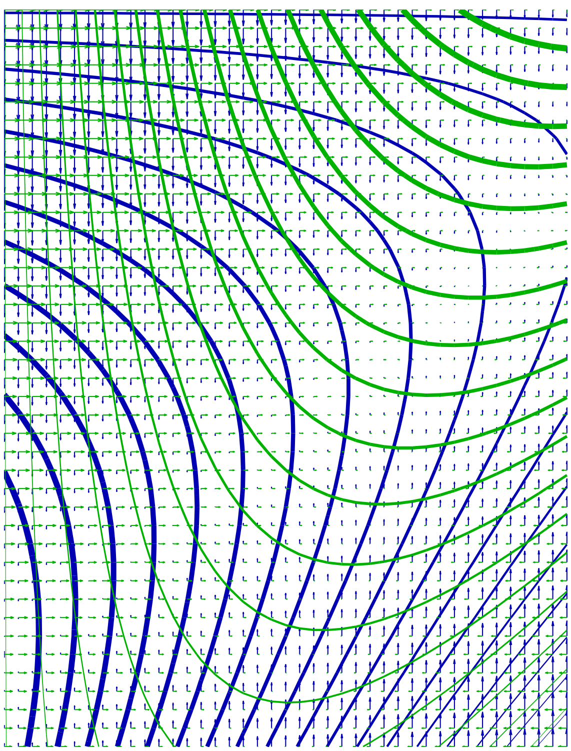

{blue row, green column}

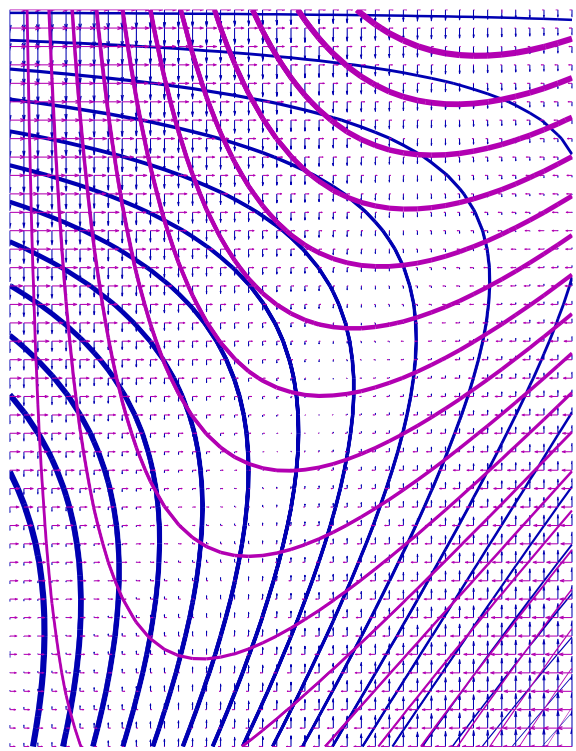

{blue row, purple column}

Figure E.1: Gradient plots of the players' payoff functions: Submodular environments (SubNP and SubP). The panels are labeled by the players' type profile: {row's type, column's type}.

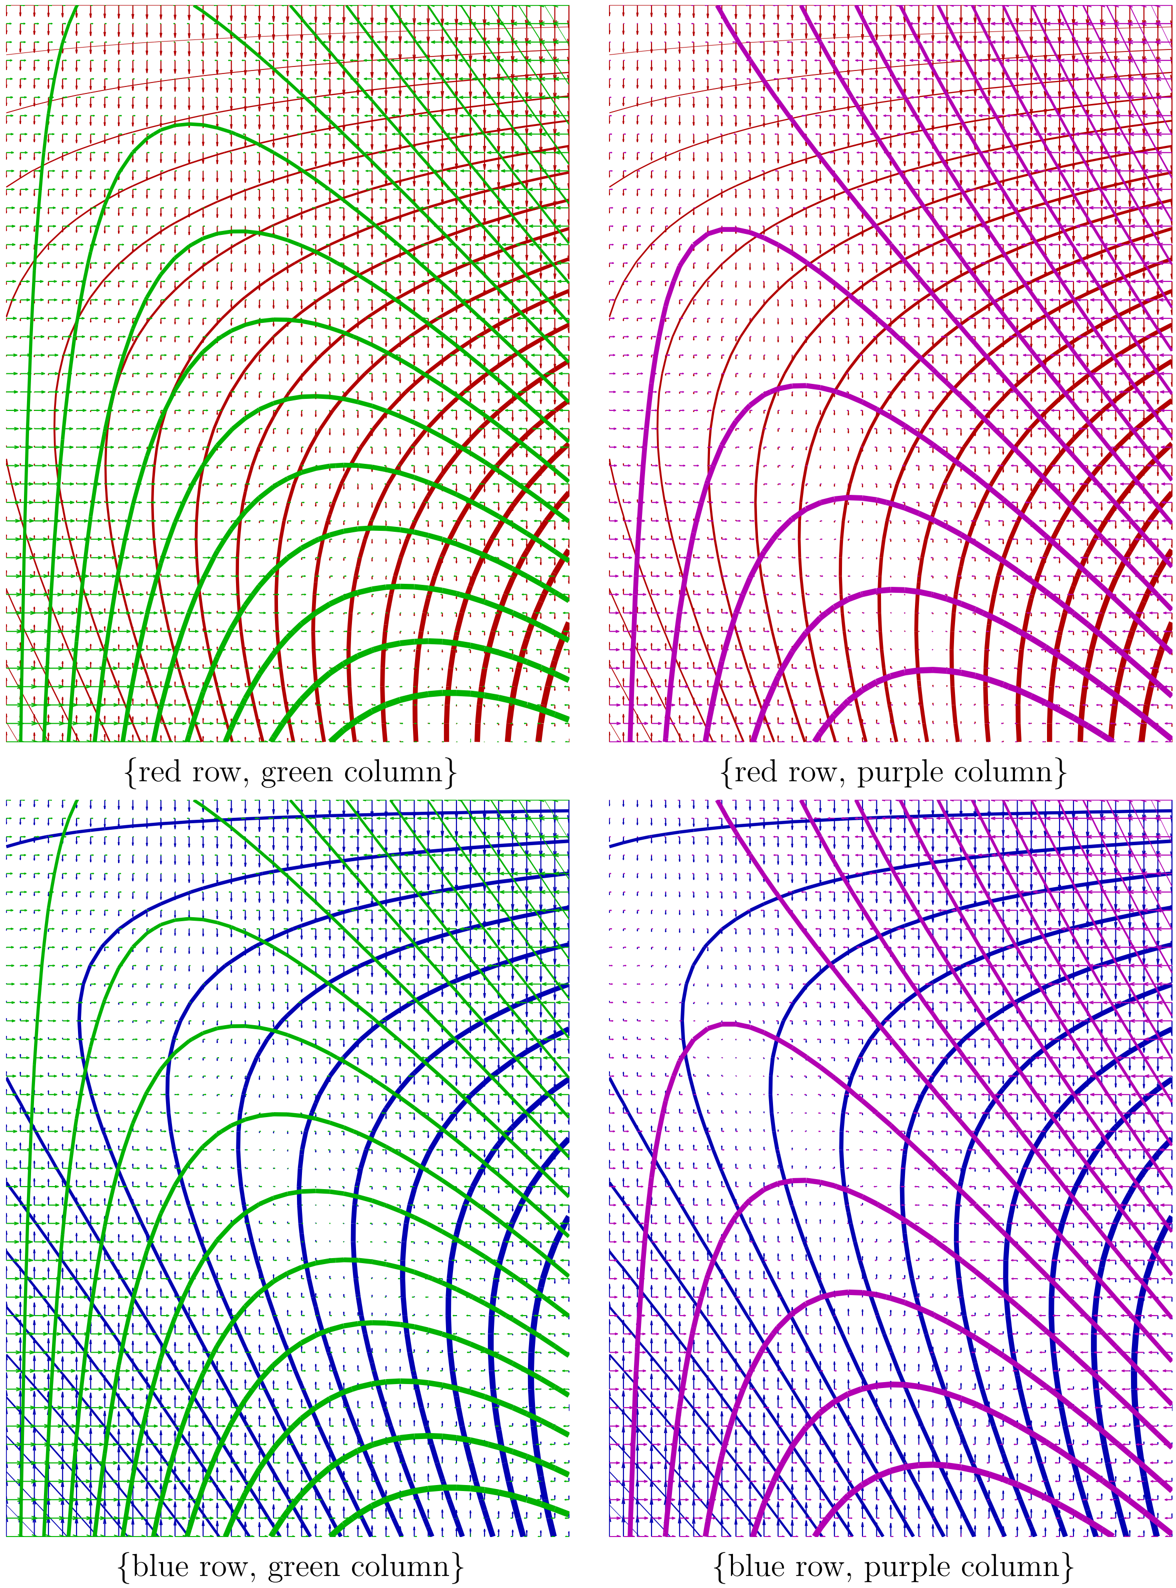

Figure E.2: Gradient plots of the players' payoff functions: Supermodular & non-potential environment (SuperNP). The panels are labeled by the players' type profile: {row's type, column's type}.

APPENDIX F

Learning Bayesian Nash Equilibrium: Computer Simulations

In this appendix, I report on computer simulations I performed to test the experimental design. The main idea is to explore the game dynamics when the players follow an established learning model capturing the principles of human learning behavior. In particular, I was interested in the following questions:

1. Would software agents behave as predicted by theory in the experiment environments? If not, what features of the design can explain that?

2. How many rounds would it take software agents to converge to the equilibrium (if they converge at all). This number served as an estimate of the required length of the experiment.

3. What average payoffs would software agents earn under different learning models?

4. What would the average payoffs be if the players all chose actions at random? How do they compare to the equilibrium payoffs? I used this and the average-payoff estimates to compute a prediction of the final average earning of human subjects and to adjust the constant coefficients of the payoff functions and determine the exchange rate between points and US dollars. I report the estimated average payoffs in Table F.1.

I programmed the software agents to implement the following learning models:

1. Replicator dynamic (RD).

2. Generalized fictitious play (GFP) with a discount factor $\gamma \in (0, 1)$. In this appendix I report the results for $\gamma = 0.5$ and $\gamma = 0.9$.

3. Fictitious play (FP), which is a special case of GFP with $\gamma = 1$.

4. Cournot best-response dynamic (BR), which is a special case of GFP with $\gamma = 0$.
5. Population fictitious play (PFP).
6. Exponentialized relative payoff sums (RPS).

Learning Model	Red Row	Blue Row	Green Column	Purple Column	Row Average	Column Average
SubNP						
Replicator dynamic	1337	788	1175	719	1063	947
PFP	1342	776	1215	711	1059	963
FP (GFP($\gamma = 1$))	1338	779	1207	716	1058	961
GFP($\gamma = 0.9$)	1339	786	1178	718	1062	948
GFP($\gamma = 0.5$)	1337	784	1169	707	1061	938
BR (GFP($\gamma = 0$))	1318	780	1134	680	1049	907
RPS ($\gamma = 0.5$, $\lambda = 0.004$)	938	563	1378	751	750	1065
Random choice	996	628	1223	756	812	989
SubP						
Replicator dynamic	1371	958	1170	722	1164	946
PFP	1385	931	1217	709	1158	963
FP (GFP($\gamma = 1$))	1373	938	1204	717	1156	960
GFP($\gamma = 0.9$)	1394	937	1201	698	1166	949
GFP($\gamma = 0.5$)	1375	941	1162	706	1158	934
BR (GFP($\gamma = 0$))	1350	926	1148	665	1138	906
RPS ($\gamma = 0.5$, $\lambda = 0.004$)	582	494	1330	725	538	1028
Random choice	681	632	1193	726	657	959
SuperNP						
Replicator dynamic	1478	818	1238	810	1148	1024
PFP	1557	844	1272	818	1200	1045
FP (GFP($\gamma = 1$))	1427	763	1221	786	1095	1003
GFP($\gamma = 0.9$)	1490	822	1244	810	1156	1027
GFP($\gamma = 0.5$)	1478	801	1237	798	1139	1017
BR (GFP($\gamma = 0$))	1458	783	1226	782	1120	1004
RPS ($\gamma = 0.5$, $\lambda = 0.004$)	774	432	890	457	603	674
Random choice	604	396	820	388	500	604

Table F.1: Average payoffs over 500 simulated games. The payoffs are rounded to the nearest integer.

F.1 Replicator dynamic

Ely and Sandholm (2005) proposed the following replicator dynamic to find Bayes-Nash equilibria:

$$\dot{x} = BR(x) - x, \tag{F.1}$$

where x is the distribution of pure strategies in the population of players.

In this experimental study, there are two different populations of pure strategies, because the row and the column players have different preferences. Thus, one population of replicator-dynamic players represents the row player's distribution of pure strategies (a real number in [0, 21]), and the other population represents the column players' distribution of pure strategies.

For the Cournot duopoly game used in the experiment, the equation above, generalized to two-population case, means that player $i = \{r, c\}$ of type type θ_i, where $\theta_r = \{r, b\}$ and $\theta_c = \{g, p\}$, plays the best reply to the average play of all players in the previous round. The dynamic becomes

$$BR_i^{\theta_i}(t+1) = \frac{b_i^{\theta_i} - f_i \frac{1}{n} \sum_{j=1}^{n} y_j(t)}{2d_i}, \tag{F.2}$$

where $y_j(t)$ is player j's play in period t, and n is the number of players ($n = 8$ in this experiment).

F.2 Fictitious-Play Learning Models

In the generalized fictitious play, player $i = \{r, c\}$, best-replies to the discounted sum of the history of the choices of other players she has been matched with:

$$BR_i^{\theta_i}(t+1) = \frac{b_i^{\theta_i} - f_i y_j(t+1)}{2d_i}, \tag{F.3}$$

where

$$y_j(t+1) = \frac{y_j(t) + \sum_{u=1}^{t-1} \gamma^u y_j(t-u)}{1 + \sum_{u=1}^{t-1} \gamma^u}, \tag{F.4}$$

where γ is the discount factor and $j = \{r, c\}$ is player i's match, who is randomly chosen at the beginning of each round according to the type distribution.

The original fictitious play corresponds to $\gamma = 1$. The Cournot best-response dynamic corresponds to $\gamma = 0$.

The above can be thought of as *individual learning*: the history of observations to which a player best-replies is the player's own memory. With *social learning*, on the other hand, the distribution of actions for the whole population is revealed after each round, as in replicator dynamic discussed above, for example.

In the population fictitious play, which is also a social type of learning, the players best-reply to the discounted sum of the average choice history of *all other* players in the population, rather than to the choice history of the players who have been matched with her so far:

$$BR_i^{\theta_i}(t+1) = \frac{b_i - f_i \frac{1}{n} \sum_{j=1}^{n} y_j(t+1)}{2d_i}, \tag{F.5}$$

where n is the number of players and $y_j(t+1)$ is as defined as above.

F.3 Exponentialized relative payoff sums

I use a nonlinear variant of the basic RPS model called the exponenialized RPS model. This model is also called the quantal-response learning model (Mookherjee & Sopher, 1997), which is a dynamic learning version of the quantal-response equilibrium model of McKelvey and Palfrey (1998). Unlike the learning models above, in RPS players are not assumed to choose best replies to the expected behavior of others. This characteristic makes the model appealing for network games and other Internet settings, where the number of players may be large and unknown and there is high uncertainty about the players' type distribution.

To implement RPS, I divide the strategy space, $[0, 21]$, into n intervals. Let $j = 0, 1, 2, 3, \ldots, n$ correspond to the strategies of choosing the number $0, \frac{21}{n}, 2\frac{21}{n}, 3\frac{21}{n}, \ldots, 21$. If player i plays strategy j in round t, denote the payoff she earns in that round by $\pi_{ij}(t)$. For all $k \neq j$, $\pi_{ij}(t) = 0$.

Define $M_{ij}(t)$ as the discounted payoff sum of player i to choose strategy j:

$$M_{ij}(t) = qM_{ij}(t-1) + \pi_{ij}(t), \tag{F.6}$$

where $q \in [0, 1]$ is the time/memory discount factor. Then the probability that player i plays strategy j in round $t+1$ is given by

$$p_{ij}(t+1) = \frac{e^{\lambda M_{ij}(t)}}{\sum_{k=0}^{n} e^{\lambda M_{ik}(t)}} \qquad \forall i, j, \tag{F.7}$$

where $\lambda \geq 0$ helps to scale up (if $\lambda \geq 1$) or scale down ($0 \leq \lambda < 1$) the relative weights of the discounted payoff sums. When $\lambda = 0$, the model degenerates into a random choice model.

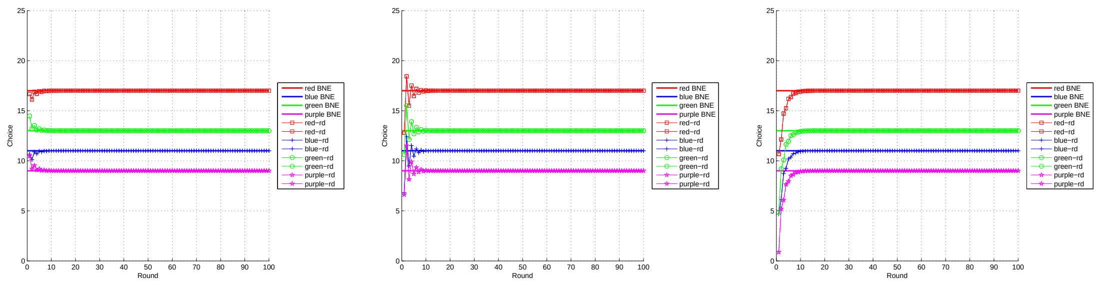

(a) SubNP (b) SubP (c) SuperNP

Figure F.1: Simulation of replicator dynamic.

In my computer simulations, $n = 100$, $\lambda = 0.004$. I also imposed an upper bound on $M_{ij}(t)$ equal to 1,000,000 to prevent the exponential terms from exploding. Since I conducted simulations before running experimental sessions, I didn't have data to estimate λ. I chose the value 0.004 by trial and error, using results from another study to guide the search (Chen & Tang, 1998). The value of 0.004 was one of few values that produced dynamics different from a constant choice (typically equal to the initial randomly chosen number) or what appeared to be random choice. Unfortunately, as Figure F.7 shows, the RPS model did not produce an interesting dynamic even with $\lambda = 0.004$ in the Cournot duopoly environment.

F.4 Simulation Results

Figures F.1–F.8 display simulated convergence dynamics for each learning model. There are eight players, each of one of the four types: red row, blue row, green column, and purple column. There are two players of each type, and the row players are randomly rematched with the column players in each round. The players' payoff functions are as described in Table 3.3, and the strategy space is $[0, 21]$ for each player. The initial choice is a random number from the strategy space.[1]

Figures F.2–F.6 show that the learning models from the FP family with relatively large γ converge to the Bayes-Nash equilibrium, and the convergence is fast. Since the equilibrium is a globally asymptotically stable steady state of the best-reply dynamic in all three environments, it is not surprising that the level of convergence is equally high in all three. The exception is the Cournot best-reply dynamic, which oscillates indefinitely, because the players do not "remember" enough history to learn

[1] As I discuss in Appendix D, the unrestricted best reply to some feasible choices may be outside the interval $[0, 21]$. However, in my computer simulations of the models from the FP family, I did not restrict the players to the strategy space except in the first round. As the figures show, the players do not choose numbers higher than 21 in more than a few cases in early rounds.

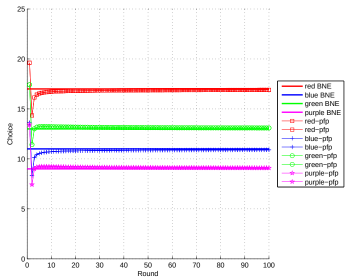

(a) SubNP

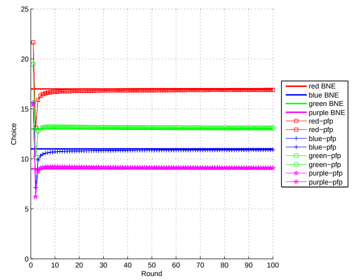

(b) SubP

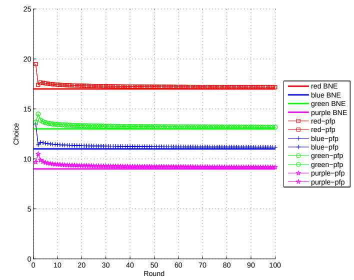

(c) SuperNP

Figure F.2: Simulation of population fictitious play.

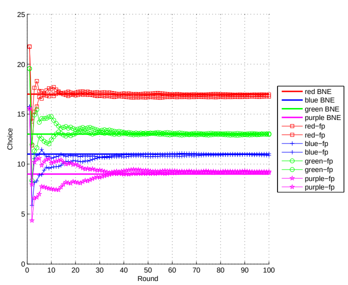

(a) SubNP

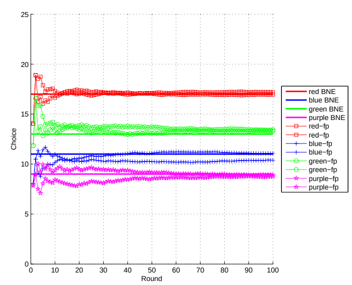

(b) SubP

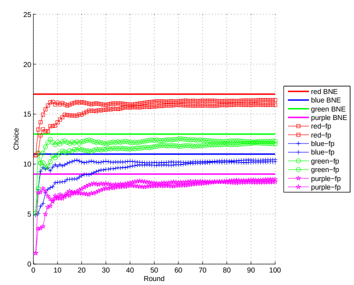

(c) SuperNP

Figure F.3: Simulation of fictitious play (equivalent to generalized fictitious play with $\gamma = 1$).

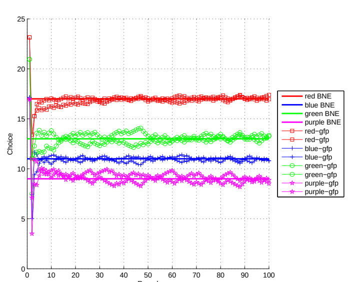

(a) SubNP

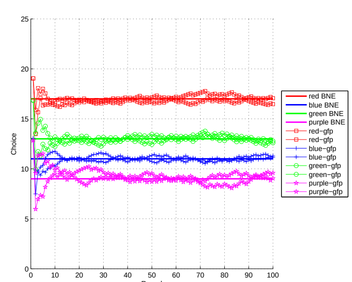

(b) SubP

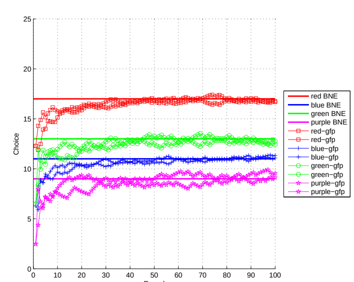

(c) SuperNP

Figure F.4: Simulation of generalized fictitious play with $\gamma = 0.9$.

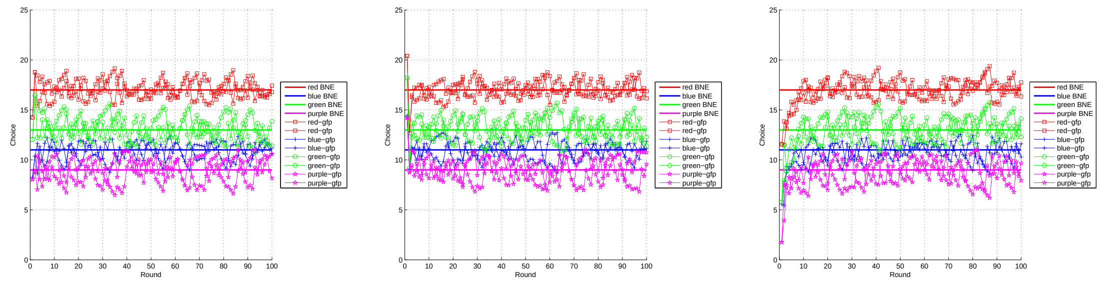

(a) SubNP (b) SubP (c) SuperNP

Figure F.5: Simulation of generalized fictitious play with $\gamma = 0.5$.

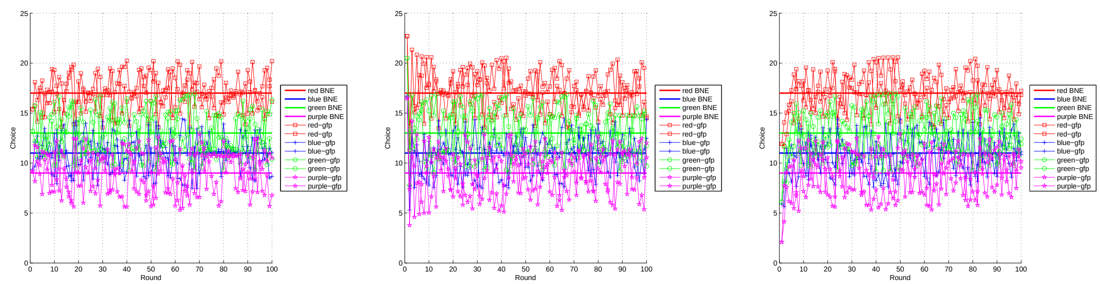

(a) SubNP (b) SubP (c) SuperNP

Figure F.6: Simulation of the Cournot best-reply dynamic (equivalent to generalized fictitious play with $\gamma = 0$).

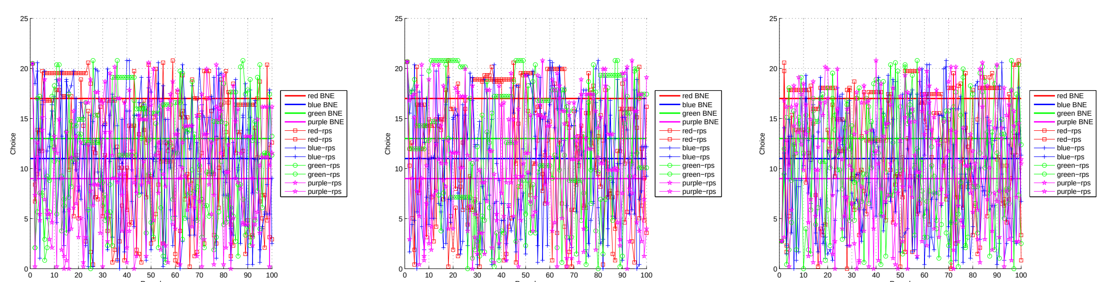

(a) SubNP (b) SubP (c) SuperNP

Figure F.7: Simulation of the exponential relative payoff sums dynamic (with $\gamma = 0.5$ and $\lambda = 0.004$).

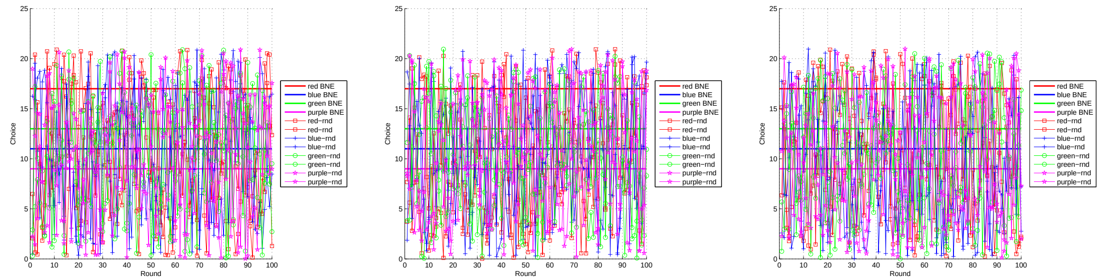

(a) SubNP (b) SubP (c) SuperNP

Figure F.8: Simulation of random play.

the type distribution of the other player, and therefore are unable to compute the appropriate best reply.

The replicator dynamic heuristic has the advantage of providing information about the type distribution in the very first round. As we see from the figures, this results in the fastest convergence to the equilibrium.

The RPS dynamic did not converge to the equilibrium. The most plausible explanation is that I failed to tune the model parameter λ appropriately, which proved to be difficult to do by trial and error. Since the probability is defined by the exponent, the players' choices are very sensitive to small changes in λ. If λ is small, the players put little weight on history and choose actions randomly from the strategy space. If λ is large, the players put all weight in the initial choice, also randomly chosen, and stick to it for the rest of the rounds. There is a very narrow interval of λ, for which the dynamic combines exploration and exploitation of the accumulated knowledge. Even then the players converge to something other than the equilibrium.[2]

[2]Letting λ grow gradually from 0 to 1 could help solve the problem, but since learning models is not the focus of this study, I did not implement the modification.

APPENDIX G

Experiment Instruction: Bayesian Information

Introduction[1]

- You are about to participate in an experiment intended to provide insight into certain features of decision-making processes. The experiment consists of 100 rounds. In each round you will interact with one of the other participants. You will earn money based on the decisions you and your match make. If you follow the instructions carefully and make good decisions, you may earn a considerable amount of money. Your decisions and earnings are confidential.

- The procedure is computerized. It includes a 15-minute tutorial with review questions, 100-round experiment, and a short questionnaire. After a start window appears, please click the "Begin Tutorial" button and follow the instructions on your screen. If you have a question, raise your hand and the experimenter will assist you.

- *During the experiment, we ask that you please do not talk to each other.* Please do not use any email client, web browser, or any other software program on the computer.

- Throughout the experiment, we use a fictitious experimental currency, called *points*. Your **total earning** will be the sum of the total number of points you earn, converted to U.S. dollars and rounded up to the nearest dollar amount. The conversion rate is

  $$\$1 = 1700 \text{ points.}$$

  In addition, you will be paid a $5 show-up fee.

[1] The original formatting of the instruction is modified to meet Rackham dissertation formatting requirements.

Types

- There are eight participants of four different types in this experiment.
- Row and Column types:
  - At the beginning of the experiment, a computer program will randomly choose half of the participants to be Row types, and the rest will be Column types.
  - If you are a Row type, you will remain a Row type for the rest of the experiment. Similarly, if you are a Column type, you will remain a Column type for the rest of the experiment.
  - A Row type will interact only with a Column type, and vice verse.
- Color types: After each participant is assigned a Row or Column type, the computer will subdivide them into color types. Each type is divided into two color types.
  - Rows: The computer will randomly choose 50% of the Rows and assign them the RED Row type, and the rest of the Rows will be the BLUE Row type.
  - Columns: Similarly, the computer will randomly choose 50% of the Columns and assign them the GREEN Column type, and the rest will be the PURPLE Column type.
  - Each participant's color type will remain the same for the first 50 rounds of the experiment.
  - After the first 50 rounds, RED Rows will become BLUE Rows and vice verse. Similarly, GREEN Columns will become PURPLE Columns and vice verse.
- Different types have different payoff tables. Please open your folder and find four payoff tables. After the computer assigns you a type, you will learn which tables are yours and how to read them in the computerized tutorial. At the end of the tutorial your understanding will be tested with review questions. You will earn 100 points for each correct answer.

Matching

- At the beginning of each round, each Row participant will be randomly matched with a Column participant. Therefore, your match may change in every round. The type of your match will *NOT* be revealed to you.

- If you are a Row, you will be equally likely to be matched with a GREEN or PURPLE Column. Similarly, if you are a Column, you will be equally likely to be matched with a RED or BLUE Row.

Choices and Payoffs

- In each round, you and your match will independently choose a number (a multiple of 0.01) from 0 to 21. Your payoff depends on your own choice and the choice made by your match. The payoff tables display payoffs for integer choices from 0 to 21. However, you are not restricted to integer choices. Also, the payoffs in the tables are rounded to the nearest integer for better readability. You will learn how to read the tables in the computerized tutorial. The tutorial includes two detailed examples and explains how to interpret the curved lines on the payoff-table sheets.

- A Row type will choose rows, whereas a Column type will choose columns in the payoff tables.

  For example, if a Row type chooses 3 and the Column type with whom she is matched chooses 6, their payoffs (in points) are written at the intersection of row 3 and column 6.

- There are two numbers written at the intersection of each row and column. **The top number in bold is always the Row type's payoff** and *the bottom number in italic is always the Column type's payoff.*

- You can compute payoffs for *any* choice pair, including non-integer choices that are not given in the table, using a tool on your input screen called the "What-If Scenario Analyzer". See Figure G.1 for a screenshot of the input screen. The analyzer is a simple calculator that allows you to type in any pair of numbers for your choice and your match's choice and view the resulting payoffs by clicking the "Display Hypothetical Payoffs" button. Also, the "What-If Scenario Analyzer" has a higher payoff precision than the payoff tables. You will practice using the analyzer in the tutorial.

- In each round, there is a recommended time limit, which is not binding. The header of your input screen will display the current round and the time remaining for making a choice. If you exceed the recommended time limit, a warning urging you to make a decision will start to flash.

- Your **total points** is the sum of your points earned in each round plus the points earned for correctly answered review questions at the end of the tutorial.

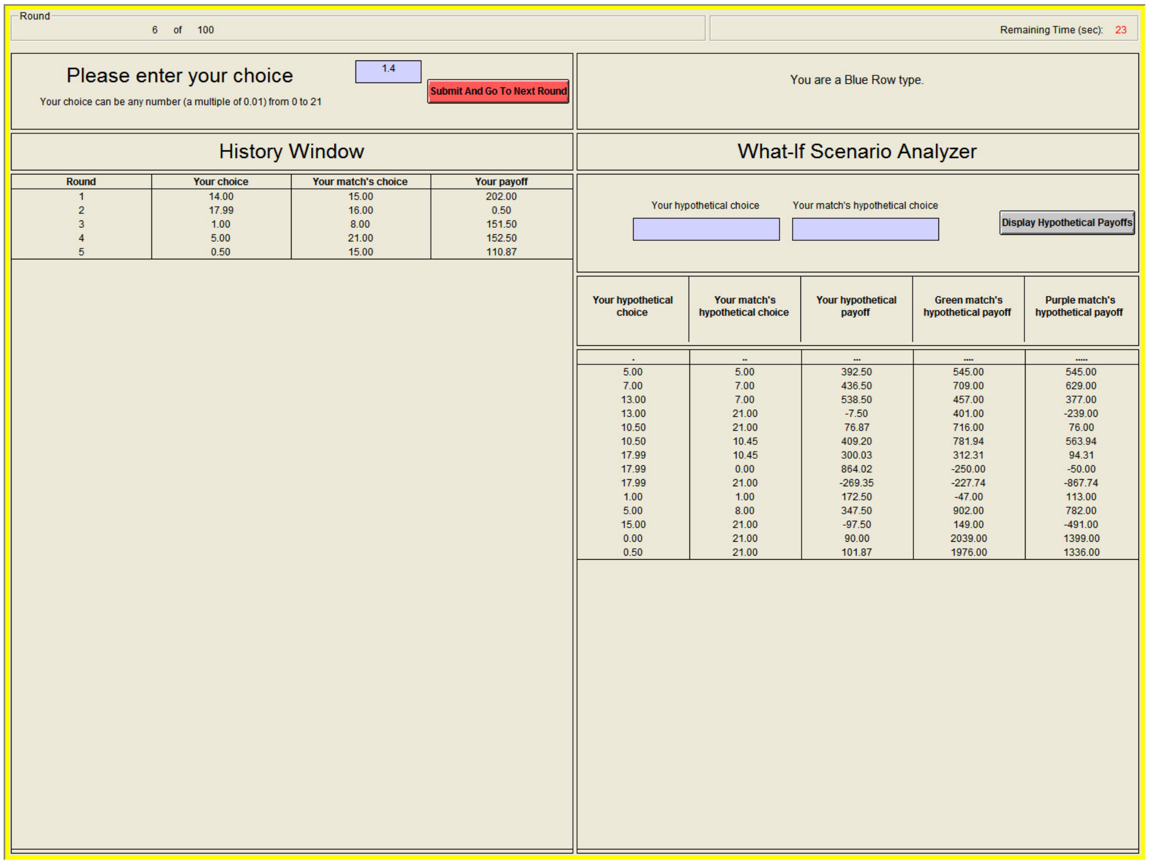

Figure G.1: Input screen.

- If you earn negative points in any given round, they will be deducted from your total points.

Feedback and History

- At the end of each round, you will be shown the results of that round, which include the following information: your choice, your match's choice, and your payoff.

- Throughout the experiment, you will be able to view the results of all previous rounds in the history window (see the screenshot on Figure G.1). The "What-If Scenario Analyzer" will also save and display the entire history of your calculations.

  Note that because your color type will change after the first 50 rounds, round 51 will start with a new choice history and an empty analyzer window.

APPENDIX H

Learning Bayesian Nash Equilibrium: Payoff Tables in Bayesian-Information Condition

Below are examples of payoff-table paper handouts that subjects received under Bayesian-information condition. In each Bayesian-information treatment, subjects each received identical sets of four payoff tables, one table for each possible type profile. For the experiment instructions, see Appendix G.

Red Row's Choice \ Green Column's Choice	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
0	-50 / -250	-50 / -41	-50 / 158	-50 / 347	-50 / 526	-50 / 695	-50 / 854	-50 / 1003	-50 / 1142	-50 / 1271	-50 / 1390	-50 / 1499	-50 / 1598	-50 / 1687	-50 / 1766	-50 / 1835	-50 / 1894	-50 / 1943	-50 / 1982	-50 / 2011	-50 / 2030	-50 / 2039
1	66 / -250	63 / -47	60 / 146	57 / 329	54 / 502	51 / 665	48 / 818	45 / 961	42 / 1094	39 / 1217	36 / 1330	33 / 1433	30 / 1526	27 / 1609	24 / 1682	21 / 1745	18 / 1798	15 / 1841	12 / 1874	9 / 1897	6 / 1910	3 / 1913
2	176 / -250	170 / -53	164 / 134	158 / 311	152 / 478	146 / 635	140 / 782	134 / 919	128 / 1046	122 / 1163	116 / 1270	110 / 1367	104 / 1454	98 / 1531	92 / 1598	86 / 1655	80 / 1702	74 / 1739	68 / 1766	62 / 1783	56 / 1790	50 / 1787
3	282 / -250	273 / -59	264 / 122	255 / 293	246 / 454	237 / 605	228 / 746	219 / 877	210 / 998	201 / 1109	192 / 1210	183 / 1301	174 / 1382	165 / 1453	156 / 1514	147 / 1565	138 / 1606	129 / 1637	120 / 1658	111 / 1669	102 / 1670	93 / 1661
4	382 / -250	370 / -65	358 / 110	346 / 275	334 / 430	322 / 575	310 / 710	298 / 835	286 / 950	274 / 1055	262 / 1150	250 / 1235	238 / 1310	226 / 1375	214 / 1430	202 / 1475	190 / 1510	178 / 1535	166 / 1550	154 / 1555	142 / 1550	130 / 1535
5	478 / -250	463 / -71	448 / 98	433 / 257	418 / 406	403 / 545	388 / 674	373 / 793	358 / 902	343 / 1001	328 / 1090	313 / 1169	298 / 1238	283 / 1297	268 / 1346	253 / 1385	238 / 1414	223 / 1433	208 / 1442	193 / 1441	178 / 1430	163 / 1409
6	568 / -250	550 / -77	532 / 86	514 / 239	496 / 382	478 / 515	460 / 638	442 / 751	424 / 854	406 / 947	388 / 1030	370 / 1103	352 / 1166	334 / 1219	316 / 1262	298 / 1295	280 / 1318	262 / 1331	244 / 1334	226 / 1327	208 / 1310	190 / 1283
7	654 / -250	633 / -83	612 / 74	591 / 221	570 / 358	549 / 485	528 / 602	507 / 709	486 / 806	465 / 893	444 / 970	423 / 1037	402 / 1094	381 / 1141	360 / 1178	339 / 1205	318 / 1222	297 / 1229	276 / 1226	255 / 1213	234 / 1190	213 / 1157
8	734 / -250	710 / -89	686 / 62	662 / 203	638 / 334	614 / 455	590 / 566	566 / 667	542 / 758	518 / 839	494 / 910	470 / 971	446 / 1022	422 / 1063	398 / 1094	374 / 1115	350 / 1126	326 / 1127	302 / 1118	278 / 1099	254 / 1070	230 / 1031
9	810 / -250	783 / -95	756 / 50	729 / 185	702 / 310	675 / 425	648 / 530	621 / 625	594 / 710	567 / 785	540 / 850	513 / 905	486 / 950	459 / 985	432 / 1010	405 / 1025	378 / 1030	351 / 1025	324 / 1010	297 / 985	270 / 950	243 / 905
10	880 / -250	850 / -101	820 / 38	790 / 167	760 / 286	730 / 395	700 / 494	670 / 583	640 / 662	610 / 731	580 / 790	550 / 839	520 / 878	490 / 907	460 / 926	430 / 935	400 / 934	370 / 923	340 / 902	310 / 871	280 / 830	250 / 779
11	946 / -250	913 / -107	880 / 26	847 / 149	814 / 262	781 / 365	748 / 458	715 / 541	682 / 614	649 / 677	616 / 730	583 / 773	550 / 806	517 / 829	484 / 842	451 / 845	418 / 838	385 / 821	352 / 794	319 / 757	286 / 710	253 / 653
12	1006 / -250	970 / -113	934 / 14	898 / 131	862 / 238	826 / 335	790 / 422	754 / 499	718 / 566	682 / 623	646 / 670	610 / 707	574 / 734	538 / 751	502 / 758	466 / 755	430 / 742	394 / 719	358 / 686	322 / 643	286 / 590	250 / 527
13	1062 / -250	1023 / -119	984 / 2	945 / 113	906 / 214	867 / 305	828 / 386	789 / 457	750 / 518	711 / 569	672 / 610	633 / 641	594 / 662	555 / 673	516 / 674	477 / 665	438 / 646	399 / 617	360 / 578	321 / 529	282 / 470	243 / 401
14	1112 / -250	1070 / -125	1028 / -10	1028 / 95	986 / 190	902 / 275	860 / 350	818 / 415	776 / 470	734 / 515	692 / 550	650 / 575	608 / 590	566 / 595	524 / 590	482 / 575	440 / 550	398 / 515	356 / 470	314 / 415	272 / 350	230 / 275
15	1158 / -250	1113 / -131	1068 / -22	1023 / 77	978 / 166	933 / 245	888 / 314	843 / 373	798 / 422	753 / 461	708 / 490	663 / 509	618 / 518	573 / 517	528 / 506	483 / 485	438 / 454	393 / 413	348 / 362	303 / 301	258 / 230	213 / 149
16	1198 / -250	1150 / -137	1102 / -34	1054 / 59	1006 / 142	958 / 215	910 / 278	862 / 331	814 / 374	766 / 407	718 / 430	670 / 443	622 / 446	574 / 439	526 / 422	478 / 395	430 / 358	382 / 311	334 / 254	286 / 187	238 / 110	190 / 23
17	1234 / -250	1183 / -143	1132 / -46	1081 / 41	1030 / 118	979 / 185	928 / 242	877 / 289	826 / 326	775 / 353	724 / 370	673 / 377	622 / 374	571 / 361	520 / 338	469 / 305	418 / 262	367 / 209	316 / 146	265 / 73	214 / -10	163 / -103
18	1264 / -250	1210 / -149	1156 / -58	1102 / 23	1048 / 94	994 / 155	940 / 206	886 / 247	832 / 278	778 / 299	724 / 310	670 / 311	616 / 302	562 / 283	508 / 254	454 / 215	400 / 166	346 / 107	292 / 38	238 / -41	184 / -130	130 / -229
19	1290 / -250	1233 / -155	1176 / -70	1119 / 5	1062 / 70	1005 / 125	948 / 170	891 / 205	834 / 230	777 / 245	720 / 250	663 / 245	606 / 230	549 / 205	492 / 170	435 / 125	378 / 70	321 / 5	264 / -70	207 / -155	150 / -250	93 / -355
20	1310 / -250	1250 / -161	1190 / -82	1130 / -13	1070 / 46	1010 / 95	950 / 134	890 / 163	830 / 182	770 / 191	710 / 190	650 / 179	590 / 158	530 / 127	470 / 86	410 / 35	350 / -26	290 / -97	230 / -178	170 / -269	110 / -370	50 / -481
21	1326 / -250	1263 / -167	1200 / -94	1137 / -31	1074 / 22	1011 / 65	948 / 98	885 / 121	822 / 134	759 / 137	696 / 130	633 / 113	570 / 86	507 / 49	444 / 2	381 / -55	318 / -122	255 / -199	192 / -286	129 / -383	66 / -490	3 / -607

Figure H.1: Submodular & non-potential environment (SubNP): {red row, green column}.

Blue Row's Choice	Purple Column's Choice: 0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
0	-200 / 200	-200 / 201	-200 / 196	-200 / 187	-200 / 172	-200 / 153	-200 / 128	-200 / 99	-200 / 64	-200 / 25	-200 / -20	-200 / -70	-200 / -124	-200 / -184	-200 / -248	-200 / -318	-200 / -392	-200 / -472	-200 / -556	-200 / -646	-200 / -740	-200 / -840
1	-161 / 200	-155 / 204	-149 / 202	-143 / 196	-137 / 184	-131 / 168	-125 / 146	-119 / 120	-113 / 88	-107 / 52	-101 / 10	-95 / -37	-89 / -88	-83 / -145	-77 / -206	-71 / -273	-65 / -344	-59 / -421	-53 / -502	-47 / -589	-41 / -680	-35 / -777
2	-132 / 200	-120 / 207	-108 / 208	-96 / 205	-84 / 196	-72 / 183	-60 / 164	-48 / 141	-36 / 112	-24 / 79	-12 / 40	0 / -4	12 / -52	24 / -106	36 / -164	48 / -228	60 / -296	72 / -370	84 / -448	96 / -532	108 / -620	120 / -714
3	-113 / 200	-95 / 210	-77 / 214	-59 / 214	-41 / 208	-23 / 198	-5 / 182	13 / 162	31 / 136	49 / 106	67 / 70	85 / 30	103 / -16	121 / -67	139 / -122	157 / -183	175 / -248	193 / -319	211 / -394	229 / -475	247 / -560	265 / -651
4	-104 / 200	-80 / 213	-56 / 220	-32 / 223	-8 / 220	16 / 213	40 / 200	64 / 183	88 / 160	112 / 133	136 / 100	160 / 63	184 / 20	208 / -28	232 / -80	256 / -138	280 / -200	304 / -268	328 / -340	352 / -418	376 / -500	400 / -588
5	-105 / 200	-75 / 216	-45 / 226	-15 / 232	15 / 232	45 / 228	75 / 218	105 / 204	135 / 184	165 / 160	195 / 130	225 / 96	255 / 56	285 / 12	315 / -38	345 / -93	375 / -152	405 / -217	435 / -286	465 / -361	495 / -440	525 / -525
6	-116 / 200	-80 / 219	-44 / 232	-8 / 241	28 / 244	64 / 243	100 / 236	136 / 225	172 / 208	208 / 187	244 / 160	280 / 129	316 / 92	352 / 51	388 / 4	424 / -48	460 / -104	496 / -166	532 / -232	568 / -304	604 / -380	640 / -462
7	-137 / 200	-95 / 222	-53 / 238	-11 / 250	31 / 256	73 / 258	115 / 254	157 / 246	199 / 232	241 / 214	283 / 190	325 / 162	367 / 128	409 / 90	451 / 46	493 / -3	535 / -56	577 / -115	619 / -178	661 / -247	703 / -320	745 / -399
8	-168 / 200	-120 / 225	-72 / 244	-24 / 259	24 / 268	72 / 273	120 / 272	168 / 267	216 / 256	264 / 241	312 / 220	360 / 195	408 / 164	456 / 129	504 / 88	552 / 43	600 / -8	648 / -64	696 / -124	744 / -190	792 / -260	840 / -336
9	-209 / 200	-155 / 228	-101 / 250	-47 / 268	7 / 280	61 / 288	115 / 290	169 / 288	223 / 280	277 / 268	331 / 250	385 / 228	439 / 200	493 / 168	547 / 130	601 / 88	655 / 40	709 / -13	763 / -70	817 / -133	871 / -200	925 / -273
10	-260 / 200	-200 / 231	-140 / 256	-80 / 277	-20 / 292	40 / 303	100 / 308	160 / 309	220 / 304	280 / 295	340 / 280	400 / 261	460 / 236	520 / 207	580 / 172	640 / 133	700 / 88	760 / 39	820 / -16	880 / -76	940 / -140	1000 / -210
11	-321 / 200	-255 / 234	-189 / 262	-123 / 286	-57 / 304	9 / 318	75 / 326	141 / 330	207 / 328	273 / 322	339 / 310	405 / 294	471 / 272	537 / 246	603 / 214	669 / 178	735 / 136	801 / 90	867 / 38	933 / -19	999 / -80	1065 / -147
12	-392 / 200	-320 / 237	-248 / 268	-176 / 295	-104 / 316	-32 / 333	40 / 344	112 / 351	184 / 352	256 / 349	328 / 340	400 / 327	472 / 308	544 / 285	616 / 256	688 / 223	760 / 184	832 / 141	904 / 92	976 / 39	1048 / -20	1120 / -84
13	-473 / 200	-395 / 240	-317 / 274	-239 / 304	-161 / 328	-83 / 348	-5 / 362	73 / 372	151 / 376	229 / 376	307 / 370	385 / 360	463 / 344	541 / 324	619 / 298	697 / 268	775 / 232	853 / 192	931 / 146	1009 / 96	1087 / 40	1165 / -21
14	-564 / 200	-480 / 243	-396 / 280	-396 / 313	-312 / 340	-144 / 363	-60 / 380	24 / 393	108 / 400	192 / 403	276 / 400	360 / 393	444 / 380	528 / 363	612 / 340	696 / 313	780 / 280	864 / 243	948 / 200	1032 / 153	1116 / 100	1200 / 43
15	-665 / 200	-575 / 246	-485 / 286	-395 / 322	-305 / 352	-215 / 378	-125 / 398	-35 / 414	55 / 424	145 / 430	235 / 430	325 / 426	415 / 416	505 / 402	595 / 382	685 / 358	775 / 328	865 / 294	955 / 254	1045 / 210	1135 / 160	1225 / 106
16	-776 / 200	-680 / 249	-584 / 292	-488 / 331	-392 / 364	-296 / 393	-200 / 416	-104 / 435	-8 / 448	88 / 457	184 / 460	280 / 459	376 / 452	472 / 441	568 / 424	664 / 403	760 / 376	856 / 345	952 / 308	1048 / 267	1144 / 220	1240 / 169
17	-897 / 200	-795 / 252	-693 / 298	-591 / 340	-489 / 376	-387 / 408	-285 / 434	-183 / 456	-81 / 472	21 / 484	123 / 490	225 / 492	327 / 488	429 / 480	531 / 466	633 / 448	735 / 424	837 / 396	939 / 362	1041 / 324	1143 / 280	1245 / 232
18	-1028 / 200	-920 / 255	-812 / 304	-704 / 349	-596 / 388	-488 / 423	-380 / 452	-272 / 477	-164 / 496	-56 / 511	52 / 520	160 / 525	268 / 524	376 / 519	484 / 508	592 / 493	700 / 472	808 / 447	916 / 416	1024 / 381	1132 / 340	1240 / 295
19	-1169 / 200	-1055 / 258	-941 / 310	-827 / 358	-713 / 400	-599 / 438	-485 / 470	-371 / 498	-257 / 520	-143 / 538	-29 / 550	85 / 558	199 / 560	313 / 558	427 / 550	541 / 538	655 / 520	769 / 498	883 / 470	997 / 438	1111 / 400	1225 / 358
20	-1320 / 200	-1200 / 261	-1080 / 316	-960 / 367	-840 / 412	-720 / 453	-600 / 488	-480 / 519	-360 / 544	-240 / 565	-120 / 580	0 / 591	120 / 596	240 / 597	360 / 592	480 / 583	600 / 568	720 / 549	840 / 524	960 / 495	1080 / 460	1200 / 421
21	-1481 / 200	-1355 / 264	-1229 / 322	-1103 / 376	-977 / 424	-851 / 468	-725 / 506	-599 / 540	-473 / 568	-347 / 592	-221 / 610	-95 / 624	31 / 632	157 / 636	283 / 634	409 / 628	535 / 616	661 / 600	787 / 578	913 / 552	1039 / 520	1165 / 484

Figure H.2: Supermodular & non-potential environment (SuperNP): {blue row, purple column}.

APPENDIX I

Experiment Instruction: Low Information

Introduction[1]

- You are about to participate in an experiment intended to provide insight into certain features of decision-making processes. The experiment consists of 100 rounds. In each round you will interact with one of the other participants. You will earn money based on the decisions you and your match make. If you follow the instructions carefully and make good decisions, you may earn a considerable amount of money. Your decisions and earnings are confidential.

- The procedure is computerized. It includes a 100-round experiment and a short questionnaire. After a start window appears, please follow the instructions on your screen. If you have a question, raise your hand and the experimenter will assist you.

- *During the experiment, we ask that you please do not talk to each other.* Please do not use any email client, web browser, or any other software program on the computer.

- Throughout the experiment, we use a fictitious experimental currency, called *points*. Your **total earning** will be the sum of the total number of points you earn, converted to U.S. dollars and rounded up to the nearest dollar amount. The conversion rate is

  $$\$1 = 1700 \text{ points.}$$

  In addition, you will be paid a $5 show-up fee.

[1] The original formatting of the instruction is modified to meet Rackham dissertation formatting requirements.

Procedure

- There are eight participants in this experiment. At the beginning of each round, each participant will be randomly matched with another participant. Therefore, your match may change in every round.

- In each round, you and your match will independently choose a number (a multiple of 0.01) from 0 to 21. Note that you are not restricted to integer choices.

- Your payoff depends on your own choice and the choice made by your match. Your payoff will be displayed on your screen at the end of each round after you and your match have made your choices.

- Some participants may have the same payoff structure as you do, while others may have different payoffs.

  For example, suppose you and your match both choose 11.11. Even though your choices are the same, your payoff may or may not be equal to the payoff of your match. You will only be shown your own payoff. Your match's payoff will *NOT* be revealed to you at any point in the experiment.

- Your payoff structure may change during the experiment. You will not be informed when that happens.

- In each round, there is a recommended time limit, which is not binding. The header of your input screen will display the current round and the time remaining for making a choice. If you exceed the recommended time limit, a warning urging you to make a decision will start to flash.

- Your **total points** is the sum of your points earned in each round. If you earn negative points in any given round, they will be deducted from your total points.

Feedback and History

- At the end of each round, you will be shown the results of that round, which include the following information: your choice, your match's choice, and your payoff.

- Throughout the experiment, you will be able to view the results of all previous rounds in the history window on your screen.

BIBLIOGRAPHY

Adams, W. J., & Yellen, J. L. (1976). Commodity bundling and the burden of monopoly. *Quarterly Journal of Economics*, *90*(3), 475–498.

Anderson, S. P., Goeree, J. K., & Holt, C. A. (2001). Minimum-effort coordination games: Stochastic potential and logit equilibrium. *Games and Economic Behavior*, *34*, 177–199.

Armantier, O., Florens, J.-P., & Richard, J.-F. (2000). Empirical game-theoretic models: Constrained equilibrium and simulations. Tech. rep., State University of New York at Stonybrook.

Aron, D. J., & Wildman, S. S. (1999). Effecting a price squeeze through bundled pricing. In Gillett, S. E., & Vogelsang, I. (Eds.), *Competition, Regulation and Convergence: Current Trends in Telecommunications Policy Research*. Lawrence Erlbaum Associates, Mahwah, New Jersey.

Arrow, K. J., & Hahn, F. H. (1971). *General Competitive Analysis*. Holden-Day, San Francisco.

Ausubel, L. M., & Cramton, P. (2002). Demand reduction and inefficiency in multi-unit auctions. University of Maryland Working Paper 96-07, revised July 2002.

Bakos, Y., & Brynjolfsson, E. (1999). Bundling information goods: Pricing, profits and efficiency. *Management Science*, *45*(12), 1613–1630.

Bakos, Y., & Brynjolfsson, E. (2000). Aggregation and disaggregation of information goods: Implications for bundling, site licensing and micropayment systems. In Hey, D., Kahin, B., & Varian, H. (Eds.), *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*. MIT Press.

Bakos, Y., & Brynjolfsson, E. (1998). Bundling information goods: Pricing, profits and efficiency. In Hurley, D., Kahin, B., & Varian, H. (Eds.), *The Economics of Digital Information Goods*. MIT Press, Cambridge, Massachusetts.

Bakos, Y., & Brynjolfsson, E. (1999). Bundling and competition on the Internet: Aggregation strategies for information goods. Working paper series, MIT Sloan School of Management.

Bergemann, D., & Valimaki, J. (2006). Information in mechanism design. Working paper. Available at SSRN: http://ssrn.com/abstract=876970.

Bertsekas, D. P. (1992). Auction algorithms for network flow problems: A tutorial introduction. *Computational Optimization and Applications*, *1*, 7–66.

Bikhchandani, S., & Mamer, J. W. (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory*, *74*, 385–413.

Birkhoff, G. (1967 (1940)). *Lattice Theory*. American Mathematical Society.

Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior*, *5*(3), 387–424. Available at http://ideas.repec.org/a/eee/gamebe/v5y1993i3p387-424.html.

Brooks, C. H., Fay, S., Das, R., MacKie-Mason, J. K., Kephart, J. O., & Durfee, E. H. (1999). Automated strategy searches in an electronic goods market: learning and complex price schedules. In *First ACM Conference on Electronic Commerce*, pp. 31–40. Available at http://doi.acm.org/10.1145/336992.337000.

Brooks, C. H., Gazzale, R., Das, R., Kephart, J., MacKie-Mason, J. K., & Durfee, E. H. (2002). Model selection in an information economy: Choosing what to learn. *Computational Intelligence*, *18*(4), 566–582.

Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton University Press.

Carbajo, J., Meza, D. D., & Seidmann, D. J. (1990). A strategic motivation for commodity bundling. *Journal of Industrial Economics*, *38*(3), 283–298. Available at http://ideas.repec.org/a/bla/jindec/v38y1990i3p283-98.html.

Chen, Y. (2003). An experimental study of serial and average cost pricing mechanisms. *Journal of Public Economics*, *87*(9-10), 2305–2335. Available at http://ideas.repec.org/a/eee/pubeco/v87y2003i9-10p2305-2335.html.

Chen, Y., & Gazzale, R. (2004). When does learning in games generate convergence to Nash equilibria? The role of supermodularity in an experimental setting. *American Economic Review*, *94*(5), 1505–1535. Available at http://dx.doi.org/10.1257/0002828043052349.

Chen, Y., & Plott, C. R. (1996). The Groves-Ledyard mechanism: An experimental study of institutional design. *Journal of Public Economics*, *59*, 335–364.

Chen, Y., Razzolini, L., & Turocy, T. (2007). Congestion allocation mechanisms for distributed networks: An experimental study. *Economic Theory*, *33*, 121–143.

Chen, Y., & Tang, F.-F. (1998). Learning and incentive-compatible mechanisms for public goods provision: An experimental study. *Journal of Political Economy*, *106*(3), 633–662. Available at http://ideas.repec.org/a/ucp/jpolec/v106y1998i3p633-662.html.

Cheng, S.-F., Leung, E., Lochner, K. M., O'Malley, K., Reeves, D. M., & Wellman, M. P. (2005). Walverine: A Walrasian trading agent. *Decision Support Systems*, *39*, 169–184.

Chuang, J. C.-I., & Sirbu, M. A. (1998). Network delivery of information goods: Optimal pricing of articles and subscriptions. In Hurley, D., B.Kahin, & Varian, H. (Eds.), *The Economics of Digital Information Goods*. MIT Press, Cambridge, Massachusetts.

Coase, R. (1960). The problem of social cost. *Journal of Law and Economics*, *III*, 1–44.

Cramton, P. (1995). Money out of thin air: The nationwide narrowband PCS auction. *Journal of Economics and Management Strategy*, *4*, 267–343.

Cramton, P. (2005). Simultaneous ascending auctions. In Cramton et al. (Cramton et al., 2005).

Cramton, P., Shoham, Y., & Steinberg, R. (Eds.). (2005). *Combinatorial Auctions*. MIT Press.

Csirik, J. A., Littman, M. L., Singh, S., & Stone, P. (2001). FAucS: An FCC spectrum auction simulator for autonomous bidding agents. In *Second International Workshop on Electronic Commerce*, Vol. 2232 of *Lecture Notes in Computer Science*, pp. 139–151. Springer-Verlag.

de Vries, S., & Vohra, R. (2003). Combinatorial auctions: A survey. *INFORMS Journal on Computing*, *15*(3), 284–309.

Dekel, E., Fudenberg, D., & Levine, D. K. (2004). Learning to play Bayesian games. *Games and Economic Behavior*, *46*(2), 282–303. Available at http://dx.doi.org/10.1016/S0899-8256(03)00121-0.

Demsetz, H. (1968). The cost of transacting. *Quarterly Journal of Economics*, *LXXXII*, 33–53.

Dixit, A. K., & Stiglitz, J. E. (1977). Monopolistic competition and optimum product diversity. *American Economic Review*, *67*(3), 297–308.

Economides, N., & Viard, B. (2004). Pricing of complementary goods and network effects. Tech. rep. 0407005, Economics Working Paper Archive at WUSTL. Available at http://ideas.repec.org/p/wpa/wuwpio/0407005.html.

Ely, J. C., & Sandholm, W. H. (2005). Evolution in Bayesian games I: Theory. *Games and Economic Behavior*, *53*(1), 83–109. Available at http://ideas.repec.org/a/eee/gamebe/v53y2005i1p83-109.html.

Erev, I., & Haruvy, E. (2008). Learning and the economics of small decisions. In Kagel, J. H., & Roth, A. E. (Eds.), *The Handbook of Experimental Economics*, Vol. 2, chap. 6. Princeton University Press. Forthcoming.

Facchini, G., van Magen, F., Borm, P., & Tijs, S. (1997). Congestion models and weighted Bayesian potential games. *Theory and Decision*, *42*, 193–206.

Farrell, J., Monroe, H. K., & Saloner, G. (1998). The vertical organization of industry: Systems competition versus component competition. *Journal of Economics and Management Strategy*, *7*(2), 143–182.

Fay, S. A. (2001). *Bundling Information Goods in Competitive Environments*. Ph.D. thesis, University of Michigan, Department of Economics.

Fay, S. A., & MacKie-Mason, J. K. (1999). Competition between firms that bundle information goods. In *27th Annual Telecommunications Policy Research Conference*. Alexandria, VA. Available at SSRN: http://ssrn.com/abstract=975733.

Fay, S. A., MacKie-Mason, J. K., & Osepayshvili, A. (2006). Competition between firms that bundle information goods. Working paper.

Fehr, E., & Tyran, J.-R. (2001). Does money illusion matter?. *American Economic Review*, *91*(5), 1239–1262. Available at http://ideas.repec.org/a/aea/aecrev/v91y2001i5p1239-1262.html.

Fehr, E., & Tyran, J.-R. (2004a). Limited rationality and strategic interaction - The impact of the strategic environment on nominal inertia. Levine's bibliography 122247000000000092, UCLA Department of Economics. Available at http://ideas.repec.org/p/cla/levrem/122247000000000092.html.

Fehr, E., & Tyran, J.-R. (2004b). Money illusion and coordination failure. IZA discussion papers 1013, Institute for the Study of Labor (IZA). Available at http://ideas.repec.org/p/iza/izadps/dp1013.html.

Fehr, E., & Tyran, J.-R. (2005). Individual irrationality and aggregate outcomes. *Journal of Economic Perspectives*, *19*(4), 43–66. Available at http://dx.doi.org/10.1257/089533005775196651.

Fischbacher, U. (2007). Z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, *10*(2), 171–178.

Fishburn, P. C., Odlyzko, A. M., & Siders, R. C. (2000). Fixed fee versus unit pricing for information goods: Competition, equilibria, and price wars. In Kahin, B., & Varian, H. R. (Eds.), *Internet Publishing and Beyond: The Economics of Digital Information and Intellectual Property*, pp. 167–189. MIT Press, Cambridge, MA.

Fudenberg, D., & Kreps, D. (1988). A theory of learning, experimentation, and equilibrium in games. Unpublished manuscript.

Fudenberg, D., & Levine, D. K. (1998). *The Theory of Learning in Games*. MIT Press.

Fudenberg, D., & Tirole, J. (1991). *Game Theory*. MIT Press.

Gigerenzer, G., & Selten, R. (Eds.). (2001). *Bounded Rationality: The Adaptive Toolbox*. MIT Press.

Greenwald, A., & Boyan, J. (2004). Bidding under uncertainty: Theory and experiments. In *Twentieth Conference on Uncertainty in Artificial Intelligence*, pp. 209–216, Banff.

Gul, F., & Stacchetti, E. (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, *87*, 95–124.

Hanson, W., & Martin, R. K. (1990). Optimal bundle pricing. *Management Science*, *36*, 155–174.

Harsanyi, J. C. (1967–1968). Games with incomplete information played by 'Bayesian' players. *Management Science*, *14*, 159–182, 320–334, 486–502.

Hirsch, M. (1985). Systems of differential equations which are competitive and cooperative II: Convergence almost everywhere. *SIAM Journal on Mathematical Analysis*, *16*, 423–439.

Hitt, L. M., & Chen, P. (2005). Bundling with customer self-selection: A simple approach to bundling low-marginal-cost goods. *Management Science*, *51*(10), 1481–1493. Available at http://dx.doi.org/10.1287/mnsc.1050.0403.

Hotelling, H. (1929). Stability in competition. *Economic Journal*, *39*, 41–57.

Johnston, J., & DiNardo, J. (1996). *Econometric Methods* (4 edition). McGraw-Hill/Irwin.

Kephart, J. O., Das, R., Brooks, C. H., Durfee, E. H., Gazzale, R. S., & MacKie-Mason, J. K. (2001). Pricing information bundles in a dynamic environment. In *Third ACM Conference on Electronic Commerce*, pp. 180–190.

Kephart, J. O., Das, R., & MacKie-Mason, J. K. (2000). Two-sided learning in an agent economy for information bundles. In *Agent-mediated Electronic Commerce, Lecture Notes in Artificial Intelligence*. Springer-Verlag, Berlin.

Kephart, J. O., Hanson, J. E., & Sairamesh, J. (1998). Price and niche wars in a free-market economy of software agents. *Artificial Life*, *4*, 1–23.

King, D. W., & Griffiths, J.-M. (1995). Economic issues concerning electronic publishing and distribution of scholarly articles. *Library Trends*, *43*(4), 713–740.

King, D. W., & Tenopir, C. (2005). Scholarly journal and digital database pricing: Threat or opportunity. Unpublished manuscript. Available at http://web.utk.edu/∼tenopir/pub/chapters.html.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Lehmann, B., Lehmann, D., & Nisan, N. (2006). Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, *55*, 270–296.

MacKie-Mason, J. K., Osepayshvili, A., Reeves, D. M., & Wellman, M. P. (2004). Price prediction strategies for market-based scheduling. In *Fourteenth International Conference on Automated Planning and Scheduling*, pp. 244–252, Whistler, BC.

MacKie-Mason, J. K., & Riveros, J. (1998). Economics and electronic access to scholarly information. In Hurley, D., Kahin, B., & Varian, H. (Eds.), *The Economics of Digital Information Goods*. MIT Press, Cambridge, Massachusetts.

MacKie-Mason, J. K., & Wellman, M. P. (2005). Automated markets and trading agents. In Tesfatsion, L., & Judd, K. (Eds.), *Handbook of Agent-Based Computational Economics*. Elsevier.

Matutes, C., & Regibeau, P. (1992). Compatibility and bundling of complementary goods in a duopoly. *Journal of Industrial Economics*, *40*(1), 37–53.

McAfee, R. P., & McMillan, J. (1996). Analyzing the airwaves auction. *Journal of Economic Perspectives*, *10*(1), 159–175.

McAfee, R. P., McMillan, J., & Whinston, M. D. (1989). Multiproduct monopoly, commodity bundling, and correlation of values. *Quarterly Journal of Economics*, *104*(2), 371–383. Available at http://ideas.repec.org/a/tpr/qjecon/v104y1989i2p371-83.html.

McKelvey, R. D., McLennan, A. M., & Turocy, T. L. (2007). Gambit: Software tools for game theory, version 0.2007.01.30. Available at http://gambit.sourceforge.net.

McKelvey, R. D., & Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, *1*(1), 9–41. Available at http://dx.doi.org/10.1023/A:1009905800005.

Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy*, *108*, 245–272.

Milgrom, P. (2003). *Putting Auction Theory to Work*. Cambridge University Press.

Milgrom, P., & Roberts, J. (1990). Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, *58*(6), 1255–1277.

Monderer, D., & Shapley, L. S. (1996a). Fictitious play property for games with identical interests. *Journal of Economic Theory*, *68*(1), 258–265. Available at http://ideas.repec.org/a/eee/jetheo/v68y1996i1p258-265.html.

Monderer, D., & Shapley, L. S. (1996b). Potential games. *Games and Economic Behavior*, *14*, 124–143.

Mookherjee, D., & Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior*, *19*(1), 97–132. Available at http://ideas.repec.org/a/eee/gamebe/v19y1997i1p97-132.html.

Nalebuff, B. J. (1999). Bundling. Working paper. Available at http://papers.ssrn.com/paper.taf?abstract_id=185193.

Nalebuff, B. J. (2000). Competiting against bundles. Working paper. Available at http://papers.ssrn.com/paper.taf?abstract_id=239684.

Novshek, W., & Sonnenschein, H. (1982). Fulfilled expectations Cournot duopoly with information acquisition and release. *Bell Journal of Economics*, *13*, 214–218.

Peters, M., & Severinov, S. (2006). Internet auctions with many traders. *Journal of Economic Theory*, *130*, 220–245.

Potters, J., & Suetens, S. (2006). Cooperation in experimental games of strategic complements and substitutes. CentER, Discussion Paper No. 2006-48.

Reeves, D. M., Wellman, M. P., MacKie-Mason, J. K., & Osepayshvili, A. (2005). Exploring bidding strategies for market-based scheduling. *Decision Support Systems*, *39*, 67–85.

Riveros, J. F. (1999). *Bundling Information Goods: Theory and Evidence*. Ph.D. thesis, University of Michigan, Department of Economics.

Rosenthal, R. W. (1973). A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, *2*, 65–67.

Rosenthal, R. W. (1993a). Bargaining rules of thumb. *Journal of Economic Behavior and Organization*, *22*, 15–24.

Rosenthal, R. W. (1993b). Rules of thumb in games. *Journal of Economic Behavior and Organization*, *22*, 1–13.

Salinger, M. A. (1995). A graphical analysis of bundling. *Journal of Business*, *68*, 85–98.

Salop, S. C. (1979). Monopolistic competition with outside goods. *Bell Journal of Economics*, *10*(1), 141–156.

Schmalensee, R. (1984). Gaussian demand and commodity bundling. *Journal of Business*, *57*(1), 211–230. Available at http://ideas.repec.org/a/ucp/jnlbus/v57y1984i1ps211-30.html.

Schuster, P., & Sigmund, K. (1983). Replicator dynamics. *Journal of Theoretical Biology*, *100*, 533–538.

Shapley, L. S. (1964). Some topics in two-person games. *Advances in Game Theory, Annals of Mathemacical Studies*, *5*, 1–28.

Simon, C., & Blume, L. (1994). *Mathematics for Economists*. W. W. Norton & Company, Inc.

Spence, A. M. (1976). Product selection, fixed costs, and monopolistic competition. *Review of Economic Studies*, *43*, 217–235.

StataCorp (2007). *Stata Longitudinal/Panel-Data Reference Manual*. Stata Press.

Taylor, P., & Jonker, L. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences*, *40*, 145–156.

Topkis, D. M. (1979). Equilibrium points in nonzero-sum n-person submodular games. *SIAM Journal of Control and Optimization*, *17*, 773–787.

Tversky, A., & Kahneman, D. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Varian, H. (1995). Pricing information goods. In *Scholarship in the New Information Environment Symposium*. Harvard Law School.

Varian, H. R. (1989). Price discrimination. In Schmalensee, R., & Willig, R. (Eds.), *Handbook of Industrial Organization*. North-Holland Press, Amsterdam.

Varian, H. R. (2000). Buying, sharing and renting information goods. *Journal of Industrial Economics*, *48*(4), 473–488. Available at http://ideas.repec.org/a/bla/jindec/v48y2000i4p473-88.html.

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, *16*, 8–37.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, *19*(3), 305–321. Available at http://dx.doi.org/10.1016/0304-4068(90)90005-T.

Vives, X. (2000). *Oligopoly Pricing: Old Ideas and New Tools*. MIT Press.

Walsh, W. E., Das, R., Tesauro, G., & Kephart, J. O. (2002). Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, Edmonton.

Weber, R. J. (1997). Making more from less: Strategic demand reduction in the FCC spectrum auctions. *Journal of Economics and Management Strategy*, *6*, 529–548.

Wellman, M. P. (2006). Methods for empirical game-theoretic analysis (extended abstract). In *Twenty-First National Conference on Artificial Intelligence*, pp. 1552–1555, Boston.

Wellman, M. P., Osepayshvili, A., MacKie-Mason, J. K., & Reeves, D. M. (2008). Bidding strategies for simultaneous ascending auctions. *The B.E. Journal of Theoretical Economics*, *8*(1). Article 27. Available at http://www.bepress.com/bejte/vol8/iss1/art27.

Wellman, M. P., Reeves, D. M., Lochner, K. M., & Vorobeychik, Y. (2004). Price prediction in a trading agent competition. *Journal of Artificial Intelligence Research*, *21*, 19–36.

Wellman, M. P., Walsh, W. E., Wurman, P. R., & MacKie-Mason, J. K. (2001). Auction protocols for decentralized scheduling. *Games and Economic Behavior*, *35*, 271–303.

Whinston, M. D. (1990). Tying, foreclosure, and exclusion. *American Economic Review*, *80*(4), 837–859. Available at http://ideas.repec.org/a/aea/aecrev/v80y1990i4p837-59.html.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.

Yee Liang, K., & Zeger, S. L. (1986). Longitudinal data analysis using generalized linear models. *Biometrika*, *73*(1), 13–22.

Zahray, W. P., & Sirbu, M. (1990). The provision of scholarly journals by libraries via electronic technologies: an economic analysis. *Information Economics and Policy*, *4*, 127–154.