

© 2016 by Rachel C. Shafer. All rights reserved.

ROBUSTNESS OF THE K -DOUBLE AUCTION UNDER KNIGHTIAN
UNCERTAINTY

BY

RACHEL C. SHAFER

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Economics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2016

Urbana, Illinois

Doctoral Committee:

Professor Steven R. Williams, Chair
Professor In-Koo Cho
Professor George Deltas
Professor Martin Perry

Abstract

This dissertation considers the robustness of private value and common value k -double auctions when those markets are populated by regret minimizers. Regret minimizing agents, unlike typical expected utility maximizers, need not commit to a single prior in their decision rule. In fact, it is a feature of the minimax regret decision rule that is not based on any prior. This makes the decision rule an interesting one for agents who face Knightian Uncertainty. A decision problem involves Knightian uncertainty if the agents know the possible outcomes but not those outcomes' probabilities – as may be the case in a new market.

This dissertation shows that in a private value k -double auction, minimax regret traders will not converge to price-taking behavior as the market grows. Similarly, in a common value auction, traders' behavior may depend on the parameter k , but does not depend on the number of other traders in the market.

The invariance of regret minimizing traders' strategies to the size of the markets they inhabit is not an accident of the sealed bid double auction institution. In fact, it is a consequence of the symmetry axiom. The final chapter in this dissertation shows that any agents in a k -double auction who use decision rules that accord with the symmetry axiom, then their bids and asks will not depend on the number of rival traders.

Soli Deo Gloria

Acknowledgments

Without counsel plans fail, but with many advisers they succeed. –Proverbs 15:22

This project exists thanks to the support of many people.

My advisor, Steve Williams, has been the gentlest and most faithful advisor that anyone could ask for. I am grateful for our frequent meetings and his thoughtful advice. Each professor on my committee offered guidance and help in keeping with his own gifts. In-Koo Cho's high standards for economic research spurred me on towards precise language and careful reasoning. George Deltas, a reliable and encouraging presence from the very start of my time at the University of Illinois, made comments that expanded my vision for my research. Marty Perry administered sage career advice and created exciting teaching opportunities for me. Their generosity made me a better scholar, teacher, and professional. The shortcomings that remain in this dissertation, of course, are my own responsibility.

Thanks to the University of Illinois Economics Department for awarding me a 2014 Summer Research Fellowship, and to Paul W. Boltz for his generous funding of the summer fellowship that supported my research in Summer 2015. These were helpful boosts at the end of a winding road.

I was also sustained by the comradeship of the men and women who traveled that road alongside me. At the start of the economics doctoral program, I learned with a stellar study group: Dongwoo Kim, Jiaying Gu, Deniz Ay, and Blake Riley. More recently, I benefited from exchanging ideas with the members of UIUC's mechanism design discussion group, particularly Juan Fung, Blake Riley, Tom Sahajdack, Kazuyuki Hashimoto, David Quigley, Kwanghyun Kim, and Chia-Ling Hsu.

In truth, my path to writing this dissertation began years before I came to Urbana-Champaign. I am grateful to the University of Tulsa's economics faculty (including Bobby Horn, Scott Carter, Greg Burn, Russel Evans, and Megumi Nakao) for inspiring me both to study economics, and to desire to teach it. Thanks especially to Steve Steib and Chad Settle, who advised my senior project (my first foray into the wonderful world of auctions!) and encouraged my pursuit of graduate studies. Their confidence in me, when I was not confident, was a great gift. I hope that I can help others in the same way that Dr. Steib and Dr. Settle helped me.

My parents, Steven Shafer and Roberta Garcia Shafer, cultivated my love of learning and equipped me for all of my future achievements. Together with my siblings, Andrew and Christiana, they created a home environment of spirited inquiry, cheerful generosity, and principled excellence. My family's limitless love is amazing and humbling.

Finally, thanks to my church (All Souls Presbyterian Church), and to my friends (too numerous to mention all by name!) for their love, prayers, and support.

Table of Contents

List of Figures	viii
Chapter 1 Introduction	1
1.1 Motivation and Background	1
1.2 The Sealed-Bid Double Auction	2
1.3 Decisions Under Knightian Uncertainty	4
1.3.1 Notation for Decision Problems Under Knightian Uncertainty	4
1.3.2 Contrasting Incomplete Information and Knightian Uncertainty	5
1.4 The Double Auction Under Knightian Uncertainty	8
1.5 Summary and Intuition of Findings	9
Chapter 2 Convergence to Price-Taking in the Private Value k- Double Auction	13
2.1 Introduction	14
2.2 Traders with Private Values	18
2.3 First approach: minimizing maximum regret	21
2.3.1 Minimax Regret defined	21
2.3.2 Minimax Regret in a k-Double Auction	22
2.3.3 Large Markets and Efficiency	23
2.4 Second Approach: minimizing expected maximum regret	25
2.4.1 Minimizing expected maximum regret in a k-Double Auction .	26
2.4.2 Minimizing Expected Maximum Regret in Large Markets . . .	28
2.5 Third Approach: minimizing maximum expected regret	30
2.5.1 Minimizing Maximum Expected Regret defined	30
2.5.2 Sufficient Conditions for Convergence to Truthful Bidding by Maximum Expected Regret Minimizers	31
2.6 Conclusion	33

Chapter 3	Information Aggregation in the Common Value k-Double Auction	35
3.1	Introduction	36
3.2	Regret Minimizers in Common Value Sealed-Bid Auctions	38
3.2.1	The Model	38
3.2.2	Vickrey Auctions: Truthful Bidding, Winner's Curse	39
3.2.3	First Price Auction: Pessimistic Bidding, Reduced Winner's Curse	39
3.3	Regret Minimizers in Common Value k -Double Auctions	42
3.3.1	The Model	42
3.3.2	Minimax Regret Bid	43
3.3.3	Minimax Regret Ask	43
3.3.4	Mutual Minimax Regret Bids and Asks for Some Values of k	43
3.4	Conclusion	44
Chapter 4	The Symmetry Axiom and Strategies Invariant to the Number of Players	45
4.1	Introduction	46
4.1.1	Motivation and Background	46
4.1.2	Related Work	48
4.2	Decision Problems Under Knightian Uncertainty	50
4.2.1	Three Decision Rules	51
4.3	Axioms for Decision Problems Under Knightian Uncertainty	55
4.3.1	Four Key Axioms	55
4.3.2	Further Discussion of Symmetry	61
4.3.3	Additional Axioms	64
4.4	Axiomatic Characterization of Decision Rules	65
4.5	Consequences of the Symmetry Axiom in k -Double Auctions	66
4.5.1	In a k -Double Auction, the Symmetry Axiom Results in Bid Invariance to Market Size	67
4.5.2	Discussion	69
4.6	Conclusion	70
Appendix A	Proofs for Chapter 2	72
Appendix B	Proofs for Chapter 3	85
Appendix C	Proofs for Chapter 4	96
References		99

List of Figures

1.1	Bidder Decision in Private-Value k -Double Auction: Game of Incomplete Information	5
1.2	Bidder Decision in Private-Value k -Double Auction: Decision Problem Under Knightian Uncertainty (Bayesian Approach)	7
1.3	Bidder Decision in Private-Value k -Double Auction: Decision Problem Under Knightian Uncertainty (Minimax Regret)	8
2.1	Bidder's Profit from bid b	20
2.2	Sellers Profit from ask a	20
2.3	Bidder's Regret given bid b and valuation v	22
2.4	The Distribution of Bids and Asks Depends on k	24
2.5	Bids and asks for traders that minimize expected maximum regret, at various market sizes	29
4.1	A Decision Problem Example	51
4.2	A Transformation of Decision Problem Example	60

Chapter 1

Introduction

1.1 Motivation and Background

An important function of markets is that they convey information through prices. In markets where each buyer has his or her own valuation for consuming the good and each seller has his or her own cost for producing it (that is, a market with “private valuations”), this aggregation of information is important because it allows the market to realize an efficient allocation of resources (Hayek, 1945). In markets where each trader attempts to estimate the unknown value of an asset (that is, a market with “common valuations”), the information aggregation is important because it reveals the asset’s true value.

Sealed-bid double auctions are models from which we can gain compelling insights into the workings of markets, and their potential to aggregate information. It is well known that no bilateral trading mechanism is efficient without outside subsidies if it is incentive compatible and individually rational (Myerson and Satterthwaite, 1983). However, the private-value k -double auction converges to efficiency quickly (Rustichini, Satterthwaite and Williams, 1994). Reny and Perry (2006) also use a sealed-bid double auction to make a step toward a strategic foundation for the

rational expectations equilibrium.

The expected utility maximizers that populate most double auction models possess and use a great deal of information. They know the number of traders in the market, and the distribution of other traders' redemption values, and they coordinate their strategies with those of other traders to reach an equilibrium. In some markets, particularly new markets, it is unlikely that traders will have precise beliefs about all of these things (Wilson, 1987; Harstad, Pekec and Tsetlin, 2008). This dissertation considers the robustness of the k -double auction when traders face more severe forms of uncertainty than the uncertainty generally accommodated in typical models.

1.2 The Sealed-Bid Double Auction

Agents submit sealed bids (denoted b_i) and asks (a^j) to the auctioneer. We assume that the submitted bids (b^1, b^2, \dots, b^m) are positive real numbers that cannot exceed the maximum valuation \bar{v} . We assume that sellers submit asks (a^1, a^2, \dots, a^n) that cannot be less than the minimum cost \underline{c} . For simplicity of notation, we will assume that the range of acceptable bids and the range of acceptable asks is $Z = [\underline{c}, \bar{v}]$, where $\underline{c} < \bar{v}$.

These bids and asks determine a single price at which all units are traded and identify which buyers and sellers will trade. The price p is set within the interval $[x, y]$ of prices such that the number of buyers whose bids exceed the price equals the number of sellers whose ask is less than the price. The exact price selected within

this interval of market-clearing prices depends on the exogenous parameter $k \in [0, 1]$:

$$p = (1 - k)x + ky. \tag{1.1}$$

Example 1 *Suppose that there are $m = 3$ buyers who submit bids ($b_1 = 4.50, b_2 = 2.12, b_3 = 7.00$) and $n = 4$ sellers who submit asks ($a_1 = 1.00, a_2 = 3.45, a_3 = 10.30, a_4 = 5.87$). Then the market clears at any price between 3.45 and 4.50. If $k = .8$, then the price set by the k -double auction is 4.29. Buyers 1 and 3 will trade with sellers 1 and 2. (Since the units of the good are identical, and all traded units are traded at the same price, it is irrelevant which buyer trades with which seller.)*

It is useful to note that the price in the market depends on the m^{th} and $m + 1^{\text{st}}$ -lowest bid or ask. Let $(z_{(1)}, z_{(2)}, \dots, z_{(m+n)}) \in Z^{m+n}$ be the ordered set of bids and asks, with $z_{(1)} < z_{(2)} < \dots < z_{(m+n)}$. In our example above, $z = (z_{(1)} = 1.00, z_{(2)} = 2.12, z_{(3)} = 3.45, z_{(4)} = 4.50, z_{(5)} = 5.87, z_{(6)} = 7.00, z_{(7)} = 10.30)$. The two values that determined the price were $z_{(3)}$ and $z_{(4)}$. In fact, the interval of market-clearing prices is always $[z_{(m)}, z_{(m+1)}]$.

In the case that $z_{(m)} = z_{(m+1)}$, the interval of market clearing prices consists of a single price, $p = z_{(m)}$. The number of sellers asking less than this price may not equal the number of buyers bidding above this price. In that case, a fair lottery may determine which of the traders on the long side of the market with bids (or asks) equal to p will be allowed to trade (Satterthwaite and Williams, 1993).

1.3 Decisions Under Knightian Uncertainty

In analyzing the market described in section 1.2, it is typical to think of the auction as a game of incomplete information (Harsanyi, 1967) and to seek a solution in the form of a Bayesian Nash equilibrium. In contrast, this paper treats the trader's situation as a decision problem under Knightian uncertainty (Knight, 1912). This section provides a notational framework for the decision rules that the rest of the paper will examine, and discusses how this paper's approach to Knightian uncertainty compares to the typical approach.

1.3.1 Notation for Decision Problems Under Knightian Uncertainty

Decision problems involve a set of acts \mathcal{A} available to the decision-maker, the set of possible states of the world \mathcal{S} , and the outcome $u \in \mathcal{U}$ that results in the state $s \in \mathcal{S}$ given the decision-maker's action $a \in \mathcal{A}$. We may find it useful to think of the decision-maker as having a payoff function $u : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$.

A decision rule specifies what a decision maker will do given a menu \mathcal{A} of possible actions. The following chapters will introduce a variety of decision rules, and later characterize some of them using axioms. Those axioms apply to decision makers' preferences \succsim . Let $\succsim_{\mathcal{A}}$ denote a preference relation over the actions available in the menu \mathcal{A} . A preference relation is defined to be a binary relation \succsim that is reflexive ($a \succsim a$ for all actions a) and transitive (if $a_1 \succsim a_2$ and $a_2 \succsim a_3$, then $a_1 \succsim a_3$) (Fishburn, 1970). From \succsim we can derive relations \succ and \sim in the usual way.

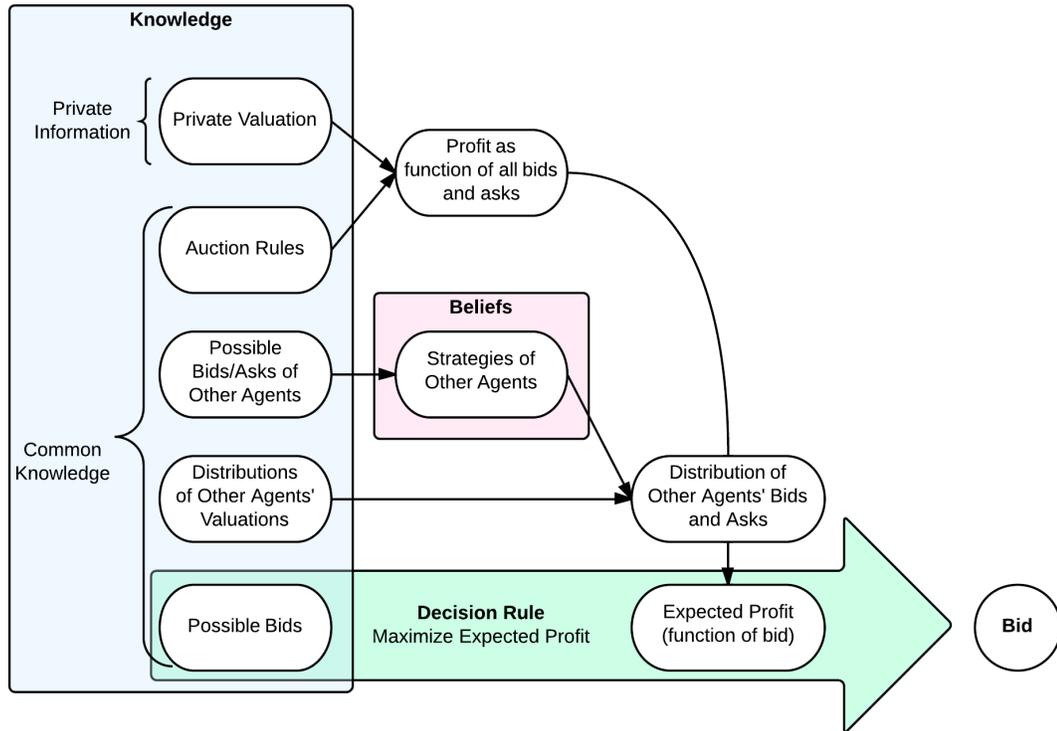


Figure 1.1: Bidder Decision in Private-Value k -Double Auction: Game of Incomplete Information

1.3.2 Contrasting Incomplete Information and Knightian Uncertainty

When the double auction is treated as an incomplete information game, each trader does not know the types of other traders. That is, in a private value auction they do not know the private redemption values of other traders. In a common value auction, they do not know the signals that other traders have observed about the asset's true value. However, the bidder does know the distributions from which the private redemption values or signals are drawn. These distributions, together with some

equilibrium belief about the traders' strategies, is the foundation for the bidder's beliefs about the distribution of others' bids and asks. Using this distribution of others' bids and asks, the bidder can calculate the expected profit of each of his own possible bids. An expected utility maximizer chooses a bid that yields the greatest expected profit. Figure 1.1 illustrates this approach to the private value double auction.

In contrast, this paper treats the agent's situation as a decision problem under ambiguity, which is also known as Knightian uncertainty after Knight (1912). In this approach, the set of possible states of the world (and the outcome in each state) is known to the decision-maker, but the probability of each state of the world is not. In our model, the traders know the range of possible types for the other traders. They know their own redemption type with certainty. However, they do not know the *distribution* of other traders' types.

Two possible approaches to this decision problem for a private value double auction under Knightian uncertainty are depicted in Figures 1.2 and 1.3. The crucial difference between these figures and Figure 1.1 is that knowledge about "Distributions of Other Agents' Valuations" has been removed. Instead, traders know the range of possible valuations of other traders – the support of the distribution, rather than the distribution itself.

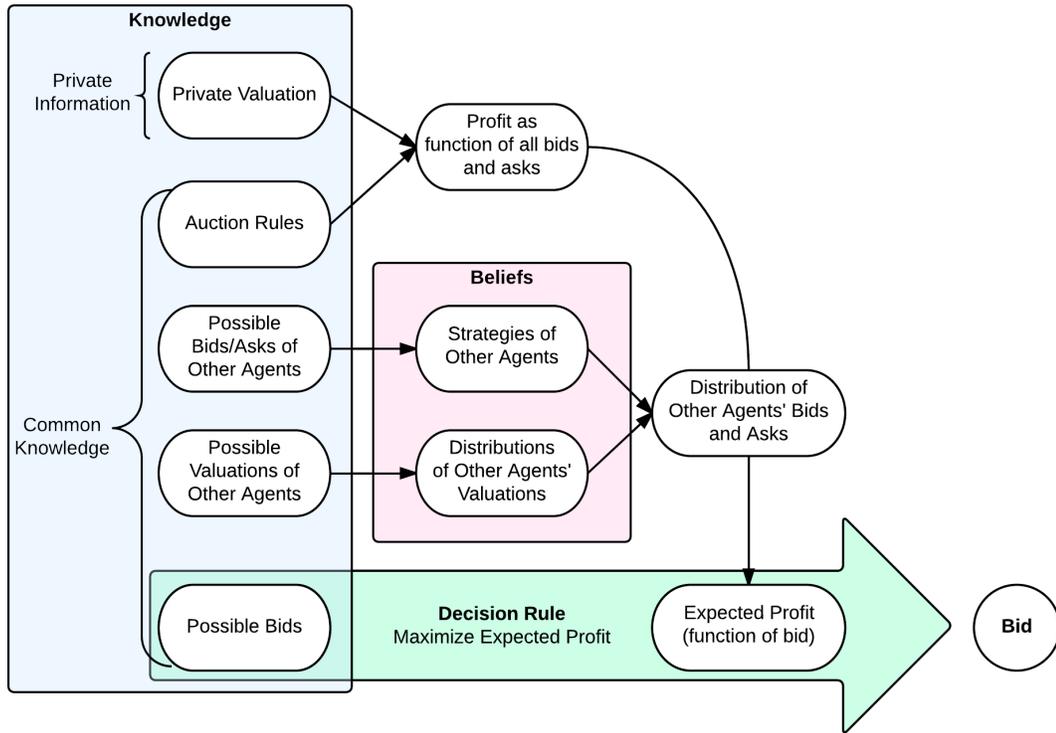


Figure 1.2: Bidder Decision in Private-Value k -Double Auction: Decision Problem Under Knightian Uncertainty (Bayesian Approach)

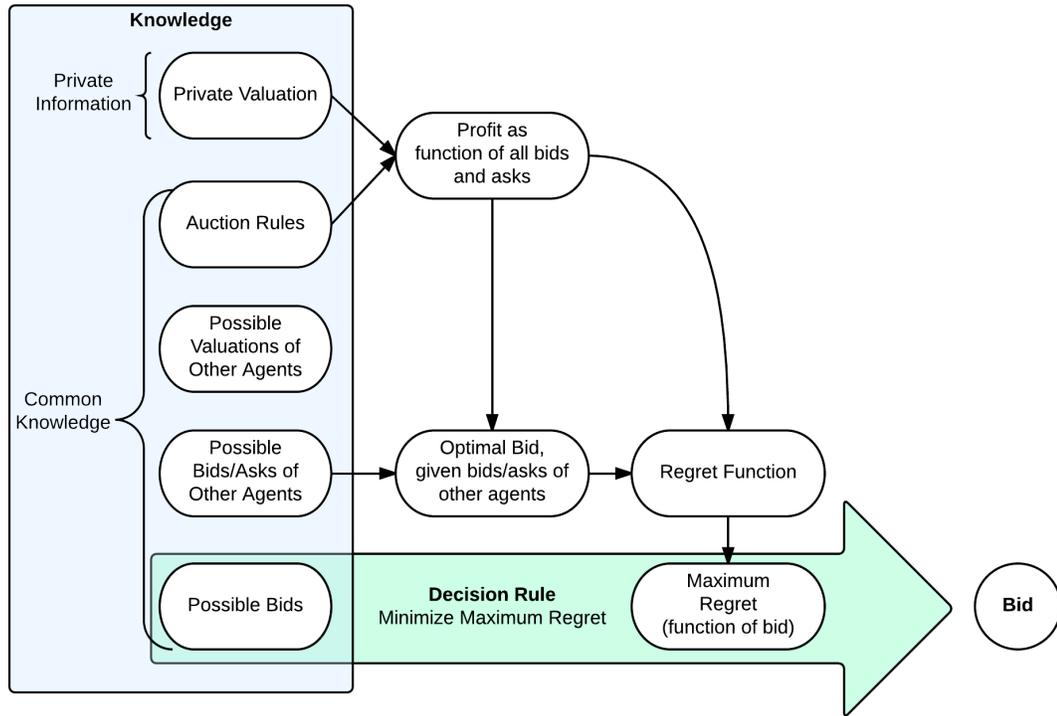


Figure 1.3: Bidder Decision in Private-Value k -Double Auction: Decision Problem Under Knightian Uncertainty (Minimax Regret)

1.4 The Double Auction Under Knightian Uncertainty

Approaching the double auction as a decision problem under Knightian uncertainty is a more general approach than studying the situation as a game. As pictured in Figure 1.2, the agent *could* choose to resolve the decision problem by selecting an equilibrium strategy given his subjective prior about the distribution of redemption values and the strategies of other traders. But it is also possible to resolve the

decision problem using another decision rule.

One example of a decision rule that does not calculate expected profit, or use a prior at all, is minimax regret. As pictured in Figure 1.3, a minimax regret bidder in a k -double auction will choose his bid by analyzing the regret function associated with each bid. No prior is used; all that is needed to make a decision is to know the set of possible outcomes.

Minimax regret is a decision rule suggested by Savage (1951) as an alternative to maxmin. Since then, it has been applied to a number of market models. Renou and Schlag (2008) have applied minimax regret to price-setting environments, for example. Hayashi and Yoshimoto (2012) have created and calibrated a risk- and regret- averse model for bidders in first price auctions. Taking another tactic, Chiesa, Micali and Zhu (2014) have studied Knightian “Self Uncertainty” in combinatorial auctions. Linhart and Radner (1989) applied minimax regret to bargaining. The work in this dissertation aligns most closely with Linhart and Radner’s approach, but it extends their results from bilateral bargaining, with only one buyer and one seller, to the sealed-bid double auction, with many buyers and many sellers. It also considers common value auctions in addition to the private valuations of Linhart and Radner (1989).

1.5 Summary and Intuition of Findings

The following chapters examine traders in auctions who face Knightian uncertainty. Private value and common value settings are considered in turn.

In Chapter 2, each trader has a private redemption value (valuation or cost). When such traders minimize maximum regret, they choose bids and asks that differ from their true redemption values, leading to inefficient outcomes for any sized market. Moreover, restricting Knightian uncertainty to the traders' beliefs about one another's strategies, while allowing them to hold beliefs about the distribution of private redemption values, does nothing to improve market efficiency. A third and final attempt at inducing regret minimizers to converge to price-taking behavior in large markets is successful, but only by endowing traders with a *set* of multiple priors. These results reinforce the power of some priors to prevent traders from realizing gains from trade.

In Chapter 3, the assets being traded have the same value to all traders, and each trader observes private signals about that true value. In the common value auction, potential regret from discovering that the asset's true value is lower than one's signal motivates minimax regret traders to submit cautious bids and asks; however, whether the price can converge to the item's true value as the market grows depends on the choice of the auction rule k and on the distribution of signals.

A common thread runs through these results. The traders' behavior does not depend on the size of the market. This is markedly different from the behavior of expected utility maximizers, whose bids and asks depend on beliefs about the number of other traders.

We can gain some intuition for this difference between regret minimizers and expected utility maximizers with a simple analogy. Think of the bids and asks of other traders as strands in a net. If there are only a few traders in the market, there

are only a few strands in the net. Such a net will have many gaps through which a trader could “fall” – that is, the trader’s bid or ask could easily turn out to be “pivotal”, one of those bids and asks that determines the market clearing price p . Then it is possible that the trader fails to maximize his profit (buyers will regret not bidding less, and sellers will regret not asking more). Whether a trader is a regret minimizer or an expected utility maximizer, the prospect of falling through a gap in the “net” of rival bids and asks influences that trader to attempt to move the price in his own favor.

Now consider what happens as more and more strands are added to the net. Someone with an expectation that new strands, when they are added, could be placed in any part of the net, would reason that as the strands in the net increase, the likelihood of falling through a gap will decrease. The gaps will become narrower and narrower, until no sizable gaps remain – so long as each space in the net has some chance of being bridged by an additional strand. In the same way, a trader who expects that his rivals’ bids and asks are distributed with full support will conclude that for very large markets, his probability of being a pivotal trader is negligible.¹ Since that is the case, the expected utility maximizer should put less and less weight on the advantages of influencing the market price, as more and more traders are added to the market.

Here is the vital difference in minimax regret. A minimax regret trader does not expect that the gaps will be filled as strand after strand is added to the net.

¹This is not a necessary condition for price-taking to be in the set of profit-maximizing actions in a private value auction, but the condition supports our intuition for why large markets induce price-taking behavior.

He cannot have such an expectation, because he has eschewed expectations entirely. Instead, the minimax regret's choice is always influenced by the possibility that a gaping hole remains in the net – that is, that he will turn out to be the pivotal trader, and therefore regret any failure to influence the price in his favor.

To explain the invariance of regret minimizing traders more formally, Chapter 4 brings attention to the axiomatic characterization of the minimax regret decision rule (as described in Stoye (2011*b*)). The invariance of minimax regret traders' strategies to the size of the market is not a special circumstance of the double auction. Rather, it stems from the decision rule's adherence to the symmetry axiom. Chapter 4 demonstrates that any decision rule that adheres to the symmetry axiom will also result in strategies that do not vary with the number of traders in the double auction.

Taken together, these results suggest some drawbacks to using minimax regret as a decision rule in models of large markets. They reveal potential problems with relying on large markets to achieve efficient allocations (for private value settings) or informative prices (for common value settings). More generally, they illuminate the importance of individuals' beliefs about the markets in which they trade.

Chapter 2

Convergence to Price-Taking in the Private Value k -Double Auction

Abstract

This paper studies a variety of forms of regret minimization as the criteria with which traders choose their bids/asks in a double auction. Unlike the expected utility maximizers that populate typical market models, these traders do not determine their actions using a single prior. The analysis proves that minimax regret traders will not converge to price-taking as the number of traders in the market increases, contrary to standard economic intuition. In fact, minimax regret traders' bids and asks are invariant to the number of other traders in the market. However, not all regret-based decision rules fail to respond to market size. Introducing priors over some part of the decision problem to minimize expected maximum regret, or multiple priors to minimize maximum expected regret, have different effects. The robustness of the sealed bid double auction is limited by the need to avoid priors that eliminate traders incentive to truthfully reveal their redemption values.

Keywords: double auctions; regret minimization; Knightian uncertainty; decision theory; mechanism design

JEL Classification Numbers: D44, D81, D82, C72

2.1 Introduction

Perfectly competitive markets are efficient only if traders act as price takers¹, behavior that can be induced in large markets if traders recognize that the size of the market attenuates each individual trader's influence. The double auction models that formally prove this familiar reasoning typically attribute a great deal of knowledge to their traders. Traders are assumed to be capable of coordinating on an equilibrium in which each trader maximizes expected utility, something that is only possible because they know the distribution of traders' bids and asks. But are these strong assumptions on traders' knowledge and capabilities necessary, or could traders who do not know the distribution of bids and asks still converge to price-taking behavior as the market grows? Our confidence in a market's robustness may depend on the answer.

This chapter replaces the expected utility maximizers that populate a conventional double auction model with regret minimizing traders. "Regret" here is the difference between one's actual payoff (a function of one's action and the realized state of the world), and the best possible payoff that could have been achieved in the realized state (Savage, 1951; Linhart and Radner, 1989). Regret minimization can be defined in a variety of ways, and this chapter examines three separate versions of

¹By "price takers," I mean that each buyer and seller truthfully reports their utility-maximizing quantity to produce or consume at a given price, rather than attempting to manipulate prices.

regret minimization. What all three versions of the regret minimizing trader have in common is that none of them determines his action by referring to a specific belief (a prior) about the distribution of other traders' bids and asks.

Because this chapter's regret minimizers do not rely on a particular prior, they are equipped to handle a type of uncertainty that conventional models do not address: Knightian uncertainty (Knight, 1912). Under Knightian uncertainty, the set of possible states of the world (and the outcome in each state) is known to the decision-maker, but the probability of each state of the world is not. For example, if a person does not trust that a coin being tossed is a fair coin, then that person faces Knightian uncertainty: the possible states of the world are known to be Heads and Tails, but the probability of each state is unknown. In this chapter, traders face Knightian uncertainty regarding other traders' strategies, and perhaps also the distribution of other traders' underlying redemption values.

An expected utility maximizer's response to Knightian uncertainty is to adopt a subjective prior, this approach can be problematic, leading us to seek an alternative approach. Each of the three problems discussed below relates to a specific version of regret minimization and a separate result in this chapter. Taken together, the chapter's three results reveals how some priors can prevent convergence to price-taking.

The first problem with relying on a single subjective prior applies when the decision maker is very unfamiliar with the decision problem. Complete ignorance cannot be adequately reflected by any prior, even a uniform prior that treats each possible outcome as equally likely, because even adopting a uniform prior asserts some

knowledge about the specification of the decision problem. If the decision maker's ignorance is so complete that he does not know which characteristics of events are relevant and which are extraneous, then his decision rule ought not to depend on the way he has chosen to specify the problem (Arrow and Hurwicz, 1972). *Minimax regret*, the first version of regret minimization that this chapter considers, is well-suited to situations of complete ignorance because it essentially accommodates all priors at once. This chapter finds that minimax regret traders do not converge to price-taking behavior.

The second problem with a single subjective prior extends to cases other than complete ignorance. Even supposing that the trader does have a sense of the distribution of the other buyers' and sellers' redemption values, the trader may face Knightian uncertainty regarding those traders' strategies. The multiplicity of Bayes Nash equilibria in a double auction makes this concern especially acute. A trader that allows for the full range of rationalizable strategies on the part of his rivals to calculate maximum regret, but then applies a prior over the rivals' valuations and costs, is said to be *minimizing expected maximum regret*. Linhart and Radner (1989) have examined this decision rule in the case of bilateral trade; the present chapter extends their analysis to larger markets, and finds that the decision rule does not induce convergence to price-taking. On the contrary, such a bidder will shade his bid more, not less, as the size of the market increases, approaching the minimax regret bid.

The third problem with a single subjective prior is that real decision makers are not always willing to commit to a single prior, even when they have a basis to do

so. De Finetti (as quoted by Dempster (1975)) explained that in many situations a decision-maker's subjective prior will only be "vaguely acceptable". Therefore "it is important not only to know the exact answer for an exactly specified initial position, but what happens changing in a reasonable neighborhood the assumed initial opinion." This justifies the use of decision rules that involve multiple priors. A well-known example of such a decision rule is maxmin expected utility with a non-unique prior (Gilboa and Schmeidler, 1989). Similarly, a decision maker can use multiple priors to *minimize maximum expected regret*. The third and final result of this chapter is that a trader who minimizes maximum expected regret may converge to price-taking behavior as the market grows – even though such a trader may not evaluate the possible bids according to a single prior, as an expected utility maximizer would. However, the set of priors must satisfy certain conditions in order for minimax expected regret traders to converge to price-taking.

Taken together, the chapter's three results indicate the significance that individual beliefs may have for the efficiency of the entire market. Traders whose decision rule is consistent with every prior fail to converge to price-taking (Theorem 1). The failure to converge suggests that restricting the priors is key to inducing price-taking behavior. But it is not enough to restrict only one aspect of the decision problem: introducing prior beliefs about redemption values but not strategies does not ensure convergence to price-taking (Theorem 2). Still, traders can converge to price-taking as long as the set of priors they consult is restricted from the priors that would prevent convergence for even expected utility maximizers (Theorem 3). Whether bidders are expected utility maximizers or regret minimizers, eliminating "bad pri-

ors” is essential for markets to function efficiently.

The following sections begin with an explanation of the traders, double auction rules, and profit functions of the model (section 2.2)). The next three sections provide formal definition of the decision rule, and analysis of the resulting outcome in the double auction, for each of the three versions of regret minimization. Section 2.6 concludes.

2.2 Traders with Private Values

There are m buyers and n sellers. Each buyer i has a valuation $v^i \in [\underline{v}, \bar{v}] \subset \mathbb{R}_+$, which is the buyer’s maximum willingness to pay for a single unit of the good. Each seller j has a cost $c^j \in [\underline{c}, \bar{c}] \subset \mathbb{R}_+$ of producing a single unit of the good. We will refer to both valuations and costs as the traders’ redemption values. Agents do not supply or demand more than one unit of the homogeneous and indivisible good.

Place these traders with private valuations into the sealed bid double auction described in section 1.2. The relationship between the trader’s profit and the outcome of the auction is straightforward. Buyer i ’s profit is $v^i - p$ if he trades and zero otherwise. Seller j ’s profit is $p - c^j$ if he trades and zero otherwise. Thus, given a trader’s redemption value, a bid or ask determines a set of possible payoffs, the realization of which depends on the bids and asks submitted in the double auction by *other* traders. Let $\zeta = (\zeta_{(1)}, \zeta_{(2)}, \dots, \zeta_{(m+n-1)})$ denote the ordered set of the bids and asks submitted by those other traders in the auction. Then if a buyer with valuation

v submits bid b in the auction, his corresponding profit function Π_B is

$$\Pi_B(b, v, \zeta) = \begin{cases} v - [(1 - k)\zeta_{(m)} + k\zeta_{(m+1)}] & \text{if } \zeta_{(m+1)} < b \\ v - [(1 - k)\zeta_{(m)} + kb] & \text{if } \zeta_{(m)} < b < \zeta_{(m+1)} \\ 0 & \text{if } b < \zeta_{(m)} \end{cases} \quad (2.1)$$

The profit function's relationship to the realization of the bids/asks $(\zeta_{(m)}, \zeta_{(m+1)})$ is pictured in figure 2.1. Since $\zeta_{(m)} < \zeta_{(m+1)}$ always, the "state space" of possible outcomes is the triangle above the 45 degree line. Two things are clear from the figure. First, the bidder trades only if he bids more than $\zeta_{(m)}$. Second, the bidder is "pivotal" (influences the market price) only if his bid b lies between $\zeta_{(m)}$ and $\zeta_{(m+1)}$.

Likewise, if a seller with cost c submits ask a in the auction, his corresponding profit function Π_A is

$$\Pi_A(a, c, \zeta) = \begin{cases} [(1 - k)\zeta_{(m)} + k\zeta_{(m+1)}] - c & \text{if } a < \zeta_{(m)} \\ [(1 - k)a + k\zeta_{(m+1)}] - c & \text{if } \zeta_{(m)} < a < \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < a \end{cases} \quad (2.2)$$

See figure 2.2.

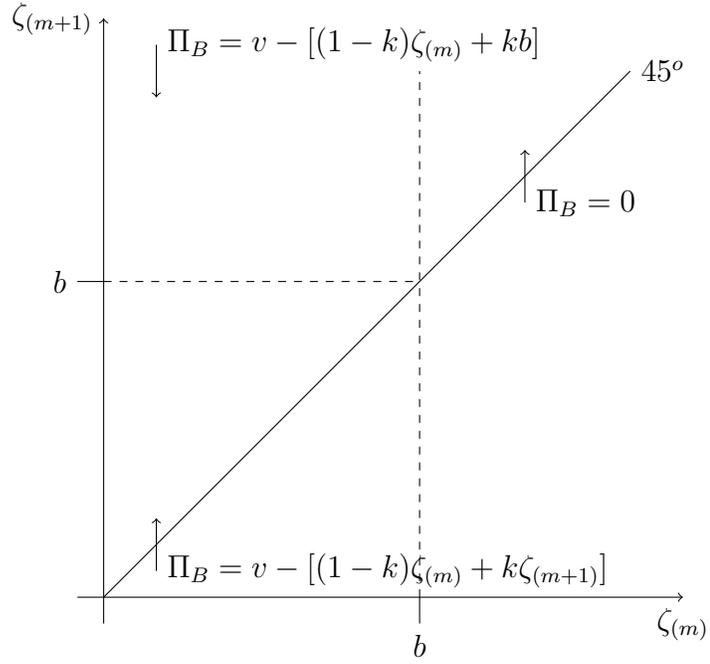


Figure 2.1: Bidder's Profit from bid b

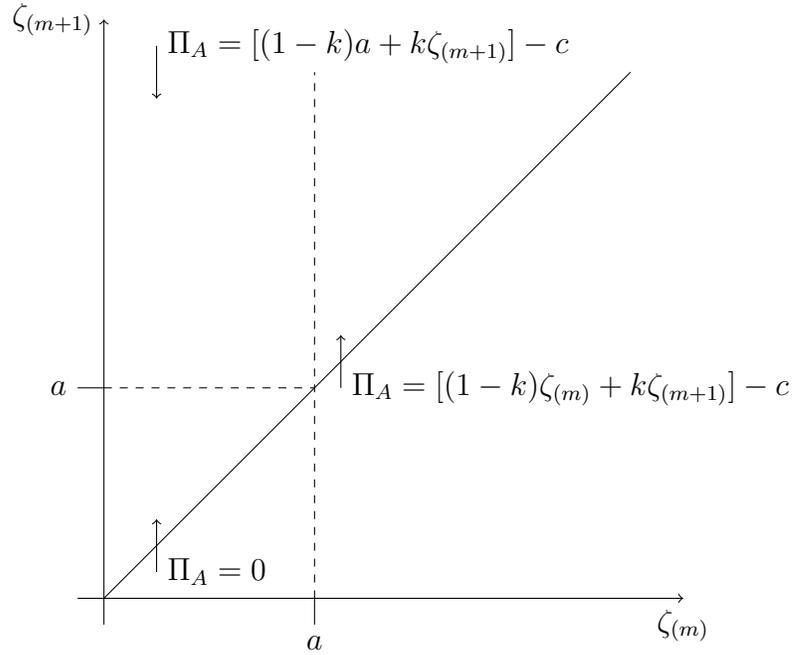


Figure 2.2: Seller's Profit from ask a

2.3 First approach: minimizing maximum regret

This section formally defines minimax regret, and shows that minimax regret traders generally do not report their true redemption values. Furthermore, they do not converge to truthful bidding no matter how large the market grows.

2.3.1 Minimax Regret defined

The action(s) minimizing maximum regret are identified by calculating the maximum regret that could be incurred under each action. The regret for a particular action in a particular state is calculated by comparing that action's payoff to the maximum possible payoff in the same state.

Definition 2 *An action a attains minimax regret if*

$$a \in \arg \min_{a \in \mathcal{A}} \max_{s \in \mathcal{S}} \{ \max_{a^* \in \mathcal{A}} u(a^*, s) - u(a, s) \} \quad (2.3)$$

From the standpoint of a person accustomed to working with expected payoffs, it may seem that minimax regret operates by choosing a “pessimistic” prior – a prior that assigns higher probability to events with very low or very high payoffs. The truth is subtly different. The decision rule does not stick to a single pessimistic prior by which each action is evaluated. Instead, a minimax regret trader evaluates each action by focusing exclusively on the state in which regret is highest *for that action*. Of course, this is equivalent to using a prior that assigns probability 1 to the event that corresponds to this extreme outcome. However, the prior that is used

to evaluate action a_1 may be very different from the prior that is used to evaluate action a_2 .

2.3.2 Minimax Regret in a k-Double Auction

The figure below shows a bidder's regret if his private valuation is v and he chooses to submit bid b . Note that the bidder's regret decreases as the rival bid $\zeta_{(m)}$ approaches his bid (since there is less regret from overbidding in that case) and again as it approaches his own valuation.

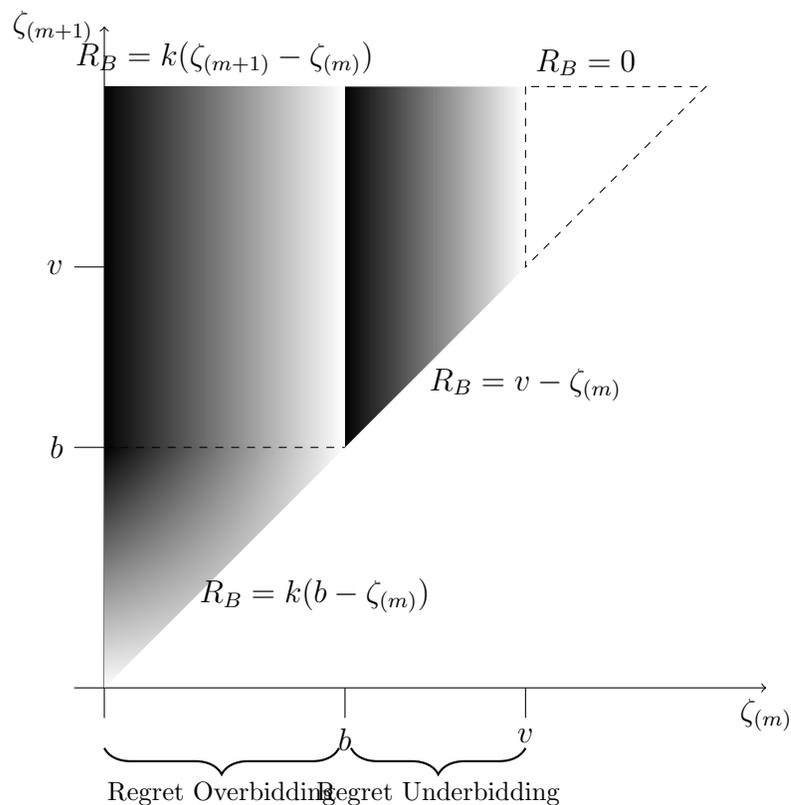


Figure 2.3: Bidder's Regret given bid b and valuation v

Theorem 1 *In a k -double auction, the bid b_i that minimizes maximum regret for a buyer with private valuation v_i is $b_i = \frac{v_i}{1+k}$. The ask a_i that minimizes maximum regret for a seller with private cost c_i is $a_i = \frac{c_i+(1-k)}{1+(1-k)}$*

The closer k is to 0, the less influence the buyer's bid has on the price, and consequently the closer the buyer minimax regret strategy will be to truthful revelation of his value. The closer k is to 1, the greater the potential influence of the buyer's bid on the price, and consequently the further the buyer minimax regret strategy will be from truthful revelation of his value. It is the opposite for a seller.

For $k = \frac{1}{2}$, these results are identical to the minimax regret strategies found by Linhart and Radner (1989) using their first approach to bilateral bargaining under incomplete information.

2.3.3 Large Markets and Efficiency

The minimax regret bids do not depend on the number of rivals. No matter how many buyers and sellers participate in the auction, a trader minimizing maximum regret under this approach will submit the same bid or ask. The strategies are also unaffected by the number of buyers relative to the number of sellers.

Figure 2.4 illustrates how this will affect the expected number of trades, the price, and the gains from trade when redemption values are uniformly distributed over $[0, 1]$ as $n, m \rightarrow \infty$. The thin lines represent the true demand and supply in the market. The thick lines show the demand and supply curves that result from aggregating the minimax regret bids and asks.

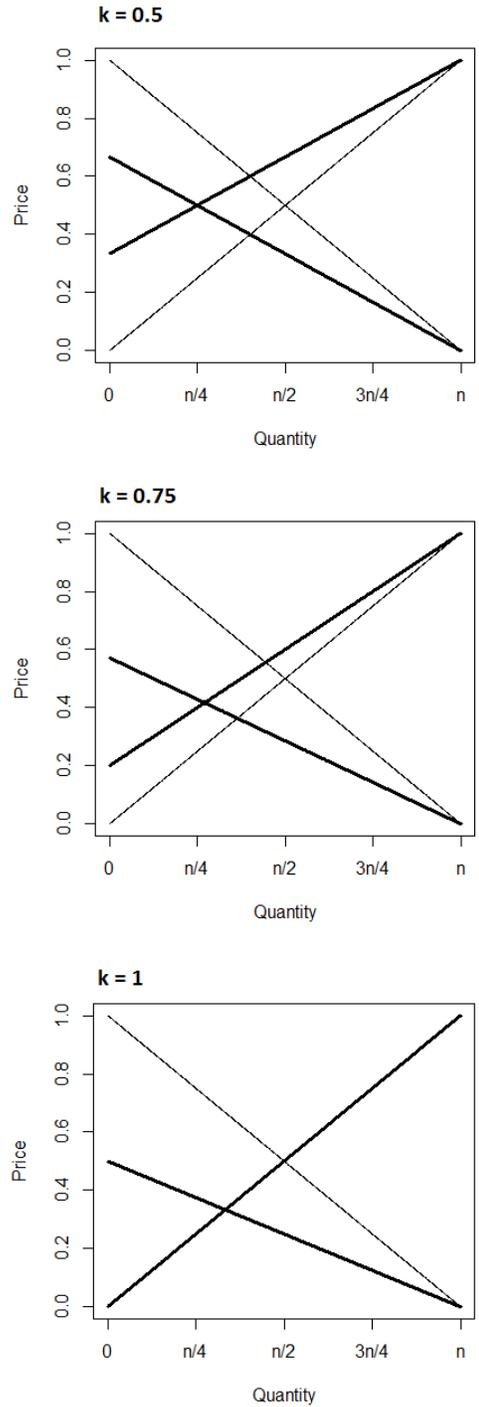


Figure 2.4: The Distribution of Bids and Asks Depends on k

Depending on k , one side of the market or the other may misrepresent their redemption values more. But whatever the value of k , the demand and supply curves meet at a quantity smaller than the quantity where the true valuations meet the true costs. Furthermore, the price may differ from the efficient price; it will favor the side of the market that has greater influence on the price.

Since the buyers and sellers do not report their true valuations/costs, some opportunities for profitable trade will be missed. Since the strategies do not converge to price-taking as the size of the market increases, the outcome will not approach efficiency either.

2.4 Second Approach: minimizing expected maximum regret

In this section, I find the bid and ask functions for traders that minimize expected maximum regret. As the name of their decision rule implies, these traders apply a prior to some part of their decision problem, unlike the minimax regret traders in the previous section. Constraining priors in part – but not all – of the decision problem clarifies the relationship between beliefs and outcomes. Although strategies are not invariant to market size, traders still do not converge to price-taking behavior, indicating the regret minimization can be troublesome if priors are unconstrained in any part of the decision problem.

2.4.1 Minimizing expected maximum regret in a k-Double Auction

This decision rule supposes that traders have some information about the trading environment, but do not know how other traders will choose to respond to that environment. Suppose that each trader knows the distribution of other sellers' costs and other bidders' valuations, unlike a minimax regret bidder. However, each trader remains in a state of Knightian uncertainty regarding the traders' strategies. Any rationalizable² strategy is considered plausible, and the trader does not wish to distinguish between probable and improbable strategies, nor to assume that all traders are coordinating on one of the auction's multiple equilibria.

When the bidder faces Knightian uncertainty about other traders' strategies but not their redemption values, then bidder i can calculate the expected maximum regret of a bid b_i in the following way. First, calculate the maximum regret conditional on the realization of the other traders' valuations and costs, $R_B(b_i|v, c)$. Then, take the expectation of maximum regret, $\bar{R}(v_i, b_i)$ given the distribution of the other trader's valuations and costs.

The intuition for why the bid that minimizes expected maximum regret is generally different from the minimax regret bid has to do with the two sources of regret for a bidder. A bidder may regret bidding too high, and winning at an unnecessarily high price. This regret occurs if enough of the other traders' bids and asks turn out to be low, so that the lower bound on the range of market clearing prices is lower

²For a seller with cost c_i , any ask $a_i \in [c_i, \bar{c}]$ is rationalizable. For a buyer with cost v_i , any bid $b_i \in [\underline{b}, v_i]$ is rationalizable.

than the bidder's bid. On the other hand, the bidder may regret bidding too low, and missing a profitable trade. This regret occurs if enough of the other traders' bids and asks are relatively high, so that the lower bound on the range of market clearing prices is greater than the bidder's bid (but less than the bidder's valuation).

Taking expectations affects the calculations of the bidders' two sources of regret differently. It is always possible under any realization of others' redemption values for the bidder's bid to be too low, since the sellers could conceivably submit asks that are higher than the bidder's bid. On the other hand, a bid can only turn out to be too high if at least one seller submitted an ask lower than v . But since no seller will submit an ask below his actual cost, this is only possible under certain realizations of others' redemption values.

The consequence is that traders' bids and asks will be closer to their redemption values under this decision rule than they would under minimax regret, as stated in the second result.

Theorem 2 *Let \tilde{F} denote the cumulative distribution function of the lowest cost among the n sellers in the market. The bid b_i that minimizes expected maximum regret for bidder i with valuation v_i satisfies*

$$\tilde{F}\left(\frac{(k+1)b_i - v_i}{k}\right) = \frac{\tilde{F}(b_i)}{1+k}$$

Such a bid b_i exists on the interval $\left[\frac{v_i}{1+k}, v_i\right]$.

Similarly, let \tilde{G} denote the cumulative distribution function of the highest valuation among the m sellers in the market. The ask a_i that minimizes expected maximum regret for seller i with cost c_i satisfies

$$\tilde{G}\left(\frac{a_i(1+(1-k))-c_i}{1-k}\right) = \frac{\tilde{G}(a_i)+(1-k)}{1+(1-k)}$$

Such a bid a_i exists on the interval $\left[c_i, \frac{c_i+(1-k)}{1+(1-k)}\right]$.

If the auction rule k is strictly between 0 and 1, then the bid that minimizes expected maximum regret will be strictly greater than $\frac{v_i}{1+k}$, and the ask that minimizes expected maximum regret will be strictly less than $\frac{c_i+(1-k)}{1+(1-k)}$. Contrast this result with the minimax regret bids and asks (when the traders do not use beliefs about the distribution of other traders' valuations and costs). Minimizing expected maximum regret results in strategies closer to so-called "sincere bidding".

2.4.2 Minimizing Expected Maximum Regret in Large Markets

This decision rule results in strategies closer to sincere bidding, but that effect diminishes as the size of the market grows. The reason that minimizing expected maximum regret results in more truthful bids and asks is that this approach puts less weight on scenarios in which it is possible to regret bidding too high or asking too little. But the more sellers there are in the market, the more likely it is that at least one seller will have a cost lower than a given bid. And the more buyers there are in the market, the more likely it is that at least one buyer will have a valuation greater than a given ask.

As the number of traders on the other side the market increases, the trader minimizing expected maximum regret misrepresents his redemption value *more*. In the limit, the trader's bid or ask converges to the fraction of his valuation or cost

that we found using the first approach.

Corollary 1 *Let $b(v; n)$ denote the bid that minimizes expected maximum regret in a k -double auction with n sellers. Then $\lim_{n \rightarrow \infty} b(v; n) = \frac{v}{1+k}$.*

Figure 2.5 demonstrate this point in the case that an equal number of buyers and sellers with redemption values uniformly distributed over $[0, 1]$ participate in a $\frac{1}{2}$ -double auction. Each bold line in the left-hand figure denotes a bidding function given a certain number of sellers. If there is only one seller, then the bidding function is significantly closer to truth-telling (the dashed line showing the function $V = v$) than it is to the minimax regret bid, $V = \frac{v}{1+k} = \frac{2v}{3}$. The bidding function approaches $\frac{2v}{3}$ rapidly as the number of sellers increases. Likewise, the right-hand figure shows how a seller will overstate his cost for any number of bidders, and the amount of overstatement increases to $\frac{2(c+1)}{3}$ as the number of buyers increases.

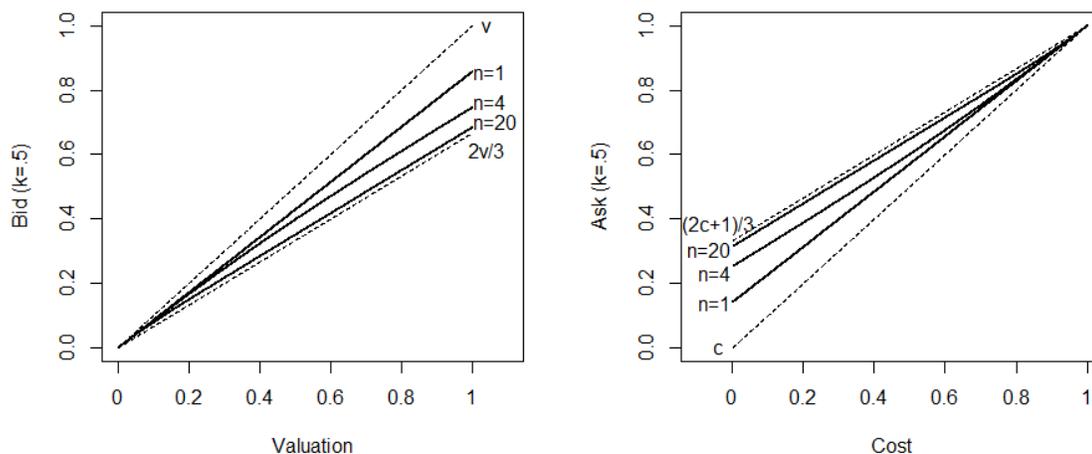


Figure 2.5: Bids and asks for traders that minimize expected maximum regret, at various market sizes

2.5 Third Approach: minimizing maximum expected regret

In this section, I find sufficient conditions for price-taking behavior when traders use multiple priors to minimize maximum expected regret. Unlike the traders examined in the previous two sections, these new regret-minimizers do not have completely unconstrained priors in any part of the decision problem. This difference is key to the possibility of convergence to price-taking behavior.

2.5.1 Minimizing Maximum Expected Regret defined

This decision rule Stoye (2011*b*) refers to as Γ -minimax regret; he defines it in this way:

Definition 3 *Let Γ denote a set of probability distributions on \mathcal{S} . An action a attains Γ -minimax regret if*

$$a \in \arg \min_{a \in \mathcal{A}} \left\{ \max_{\pi \in \Gamma} \left\{ \int \max_{a^* \in \mathcal{A}} u(a^*, s) - u(a, s) d\pi \right\} \right\} \quad (2.4)$$

This decision rule bridges the gap between expected utility maximization and minimax regret, via the choice of the set of priors Γ . If Γ includes all possible priors, then the prior(s) π that will maximize the expected regret of action a will be the prior(s) assigning probability 1 to the event that $u(a, s) = \min_{s \in \mathcal{S}} u(a, s)$. Then minimax expected regret will correspond to minimax regret. On the other hand, if Γ is a singleton π , then the maximum expected regret of each action is simply the

expected payoff under π . Then minimax expected regret will correspond to expected utility maximization.

2.5.2 Sufficient Conditions for Convergence to Truthful Bidding by Maximum Expected Regret Minimizers

Convergence to price-taking under this decision rule will depend on which priors the trader includes in his set of priors Γ . This is clear from the range of decision rules that are included in minimax expected regret. Minimax expected regret includes minimax regret, which does not induce convergence to price-taking, when all priors are included in Γ . It also includes expected utility maximization, which can induce convergence to price-taking, when Γ is a singleton. The conditions on Γ that allow for convergence is the subject of this section.

We introduce some additional notation here, in order to discuss clearly the possibility of convergence to price-taking under a set of priors Γ . Convergence will take place (or fail) as the market grows, so we must specify how the market grows, as well as the prior(s) that the agent applies to each market.

Let $\{(m_i, n_i)\}_{i=1}^{\infty}$ be a sequence of markets. Market i has m_i buyers and n_i sellers. Let $\Gamma = \{\Gamma_i\}_{i=1}^{\infty}$ be the sequence of the bidder's set of priors over the rival bids and asks. Γ_i is the set of priors over ζ for market i . A typical member of Γ is $G_\gamma = \{G_{\gamma,i}\}_{i=1}^{\infty}$ where $G_{\gamma,i} \in \Gamma_i$ is a joint distribution of the m_i^{th} and $m_i + 1^{th}$ order statistics in the market of size (m_i, n_i) .

Theorem 3 *Suppose that the following conditions hold for $\Gamma = \{\Gamma_i\}_{i=1}^{\infty}$:*

1. For each sequence of priors $\{G_{\gamma,i}\} \in \{\Gamma_i\}_{i=1}^\infty$, for every $\epsilon \in (0, v)$, there exists $N(\epsilon, G_\gamma) \in \mathbb{N}$ such that for all $i \geq N(\epsilon, G_\gamma)$:

$$\begin{aligned} \int u(v - \epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ > \int u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \end{aligned} \quad (2.5)$$

for all $b' < v - \epsilon$. That is, under each prior $\{G_\gamma\}_{i=1}^\infty \in \{\Gamma_i\}_{i=1}^\infty$, the utility-maximizing bid converges to v over the sequence of markets.

2. There exists a well-defined function

$$\bar{N}(\epsilon) = \max_{G_\gamma \in \Gamma} \{N(\epsilon, G_\gamma)\} \quad (2.6)$$

Then the bid that minimizes maximum expected regret under $\{\Gamma_i\}_{i=1}^\infty$ converges to truthful bidding over the sequence of markets $\{(m_i, n_i)\}_{i=1}^\infty$.

This Theorem states that *if* the growth of the market, and the priors over the distribution of bids and asks as the market grows, are such that an expected utility maximizer would converge to truthful bidding under each of the priors in the set (and the priors are bounded away from any priors which would fail to satisfy that condition), then a minimax expected regret bidder will *also* converge to price-taking behavior. Note that the growth of the market has been purposefully left undetermined, as has been the ratio of buyers to sellers in the limit.

These are sufficient conditions for convergence to truthful bidding by traders that minimize expected maximum regret. Are these conditions “easy” or “hard” to

satisfy? Some examples of straightforward priors easily satisfy the conditions.

For example, if the trader believes that all of the bids and asks are iid draws from some distribution $f(\cdot)$ which is bounded away from zero, then the regret minimizing bid will approach truthful bidding as the number of other bidders becomes large. From the theorem above, we can therefore conclude that a maximum expected regret minimizer will converge to truthful bidding in any market in which the number of bidders increases without bound, so long as each prior f in the set of priors Γ satisfies $f(x) > \epsilon$ for some positive ϵ , for all x in the range of possible bids and asks.

Lemma 1 *Let $\{(m_i, n_i)\}_{i=1}^{\infty}$ be a sequence of markets in which the m_i buyers approaches infinity. Let each $G_\gamma = \{G_{\gamma,i}\}_{i=1}^{\infty}$ in Γ be a joint distribution of the m_i^{th} and $m_i + 1^{\text{th}}$ order statistics in which all bids and asks are treated as $(m_i + n_i - 1)$ iid draws from a distribution $f_\gamma(x)$, where $f_\gamma(x) > \epsilon > 0$.*

Then the bid that minimizes the maximum expected regret will approach truthful bidding as $i \rightarrow \infty$.

2.6 Conclusion

This exploration of regret minimizing traders' behavior in k -double auctions suggests that including even one "bad" prior can wreak havoc on a trader's tendency to converge to price-taking behavior. If permitted to take into account any and all such pathological priors, as in minimax regret, then traders will misrepresent their redemption values, and never adjust their bids and asks in response to the market. Restricting traders' beliefs only regarding the other traders' valuations and costs,

but imposing no beliefs or equilibrium condition on traders' strategies beyond rationalizability, does nothing to improve market outcomes. Minimax expected regret using multiple priors can induce convergence to price-taking, if the "bad" priors are avoided.

Chapter 3

Information Aggregation in the Common Value k-Double Auction

Abstract

This chapter studies minimax regret as the criteria with which traders choose their bids/asks in auctions with common values. The analysis proves that minimax regret traders change their bids as the number of traders in the market increases. In fact, minimax regret traders' bids and asks are invariant to the number of other traders in the market. Whether traders can avoid the winner's curse depends on the auction rule and the distribution of private signals.

Keywords: double auctions; regret minimization; Knightian uncertainty; decision theory; mechanism design; common values

JEL Classification Numbers: D44, D81, D82, C72

3.1 Introduction

The previous chapter dealt with sealed-bid auctions in which each trader has a private redemption value. In this chapter, we turn our attention to common value auctions. In a pure common value auction, the item(s) traded have the same value to all trader in the market, but this true value is unknown to the traders. Instead, each trader observes a private signal and estimates the true value.

An important question is how well the market will aggregate information – that is, whether the market price will converge to the true value. The rational expectations hypothesis (Lucas, 1978) assumes that the price reflects all of the information in the market. Since that idea’s introduction to economics and finance, the search has been on for micro foundations to support it. Radner (1979) shows that if there is a finite number of alternative states of initial information, a rational expectations equilibrium is generic. Another step toward establishing such a micro foundation is Reny and Perry (2006). In that model, traders have interdependent values and affiliated private information. In a sealed bid double auction, the price converges to the asset’s true value as the market grows. This convergence stems from individual bidders’ diminishing incentive to attempt to manipulate the price as the market grows. That incentive diminishes in larger markets because the probability of being pivotal decreases in larger markets. The mechanism by which competition in this market achieves convergence of the price to the true value strongly resembles the mechanism by which a private-value double auction achieves convergence of the market allocation to efficiency.

This chapter asks whether minimax regret traders will bid (and ask) in such a

way that the resulting price(s) in an auction (or double auction) will converge to the good's true value. Minimax regret traders in common value auctions have been studied before: Hayashi and Yoshimoto (2012) have created and calibrated a risk- and regret- averse model for bidders in first price auctions. However, these traders have not been studied in the double auction institution.

One of the contributions of this chapter is the introduction of a solution concept that is similar to an equilibrium, but that is suitable for traders facing Knightian uncertainty. Since equilibrium concepts generally involve correct beliefs about other players, defining such a solution concept requires that Knightian Uncertainty extend to only part of the decision problem. It is important to distinguish between the different types of information that a trader playing the Bayes Nash equilibrium in this game would use to determine his or her bid. The trader could be in a state of Knightian uncertainty regarding the true value of the item (in which case, the distribution of ω is unknown), or of the distribution of signals (in which case, the distribution of s is unknown), or of the strategies of the other players, or of some other aspect of the environment, such as the number of other bidders. In order for traders to play an "equilibrium" strategy, they must have some beliefs about the strategies of other players. It is still possible for traders to face Knightian uncertainty regarding the distribution of other trader's redemption values, while still holding beliefs about what those traders will do given their redemption values. This is the approach to "mutually minimizing maximum regret" defined in this chapter.

The following section sets up a pure common value model, and begins with examining one-sided auctions with minimax regret bidders. In a common value Vickrey

auction, the minimax regret bid is to bid one's signal, and there is a winner's curse. In a common value first price sealed bid auction, the minimax regret bid is reduced from one's signal, and it is possible that the bidders will avoid the winner's curse. In fact, there is a "minimax regret equilibrium" in which the bidders make the most "pessimistic" bids (that is, the bid equals the lowest possible value of the asset, given the observed signal) and the price will converge to the true value of the asset as the number of traders grows.

We next turn our attention to the k -double auction. In a common value k -double auction, it is possible to induce both sides of the market to truthfully reveal their signals when $k = \frac{1}{2}$. In this case, if the number of buyers and the number of sellers converged to equal amounts as the market grows, then the market price would converge to the value of the median signal.

3.2 Regret Minimizers in Common Value Sealed-Bid Auctions

3.2.1 The Model

There are m buyers for an item. The true value of the item is $\omega \in \Omega$. This true value is unknown to the bidders. Instead, each bidder observes a signal $s_i \in (\omega - \epsilon, \omega + \epsilon)$.

As a minimax regret decision maker who does not use a prior over the distribution of his own or other signal, therefore, each bidder knows only that the true value ω is within the range $(s_i - \epsilon, s_i + \epsilon)$. It follows that any other bidders' signals could be

$s_j \in (s_i - 2\epsilon, s_i + 2\epsilon)$, depending on the state of the world.

3.2.2 Vickrey Auctions: Truthful Bidding, Winner's Curse

Theorem 4 *In a common value Vickrey auction, the minimax regret bid is $b_i = s_i$.*

This minimax regret strategy does not change with the number of bidders.

As a result of this truthful bidding, the price will be equal to the second-highest signal. As long as the distribution of signals is such that the second-order statistic approaches the maximum possible value for the signals ($\omega + \epsilon$), the price will converge to $\omega + \epsilon$ as the number of bidders increases. These minimax regret bidders do nothing to avoid the winner's curse in a Vickrey auction.

3.2.3 First Price Auction: Pessimistic Bidding, Reduced Winner's Curse

Lemma 2 *In a common value first price sealed bid auction, a bidder who observed signal s_i has the following maximum regret function:*

$$\sup R = \max\{b_i - \inf_{j \neq i} \{ \max b_j \}, b_i - s_i + \epsilon, s_i + \epsilon - b_i\} \quad (3.1)$$

It is clear from the maximum regret function that the bidder's minimax regret bid depends on the lowest possible bid that the other bidders could place in the first price auction. One way to approach this would be to define the set of feasible bids (perhaps, any bid between zero and the highest possible value of the asset), and to always take the lowest value in that set as the lowest possible bid.

Lemma 3 *Suppose that the minimax regret bidder in the first price auction always considers a set of possible rival bids with minimum value of zero (no matter what signal he observes). Then the bid which minimizes his maximum regret is*

$$b_i = \begin{cases} s_i & \text{if } s_i \leq \epsilon \\ \frac{s_i + \epsilon}{2} & \text{if } s_i > \epsilon \end{cases} \quad (3.2)$$

The above approach does not take into account the bidder's information about other bidders' signals that is provided by his own signal. However, a bidder who observes the signal s_i knows that the other bidders' signals cannot possibly be less than $s_i - 2\epsilon$ or greater than $s_i + 2\epsilon$. Therefore, an alternative approach readily suggests itself. Instead of supposing that the other bidders could bid the lowest feasible bid, no matter what signals they can have observed, the bidder might instead suppose that the other bidders are using the same strategy that he himself is using, and that their bids will only differ from his based on the difference in the signals that they observe. It follows that the lowest rival bid that the bidder could observe, given his own signal s_i and his bidding function $b(\cdot)$, would be $b(s_i - 2\epsilon)$.

Definition 4 *Let $b = (b_1, b_2, \dots, b_n)$ be a vector of bidding functions in the first price auction. If for each bidder i , their bid minimizes the maximum regret function given the bidding functions of the other bidders,*

$$b_i = \arg \min \max \{ b_i - \max_{j \neq i} b_j(s_i - 2\epsilon), b_i - s_i + \epsilon, s_i + \epsilon - b_i \} \quad (3.3)$$

then the bidders may be said to be mutually minimizing maximum regret.

This alternative approach can be thought of as subjecting the distribution of signals to Knightian uncertainty, but allowing bidders to form beliefs about the strategies of their rivals.

The notion of finding a profile of strategies such that each trader is making his or her best response to the others has some resemblance to an equilibrium concept. The difference that distinguishes my approach from an equilibrium is that traders do not use one particular belief for the distribution of other traders' signals when they choose their best response – much less an accurate belief about that distribution, which is typically part of the equilibrium in a game of incomplete information. In other words, traders do not have correct beliefs about the behavior of the “player” Nature. However, they do have correct beliefs about the strategies of fellow-traders. This approach can be justified if Nature is inscrutable in a way that other traders are not. A trader may find it easy to imagine that other traders make decisions in a manner similar to himself, for example; but the processes that determine how traders' information relates to the asset's underlying value may be far more complex.

Theorem 5 *If each bidder believes that the other bidders will reduce their bid from their own signal by the maximum possible error,*

$$b(s) = s - \epsilon, \tag{3.4}$$

Then this same bidding function will minimize the bidder's own maximum regret. Thus, bidders who reduce their bids from their signals by the maximum error ϵ will mutually minimize maximum regret.

When each bidder uses this minimax equilibrium strategy, the bidder with the highest signal wins the item and pays the price $p = s_i - \epsilon$. If it is the case that the first order statistic of the signals converges to the highest possible signal $\omega + \epsilon$ as the number of bidders grows, then the price converges to the true value of the asset ω . The bidders do not experience a winner's curse.

3.3 Regret Minimizers in Common Value

k-Double Auctions

We are now prepared to turn our attention to a more complicated setting: the *k*-double auction.

3.3.1 The Model

As in the previous section, let the true value of the asset be ω , and let each trader draw a signal $s_i \in (\omega - \epsilon, \omega + \epsilon)$.

Let there be m sellers and n buyers. Each seller possesses 1 unit of the asset.

3.3.2 Minimax Regret Bid

Lemma 4 Let $\check{\zeta}_{(m)}$ denote the lowest possible value of $\zeta_{(m)}$ given the signal s_i observed by the bidder. Then the minimax regret bid is

$$b = \begin{cases} s & \text{if } s \leq \frac{\epsilon}{k} + \check{\zeta}_{(m)} \\ \frac{s + \epsilon + k\check{\zeta}_{(m)}}{1+k} & \text{if } s > \frac{\epsilon}{k} + \check{\zeta}_{(m)} \end{cases} \quad (3.5)$$

Note that the buyer's bid will be equal to or less than his signal.

3.3.3 Minimax Regret Ask

Lemma 5 Let $\hat{\zeta}_{(m+1)}$ denote the highest possible value of $\zeta_{(m+1)}$ given the signal s_i observed by the seller. Then the minimax regret ask is

$$a = \begin{cases} s & \text{if } s \geq \hat{\zeta}_{(m+1)} - \frac{\epsilon}{1-k} \\ \frac{s - \epsilon + (1-k)\hat{\zeta}_{(m+1)}}{1+(1-k)} & \text{if } s < \hat{\zeta}_{(m+1)} - \frac{\epsilon}{1-k} + \check{\zeta}_{(m)} \end{cases} \quad (3.6)$$

3.3.4 Mutual Minimax Regret Bids and Asks for Some Values of k

Theorem 6 The following bid and asks functions mutually minimize maximum regret.

1. If $k = 0$: $b = s$, and $a = s + \frac{\epsilon}{2}$.
2. If $k = 1$: $b = s - \frac{\epsilon}{2}$, and $a = s$.

3. If $k = \frac{1}{2}$: $b = s$, and $a = s$.

It is interesting that in a k -double auction with $k = \frac{1}{2}$, both the buyers and sellers find that truthful bidding minimizes their maximum regret. With truthful bidding in a common value auction, if the numbers of buyers and sellers is kept equal as the market grows, the price will converge to the median signal – which for many distributions would be the true value of the asset.

3.4 Conclusion

In general, the results of these pure common value models are not in line with the experimental literature. In the laboratory as well as in the field, economists have observed a winner’s curse. Bidders are not as “pessimistic” as the minimax regret bidders in a first price auction. This should make us cautious about taking the results of the model at face value. Nevertheless, it is instructive that minimax regret bidders, although they do not use a prior over the distribution of signals, can be induced to submit their true signals. The results for k -double auctions indicate that balancing the influence of the buyers and sellers on the price can be important in this setting for soliciting the true signals of both buyers and sellers.

Chapter 4

The Symmetry Axiom and Strategies Invariant to the Number of Players

Abstract

We consider whether competitive pressures can induce traders to truthfully report their private redemption valuations under Knightian uncertainty. Traders face Knightian uncertainty if they know the possible outcomes of each available action, but do not know each outcome's probability. Such uncertainty may motivate use of a decision rule other than expected utility maximization. Two such alternative decision rules are maxmin and minimax regret. Stoye's (2011) axiomatic characterization of these decision rules reveals that there is one axiom that maxmin and minimax regret share, and that distinguishes them from Bayes rule: the axiom of symmetry.

We find that if agents use decision rules that accord with the symmetry axiom, then their strategies will be invariant to the number of other players in the game. Consequently, a market populated by traders that follow the symmetry axiom will not converge to efficiency as the market grows.

Keywords: double auctions; regret minimization; Knightian uncertainty; deci-

sion theory; mechanism design

JEL Classification Numbers: D44, D81, D82, C72

4.1 Introduction

4.1.1 Motivation and Background

A crucial aspect of the perfect competition model is that buyers and sellers are price-takers who truthfully report their utility-maximizing quantity to produce or consume at the market price, and do not attempt to manipulate prices. This is important because when agents act as price-takers, the market outcome is efficient.

The justification for assuming price-taking behavior by strategic agents is that in a market with many buyers and sellers, each individual is unlikely to affect the terms of trade, and if he were to do so, his effect would be negligible. Thus, agents that make their choices strategically to maximize their individual profit will find that they cannot manipulate the market price in their favor by misrepresenting their preferences. The same agent who would curtail supply if he had a monopoly will, when competing with many other sellers, produce the quantity at which his marginal cost equals the market price, taking the market price as given. For this reason, efficiency can be predicted for a market populated by many agents, each of whom believes that the size of the market attenuates the influence of individual traders of similar size. The greater the number of buyers and sellers, the closer the market outcome to efficiency.

A double auction market provides a model where the above reasoning can be proven formally (Wilson, 1987). A complete model of such a market must specify the information, preferences, and behavior of its agents. Typically, these models assume that agents are strategic and capable of coordinating on an equilibrium in which each agent maximizes von Neumann Morgenstern expected utility, calculated based on prior beliefs which are common knowledge to all market participants. If we take a broader view of possible ways to characterize agents in these models, we can think of the agent as confronted with a decision problem that he may or may not use prior beliefs to solve.

Given the importance of market efficiency, it is worthwhile to examine the extent to which efficient market outcomes depend on the market participants' information and capabilities. We want to better understand what *sort* of agent will act as a price-taker in a large double auction in order to identify the conditions under which we can predict an efficient outcome for a market.

Efficiency of some double market institutions has been well established for certain types of agents. For example, we know that bidders who play Bayesian Nash Equilibrium strategies in a k -double auction will converge to efficiency as the market grows large (Rustichini, Satterthwaite and Williams, 1994). However, the Bayesian Nash equilibrium solution concept brings with it significant limitations.

It may not be realistic to assume that traders have the capacity to calculate and coordinate on a Bayesian Nash Equilibrium. One important reason for this is that the traders may face Knightian uncertainty – uncertainty about the underlying distribution of other traders' types, along with uncertainty about the realization of

those types (Luce and Raiffa, 1957). A Bayesian trader's response to such Knightian uncertainty is to adopt a prior, but we may have good reason to reject this approach. Even adopting a uniform prior asserts some knowledge about a decision problem, knowledge that the decision maker may not have: the knowledge of which states of the world are relevant and how they relate to the decision-makers payoffs. If the decision-maker's ignorance is so complete that he does not know what the relevant characteristics of events are, then his decision rule ought not depend on the way he has chosen to specify the problem (Arrow and Hurwicz, 1972). A decision rule that reflects this kind of uncertainty will be significantly different from Bayesianism.

Since alternatives to Bayesian traders have these attractions, it is worthwhile to consider carefully what we are giving up when we give up Bayesianism. We can take alternative decision rules (such as minimax regret and maxmin) on a case-by-case basis, but the conclusions that follow are then very narrow. Instead, this chapter investigates the axioms that may characterize agents' decision rules. Considering the axioms that underpin the decision rules allows us to draw conclusions about whole classes of decision rules.

4.1.2 Related Work

This chapter is related to the literature on decision-making under Knightian uncertainty, and more particularly to work that applies alternatives to expected utility maximization to bargaining and auctions.

Linhart and Radner (1989) have studied bilateral bargaining for minimax regret agents. They find that players minimizing maximum regret will realize expected gains

from trade equal to half of the potential expected gains from trade. In contrast, the Bayesian Nash equilibria in the same setting yield expected gains from trade ranging from 0 in the no-trade equilibrium, to 84.4% in the “second-best” equilibrium that can be achieved under incomplete information (p. 173).

One application of our main result is that the per capita expected gains from trade realized by minimax regret agents do not improve if we change the setting from the bilateral case to a larger market. In contrast, the “second best” Nash equilibrium converges to efficiency as market size increases (Aldo Rustichini, Mark A. Satterthwaite and Steven R. Williams, 1994, Theorem 3.2), as per capita market inefficiency is $O(1/m^2)$, where m is the number of buyers.

Thus, you could say that we are faced with a tradeoff. On the one hand, the minimax regret approach has the advantage of offering a single and straightforward recommendation to the decision-maker. On the other hand, this same recommendation is made in all markets regardless of the market’s size! If this offends our intuition of how we think traders ought to behave in large markets, or how we expect traders actually do behave, then we may conclude that the large double market is one of the cases in which minimax regret simply does not seem plausible.

Finding a limitation in the usefulness of minimax regret does not negate that decision rule’s value. It is certainly legitimate to vary one’s choice of decision rules depending on the type of decision being made. However, we may have a broader goal in mind – a goal to understand what assumptions on human behavior do the best job of generating intuitive results for both bilateral bargaining *and* large market settings. It is with this goal in mind that this chapter investigates the characteriza-

tion and consequences of a variety of decision rules. Stoye (2011b) provides a helpful overview of decision rules for Knightian uncertainty; his characterizations of Bayes Rule, maxmin, and minimax regret, are the basis of this chapter’s approach.

4.2 Decision Problems Under Knightian Uncertainty

This section will consider decision rules from a more general perspective than the double auction. We will return to the double auction in section ?? Decision problems involve a set of acts \mathcal{A} available to the decision-maker, the set of possible states of the world \mathcal{S} , and the outcome $u \in \mathcal{U}$ that results in the state $s \in \mathcal{S}$ given the decision-maker’s action $a \in \mathcal{A}$. We may find it useful to think of the decision-maker as having a payoff function $u : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$.

Example 5 *Suppose that a decision-maker has a choice between two actions, a_1 and a_2 , whose payoffs depend on the realization of random variable s that has support $[0, 1]$ and unknown distribution:*

$$u(a_1, s) = \begin{cases} 5 & \text{if } 0 \leq s < \frac{1}{2} \\ 1 & \text{if } \frac{1}{2} \leq s < 1 \end{cases} \quad (4.1)$$

$$u(a_2, s) = \begin{cases} 1.5 & \text{if } 0 \leq s < \frac{1}{3} \\ 2.5 & \text{if } \frac{1}{3} \leq s < \frac{2}{3} \\ 3.5 & \text{if } \frac{2}{3} \leq s < 1 \end{cases} \quad (4.2)$$

In Figure 4.1, the payoff of a_1 is shown in a solid line and the payoff of a_2 is shown in a dashed line.

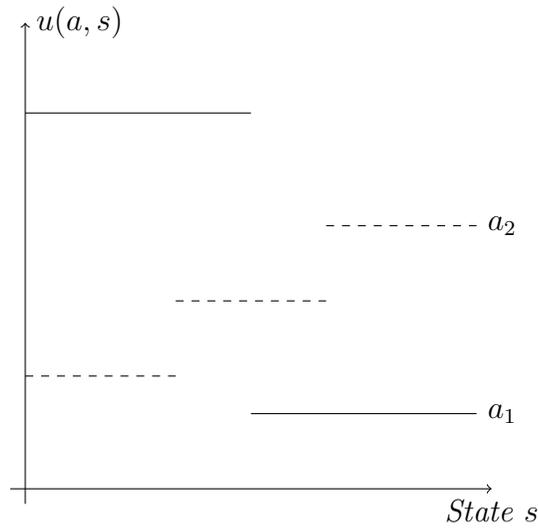


Figure 4.1: A Decision Problem Example

Section 4.2.1 explores three decision rules, and considers how each one handles this straightforward decision problem.

4.2.1 Three Decision Rules

A decision rule specifies what a decision maker will do given a menu \mathcal{A} of possible actions. In the following sections, we will characterize various decision rules using axioms that apply to decision makers' preferences \succsim . Let $\succsim_{\mathcal{A}}$ denote a preference relation over the actions available in the menu \mathcal{A} . A preference relation is defined

to be a binary relation \succsim that is reflexive ($a \succsim a$ for all actions a) and transitive (if $a_1 \succsim a_2$ and $a_2 \succsim a_3$, then $a_1 \succsim a_3$) (Fishburn, 1970). From \succsim we can derive relations \succ and \sim in the usual way.

Expected Utility

Maximizing expected utility depends on a prior that assigns a probability to each state of the world.

Definition 6 *Let π be a probability distribution on \mathcal{S} . An action a maximizes expected utility under prior π if*

$$a \in \arg \max_{a \in \mathcal{A}} \int u(a, s) d\pi. \quad (4.3)$$

Example 6 (continued) *An expected utility maximizer's preferences over a_1 and a_2 depend, of course, on his subjective prior concerning the probability of the possible realizations of s . If, for example, the decision maker has a uniform prior (treating all possible values of s as equally likely) then the expected payoff of a_1 is 3 and the expected payoff of a_2 is 2.5, so a_1 would be preferred.*

Maxmin

This decision rule was first proposed by Wald (1945).

Definition 7 *An action a attains maxmin payoff if*

$$a \in \arg \max_{a \in \mathcal{A}} \min_{s \in \mathcal{S}} u(a, s) \quad (4.4)$$

Example 7 (continued) Under *maxmin*, $a_2 \succsim a_1$. The minimum payoff under a_2 is 1.5, which is higher than the minimum payoff under a_1 of 1.

Minimax Regret

Minimax regret was suggested by Savage (1951) as an alternative to *maxmin*.

Definition 8 An action a attains *minimax regret* if

$$a \in \arg \min_{a \in \mathcal{A}} \max_{s \in \mathcal{S}} \{ \max_{a^* \in \mathcal{A}} u(a^*, s) - u(a, s) \} \quad (4.5)$$

The action(s) minimizing maximum regret are identified by calculating the maximum regret that could be incurred under each action. The regret for a particular action in a particular state is calculated by comparing that action's payoff to the maximum possible payoff in the same state.

Example 8 (continued) In our example decision problem, the optimal action in each state of the world is:

$$a^* = \begin{cases} a_1 & \text{for } s < \frac{1}{2} \\ a_2 & \text{for } s \geq \frac{1}{2} \end{cases} \quad (4.6)$$

Therefore, the regret function $R(a, s)$ for each action is:

$$R(a_1, s) = u(a^*(s), s) - u(a_1, s) = \begin{cases} 0 & \text{for } s < \frac{1}{2} \\ 1.5 & \text{for } \frac{1}{2} \leq s < \frac{2}{3} \\ 2.5 & \text{for } \frac{2}{3} \leq s < 1 \end{cases} \quad (4.7)$$

$$R(a_2, s) = u(a^*(s), s) - u(a_2, s) = \begin{cases} 3.5 & \text{for } s < \frac{1}{3} \\ 2.5 & \text{for } \frac{1}{3} \leq s < \frac{1}{2} \\ 0 & \text{for } \frac{1}{2} \leq s < 1 \end{cases} \quad (4.8)$$

Thus, when we apply the minimax regret decision rule to our example, $a_1 \succsim a_2$. The maximum regret when a_2 is chosen is 3.5. The maximum regret when a_1 is chosen is 2.5.

From the standpoint of a person accustomed to working with expected payoffs, it may seem that the maxmin and the minimax regret decision rules operate by choosing a “pessimistic” prior – a prior that assigns higher probability to events with very low or very high payoffs. The truth is subtly different. These decision rules do not stick to a single pessimistic prior by which each action is evaluated. Instead, these decision rules evaluate each action by focusing exclusively on the state in which the payoff is lowest or the regret is highest *for that action*. Of course, this is equivalent to using a prior that assigns probability 1 to the event that corresponds to this extreme outcome. However, the prior that is used to evaluate action a_1 may

be very different from the prior that is used to evaluate action a_2 .

4.3 Axioms for Decision Problems Under Knightian Uncertainty

4.3.1 Four Key Axioms

Stoye (2011) supplies a useful axiomatic analysis of decision rules under Knightian uncertainty. Completeness, von Neumann Morgenstern independence, independence of irrelevant alternatives, and symmetry are four key axioms in his formulation.

Completeness

Definition 9 For any actions a_1, a_2 , and menu \mathcal{A} containing $\{a_1, a_2\}$: $a_1 \succsim_{\mathcal{A}} a_2$ or $a_2 \succsim_{\mathcal{A}} a_1$.

In other words, any two payoff functions can be compared with one another.

von Neumann Morgenstern Independence

Definition 10 $a_1 \succsim a_2 \Leftrightarrow \delta a_1 + (1 - \delta)a_3 \succsim \delta a_2 + (1 - \delta)a_3$ for all $\delta \in [0, 1]$.

In defense of this axiom, Stoye puts forward the following example. Suppose that a person is considering two payoff functions, a_1 and a_2 . If the decision maker is told that his choice will only be enacted with probability δ , and that otherwise he will get payoff function a_3 , will that have any effect on his preferences? The independence axiom says no – even if a_3 provided a hedge against a_2 .

Because independence is foundational to expected utility theory, this independence condition seems very reasonable – especially to a person accustomed to evaluating alternatives based on expected payoffs. Human behavior does not always conform to this axiom, however. This is demonstrated in the Allais paradox.¹

Abandoning independence in order to conform to the preferences revealed in the Allais paradox is costly. If a decision rule does not satisfy independence (but is an ordering that respects stochastic dominance), then it will fail “sequential coherence” (Seidenfeld, 1988, p. 281)

Given its importance for guaranteeing coherent choices, independence seems to be a reasonable axiom in the context of an auction.

Independence of Irrelevant Alternatives

In decision theory, Marschak and Radner (1954) formulated Independence of Irrelevant Alternatives (IIA-RM) in the following way.²

Definition 11 *If $a_1 \succ_{\{a_1, a_2\}} a_2$, then for any menu \mathcal{A} containing $\{a_1, a_2\}$, $a_1 \succ_{\mathcal{A}} a_2$.*

¹Maurice Allais constructed two decision problems, each a choice between two lotteries. A decision-maker is first asked to choose between the following two lotteries: lottery A yields a payoff of 100 with certainty; lottery B yields a payoff of 500 with probability .1, 100 with probability .89, and zero with probability .01. In the second decision problem, lottery C yields a payoff 100 with probability .11 and zero with probability .89; lottery D yields a payoff of 500 with probability .1, and zero with probability .9. Faced with these two decision problems, decision-makers often prefer lottery A over lottery B but prefer lottery D over lottery C. In fact, lotteries C and D are the same as lotteries A and B, respectively, with the simple change that an 89% chance of winning 100 is replaced with an 89% chance of winning nothing. Thus, preferring A to B but preferring lotteries D to C violates von Neumann Morgenstern Independence. One way to interpret this inconsistency of revealed preferences is that people over-value complete certainty.

²There are multiple versions of Independence of Irrelevant Alternatives. The title has been given to conditions used in social welfare theory, decision theory, and probabilistic choice (Ray, 1973). Arrow (1987) described the difference between IIA-RM and his definition of a related idea that he used in his famous impossibility theorem in this way: IIA-RM “refers to variations in the set of opportunities, mine to variations in the preference orderings.”

The appeal of Independence of Irrelevant Alternatives is easily illustrated. It would be strange for a decision maker to select apple pie from a menu of apple and blueberry, but switch his order to blueberry once he is informed that cherry pie is also available. On the other hand, it may be reasonable to violate IIA if the existence of alternatives reveals information (Luce and Raiffa, 1957, p. 288).

In the context of an auction, the existence of alternatives is unlikely to reveal information about the rival bids. One might consider, however, the possibility that strategic considerations could cause a bidder to violate IIA. If a bid becomes available or unavailable, it is worthwhile to consider the possibility that knowing this could influence the strategy of other bidders and therefore the expected payoff of the bidder's bid.

For its role in providing a foundation in evaluating outcomes based on expected utility, however, Independence of Irrelevant Alternatives is unlikely to be rejected in an auction setting.

Symmetry

The axiom of symmetry is motivated by a desire not to give undue weight to any state of the world. The agent is unwilling to express preferences that would imply a knowledge of the relative likelihood of separate events. Therefore, a preference for one action over another should be unaffected if the outcome attributed to a certain state – or set of states – is swapped with the outcome of another state. Essentially, the symmetry axiom requires that the decision rule not be manipulable by changes in the way that the outcomes are associated with states of the world. A change to

the list of states can be conceived as a function $\psi : \mathcal{S} \rightarrow \mathcal{S}$. The definition below specifies what sort of changes to the state space are permissible.

Condition 12 $\psi : \mathcal{S} \rightarrow \mathcal{S}$ preserves the profile of outcomes if $\psi(s_1) = \psi(s_2) \Rightarrow u(a, s_1) = u(a, s_2)$ for all actions $a \in \mathcal{A}$.

We make a few observations about this way of transforming the set of states before introducing the axiom of symmetry.

First, any simple re-labeling of states preserves the profile of outcomes. For instance, in our example with state space $\mathcal{S} = [0, 1]$, the transformation

$$\psi(s) = 1 - s \tag{4.9}$$

is a bijection that clearly preserves the profile of outcomes. In fact, any bijection $\psi : \mathcal{S} \rightarrow \mathcal{S}$ will satisfy the above definition, simply because $\psi(s_1) = \psi(s_2)$ implies that $s_1 = s_2$.

Second, in some decision problems, the measure of some events could be expanded while the measure of other events could be reduced. This is possible for a decision problem in which multiple states of the world have identical outcomes, as in example 2 above. A “constant payoff event” (that is, a set of states, all of which yield the same payoff under each action available to the agent) can be swapped with a constant payoff event that has a very different measure. Such a transformation is described below for our decision problem example.

Finally, as long as a transformation ψ preserves the profile of outcomes, we cannot “lose” any event completely. That is, if there is some state of the world s , then some

state of the world $\tilde{s} = \psi(s)$ so that $u(a, \tilde{s}) = u(a, \psi(s))$. Most notably, the minimum payoff and the maximum payoff achieved by each action will remain unchanged.

Armed with this understanding of preserving the profile of outcomes, we supply a formal definition of the symmetry axiom.

Definition 13 *Consider an agent's preferences over actions \mathcal{A} . Let $\psi : \mathcal{S} \rightarrow \mathcal{S}$ be any function that preserves the profile of outcomes. Define a transformation of the payoff of each act $a_i \in \mathcal{A}$ in the following way: $u'(a_i, s) \equiv u(a_i, \psi(s))$. Then $\succsim_{\mathcal{A}}$ satisfies the axiom of symmetry if: $a_1 \succsim a_2$ in the original problem $\Leftrightarrow a_1 \succsim a_2$ in the transformed decision problem.*

Example 14 *Consider the decision problem from the previous example. We will demonstrate an application of the symmetry axiom for this decision problem.*

Take two disjoint events E and F , each of which yields constant profit under a_1 and a_2 . Let E be the event that $s \in [0, \frac{1}{3}]$. For any $s \in E$, $u(a_1, s) = 5$ and $u(a_2, s) = 1.5$. Let F be the event that $s \in [\frac{1}{2}, \frac{2}{3}]$. For any $s \in F$, $u(a_1, s) = 1$ and $u(a_2, s) = 2.5$.

Consider the effect of exchanging the payoffs under events E and F for each action. This can be accomplished by any number of transformations ψ . For example,

$$\psi(s) = \begin{cases} \frac{1}{2} & \text{if } 0 \leq s < \frac{1}{3} \\ 0 & \text{if } \frac{1}{2} \leq s < \frac{2}{3} \\ s & \text{otherwise} \end{cases} \quad (4.10)$$

It is straightforward to verify that this transformation preserves the profile of outcomes.

In effect, we now have a new decision problem, where the payoffs of a_1 and a_2 are

$$u'(a_1, s) = \begin{cases} 1 & \text{if } 0 \leq s < \frac{1}{3} \\ 5 & \text{if } \frac{1}{3} \leq s < \frac{2}{3} \\ 1 & \text{if } \frac{2}{3} \leq s < 1 \end{cases} \quad (4.11)$$

$$u'(a_2, s) = \begin{cases} 2.5 & \text{if } 0 \leq s < \frac{1}{2} \\ 1.5 & \text{if } \frac{1}{2} \leq s < \frac{2}{3} \\ 3.5 & \text{if } \frac{2}{3} \leq s < 1 \end{cases} \quad (4.12)$$

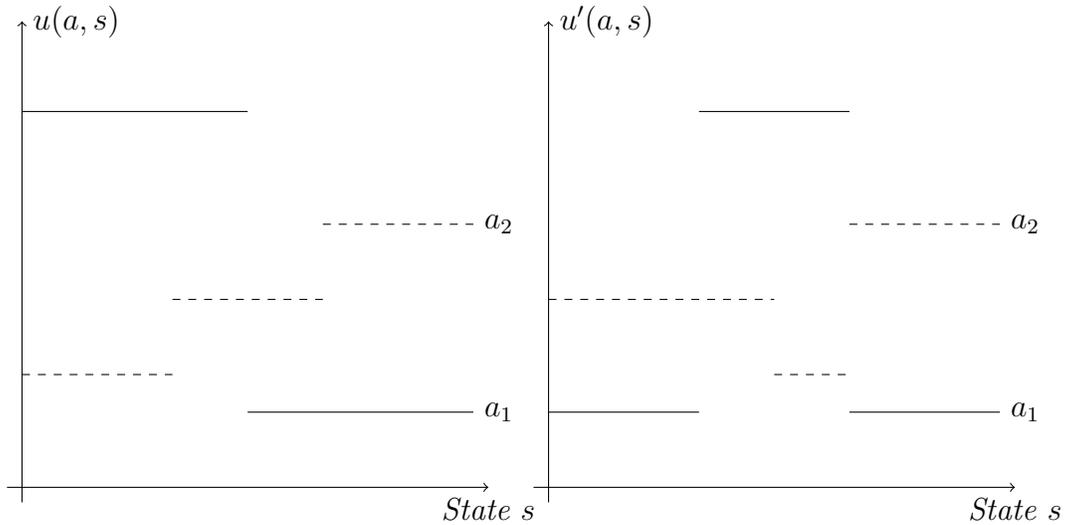


Figure 4.2: A Transformation of Decision Problem Example
 Figure 4.2 illustrates this transformation. The payoffs of action a_1 are shown in

solid lines; payoffs of action a_2 are shown in dashed lines. Symmetry requires that if $a_1 \succsim a_2$ in the decision problem pictured on the left, those same preferences hold in the decision problem pictured on the right.

In each decision problem, a_2 maximizes the minimum payoff, and a_1 minimizes maximum regret. This consistency is to be expected, since those decision rules satisfy the axiom of symmetry.

However, consider the case of a utility maximizer. It is obvious that his choice between a_1 and a_2 depends on his prior, and that even if $a_1 \succsim a_2$ in the first decision problem it need not follow that $a_1 \succsim a_2$ in the transformed decision problem. For example, if he applies the uniform prior to the transformed decision problem, then the expected profit of a_1 is $2\frac{1}{3}$ and the expected profit of a_2 is $2\frac{2}{3}$, so that $a_2 \succ a_1$, even though $a_1 \succsim a_2$ under the uniform prior in the original decision problem. Applying the uniform prior violates the axiom of symmetry in this example.

Section 4.3.2 discusses the difference between symmetry and the principle of insufficient reason.

4.3.2 Further Discussion of Symmetry

It may be surprising that a Bayesian decision-maker using uniform prior will fail to satisfy symmetry, since a uniform prior assigns the same probability to each state of the world. Take for example the following decision problem.

We might imagine that one could guarantee that one was treating each state “equally” if the preferences over a_1 and a_2 were unchanged when the outcome of two states were swapped, as in the following decision problem.

	s_1	s_2	s_3
a_1	3	3	9
a_2	6	6	0

Table 4.1: A Decision Problem with Three States

	s_1	s_2	s_3
a_1	3	9	3
a_2	6	0	6

Table 4.2: Swapping Columns in a Decision Problem

The second decision problem can be derived from the first by a bijection on the states of the world, $\psi : \mathcal{S} \rightarrow \mathcal{S}$ such as:

$$\psi(s_i) = \begin{cases} s_1 & \text{if } i = 1 \\ s_2 & \text{if } i = 3 \\ s_3 & \text{if } i = 2 \end{cases} \quad (4.13)$$

Note that under a uniform prior (where the probability of $s_i = \frac{1}{3}$ for all s_i), the expected utility of a_1 is 5 and the expected utility of a_2 is 4 in each decision problem. In fact, relabeling the states of the world (when there is a finite countable number of states) will not change the expected utility of any action, under a uniform prior. (On the other hand, if the number of states is uncountable, as when the state of the world is a random variable $z \in [0, 1]$, then we must be careful not to change the measure of any event by our transformation.)

The above example outlines how a uniform prior can select an action consistently even after a careful relabeling of states. However, adoption of a uniform prior asserts

a confidence in the agent’s understanding of the structure of the decision problem – confidence that an agent faced with Knightian uncertainty may not possess. Arrow and Hurwicz (1972) described the difficulty eloquently: “a state of nature is a complete description of the world. But how we describe the world is a matter of language, not of fact.” After all, if in the decision problem above, the agent were to discover that states of the world s_1 and s_2 correspond to the descriptions “the next vehicle is a red bus” and “the next vehicle is a blue bus,” respectively, whereas s_3 corresponds to the description of the world “the next vehicle is a taxi or a shuttle,” the decision-maker might want to change the structure of the decision table! He could do so, while still preserving the profile of payoffs, with a transformation such as the following:

$$\psi(s_i) = \begin{cases} s_1 & \text{if } i = 3 \\ s_3 & \text{otherwise} \end{cases} \quad (4.14)$$

This function ψ preserves the profile of outcomes. It swaps the constant payoff event $\{s_1, s_2\}$ with the constant payoff event $\{s_3\}$.

Applying ψ , we have a new decision problem:

	s_1	s_2	s_3
a_1	9	3	3
a_2	0	6	6

Table 4.3: Duplicating and Combining Columns in a Decision Problem

Effectively, we have combined two columns in the decision problem, while simultaneously duplicating another column.

Thus, we can think of the axiom of symmetry as allowing us to relabel columns, and also to combine or duplicate columns in a decision problem. In a decision problem with a finite and countable number of states, such as the example discussed here, this power to combine or duplicate columns has some limitations, if it can only be accomplished via a transformation $\psi : \mathcal{S} \rightarrow \mathcal{S}$ as we defined in Condition 12. However, this simple example demonstrates that even in a decision problem involving just three states of the world and two constant-payoff events, Bayes Rule may run into serious problems with the axiom of symmetry. Even when using a uniform prior, Bayes Rule and the axiom of symmetry are incompatible.

4.3.3 Additional Axioms

The following axioms are also key to Stoye's formulation, although he does not emphasize them as he does the axioms above.

Definition 15 *Mixture Continuity.* For any acts a_1, a_2, a_3 and menu \mathcal{A} containing $\{a_1, a_2, a_3\}$: The sets $\{\delta \in [0, 1] : \delta a_1 + (1 - \delta)a_3 \succsim_{\mathcal{A}} a_2\}$ and $\{\delta \in [0, 1] : a_2 \succsim_{\mathcal{A}} \delta a_1 + (1 - \delta)a_3\}$ are closed.

Definition 16 *Ambiguity Aversion.* For any acts a_1, a_2 and menu \mathcal{A} containing $\{a_1, a_2\}$, $a_1 \sim_{\mathcal{A}} a_2$ implies $\delta a_1 + (1 - \delta)a_2 \succsim_{\mathcal{A}} a_2$ for all $\delta \in (0, 1)$.

4.4 Axiomatic Characterization of Decision Rules

Consider an ordering over the actions in action set \mathcal{A} that satisfies mixture continuity and ambiguity aversion. Then:

- (i) An ordering satisfies completeness, von Neuman Morgenstern independence and independence of irrelevant alternatives if and only if it is Bayes' Rule (Fishburn, 1970).
- (ii) An ordering satisfies completeness, independence of irrelevant alternatives, and symmetry if and only if it is maxmin (Milnor, 1954).
- (iii) An ordering satisfies completeness, von Neuman Morgenstern independence, and symmetry if and only if it is minimax regret (Stoye, 2011*b*).

Symmetry, then, is what distinguishes Bayes rule from maxmin and minimax regret; Bayes rule does not satisfy symmetry, while the other two decision rules do satisfy symmetry. So, in order to better understand what we give up when we give up Bayesianism and turn to either maxmin, minimax regret, or any other decision rule that satisfies symmetry, we turn to a study of the symmetry axiom and what its introduction implies for agents' behaviors in an auction market.

4.5 Consequences of the Symmetry Axiom in k -Double Auctions

We begin with some general observations on applying the axiom of symmetry to a model of a double auction.

First, it is well-understood that the symmetry axiom excludes the use of priors. This has important implications for how a bidder will choose his bid in a double auction. Economists are accustomed to assuming that bidders will maximize their expected profit (or expected utility) according to their beliefs about the distribution of rival bids. If instead a bidder is unwilling or unable to calculate expected profit, then his decisions will be based on the various possible payoffs given a certain action, but not on how likely it is that each payoff will be realized.

Second, symmetry may offend our intuition about the relative “size” of events in a double auction. This is a direct result of the first point.

Third, it may (frequently) be the case that the bidders in an auction have enough knowledge of the *market environment* that symmetry is unreasonable. The relative “size” of events may be something that we expect bidders to have a basis for a subjective prior, even if they are not familiar with the details of a particular auction setting.

Fourth, we cannot deem symmetry to be “reasonable” or “unreasonable” for a particular *market institution*. The symmetry axiom was formulated for decision problems with *complete uncertainty*. Such a radical form of uncertainty may be rare, but when it holds, requiring symmetry is reasonable. Since we do not believe that the

institution necessarily implies anything about the bidder's knowledge or ignorance of the the proper description of the states of the world, we cannot make any assertions about the appropriateness of requiring symmetry in a particular market institution.

4.5.1 In a k -Double Auction, the Symmetry Axiom Results in Bid Invariance to Market Size

Suppose that a trader in a k -double auction selects his bid or ask using a decision rule that satisfies the symmetry axiom. We can prove that such a trader will choose the *same* action *regardless of the size of the market*.

In order to show that this is the case, we need to define events in such a way that swapping the consequences of disjoint constant-profit events will transform the profit function of a given bidding strategy to the profit function of the *same* bidding strategy given a different number of bidders. With this in mind, we define the state of the world so that it includes information applicable to two hypothetical markets of different sizes: one with m buyers and n sellers, and the other with m' buyers and n' sellers. Let the state of the world be

$$s = (\zeta_{(1)}, \zeta_{(2)}, \dots, \zeta_{(m+n-1)}, \xi_{(1)}, \xi_{(2)}, \dots, \xi_{(m'+n'-1)})$$

where $\zeta_{(1)} < \zeta_{(2)} < \dots < \zeta_{(m+n-1)}$ are the submitted bids and asks of $m - 1$ buyers and n sellers, and $\xi_{(1)} < \xi_{(2)} < \dots < \xi_{(m'+n'-1)}$ are the submitted bids and asks of $m' - 1$ buyers and n' sellers.³ Note that in a double auction with m buyers, what

³We do not consider how these bids and asks may be derived from bid and ask functions and underlying distributions of valuations and costs, because we are not concerned with finding an equilibrium.

matters to a bidder's payoff are the m^{th} and $m + 1^{\text{th}}$ highest bids/asks. We may therefore abbreviate the state of the world's description:

$$s = (\zeta_{(m)}, \zeta_{(m+1)}, \xi_{(m')}, \xi_{(m'+1)}).$$

We define $u_i(z)$ as the trader's utility from placing the bid b_i against the $m - 1$ other buyers and n sellers whose actions result in $\zeta_{(m)}$ and $\zeta_{(m+1)}$. We assume that this payoff is a strictly increasing function of the trader's profit.

Let \succsim be the bidder's preferences over a menu of payoff functions $\{u_i\}_{i \in M}$. Application of the symmetry axiom requires that for any menu of payoff functions $\{u'_i\}_{i \in \mathcal{A}}$ that can be derived from swapping the payoffs of constant-profit events, the same preferences must hold. We will transform the payoff functions in such a way that the new payoff functions $\{u'_i\}_{i \in \mathcal{A}}$ are the payoffs from the bidder's profit when he is bidding against the $m' + n' - 1$ rivals that place bids $(\xi_{(1)}, \xi_{(2)}, \dots, \xi_{(m'+n'-1)})$.

This set-up is not intended to model a situation in which the bidder does not know how many bidders he will face. It may be more helpful to imagine instead that the bidder has been placed in a room with a certain number of traders, and that another auction, being conducted in the next room, has a different number of traders. We will show that in such a scenario, his behavior would be the same no matter which room he was placed in.

Lemma 6 *For any $\zeta \in Z^{m+n-1}$, there exists $\xi \in Z^{m'+n'-1}$ such that $u(b, \zeta) = u(b, \xi)$ for any possible bid $b \in Z$.*

Theorem 7 *Suppose that a bidder's preferences over possible bidding strategies \succsim^S satisfies the symmetry axiom. Consider a menu of possible bids \mathcal{A} containing at least*

two possible actions, bids b_1 and b_2 . If $b_1 \succ_{\mathcal{A}}^S b_2$ when there are n sellers and $m - 1$ rival bidders, then $b_1 \succ_{\mathcal{A}}^S b_2$ when there are n' sellers and $m' - 1$ rival bidders. In fact, all of the preferences over actions in \mathcal{A} remain unchanged regardless of the number of traders in the market.

The proof is given in the appendix. The intuition behind this result is quite simple. Any bidding strategy that satisfies symmetry does *not* depend on the probability that the bid will win - something that changes with the number of bidders. It depends, instead, on extreme outcomes - for example, the best and/or the worst payoffs possible given his bid, or perhaps the maximum possible regret - something that does *not* change with the number of bidders. Therefore, it is clear that the number of bidders will not affect the bidding strategy.

4.5.2 Discussion

The fact that symmetry requires that bidder's strategies be invariant to the size of the market has important implications for the market's convergence to efficiency.

The proofs for convergence of double auctions to efficiency which assume that bidders play their Bayesian Nash Equilibrium strategies usually prove that the outcome converges to efficiency because when there are more traders in the market, the advantage of misrepresenting one's own true redemption value diminishes. Each trader's strategy converges to truthful reporting, so the outcome of the auction converges to the efficient outcome.

In contrast, if the market is populated by traders whose strategies are not based on these considerations, but instead are selected to satisfy symmetry, then increasing

the size of the market will not matter. As long as the “best case” and “worst case” outcomes remain unchanged, the asks and bids of such traders will remain the same. If they act as price-takers even in the bilateral case, then they will achieve efficiency in large markets as well. But increasing the size of the market will not bring their behavior closer to price-taking. This “negative result” concerning decision rules that satisfy the axiom of symmetry sharpens our understanding of the necessary conditions for convergence to price-taking behavior.

4.6 Conclusion

These results sharpen our understanding of the beliefs necessary for a market to converge to efficiency as the number of traders in the market grows. The theorem demonstrates that a trader with *unrestricted* beliefs about the distribution of rival bids and asks will not change his behavior even if he knows that the size of the market has changed. This reinforces the lessons learned in Chapters 2 of this dissertation. Unrestricted beliefs prevent convergence to price taking in private value double auctions. Unrestricted beliefs will result in invariant strategies in common value double auctions, as well.

Moreover, this insight into the nature of minimax regret and maxmin – and any other decision rule that satisfies the axiom of symmetry – has implications beyond the double auction. Not only in the double auction, but in any game for which the number of players does not change the “profile of outcomes”, such decision rules will choose strategies invariant to the number of players. This result could be applied to

the public contribution games, in which the size of the group is typically thought to affect the group's chances of funding a public good (Chamberlin, 1974).

Appendix A

Proofs for Chapter 2

Theorem 1 *In a k -double auction, the bid b_i that minimizes maximum regret for a buyer with private valuation v_i is $b_i = \frac{v_i}{1+k}$. The ask a_i that minimizes maximum regret for a seller with private cost c_i is $a_i = \frac{c_i+(1-k)}{1+(1-k)}$*

Proof: Buyer i 's profit is his valuation minus the price if he wins a unit of the good, and zero if he does not win. In the case that he wins, the price that he pays will be $k\zeta_{(m+1)} + (1-k)\zeta_{(m)}$ if his own bid is greater than $\zeta_{(m+1)}$, or $kb_i + (1-k)\zeta_{(m)}$ if his own bid is between $\zeta_{(m)}$ and $\zeta_{(m+1)}$:

$$\Pi_B = \begin{cases} v_i - (k\zeta_{(m+1)} + (1-k)\zeta_{(m)}) & \text{if } \zeta_{(m+1)} < b_i \\ v_i - (kb_i + (1-k)\zeta_{(m)}) & \text{if } \zeta_{(m)} < b_i < \zeta_{(m+1)} \\ 0 & \text{if } b_i < \zeta_{(m)} \end{cases} \quad (\text{A.1})$$

For any set of rival bids and offers ζ , the supremum of buyer i 's possible profit is

$$\Pi_B^* = \begin{cases} v_i - \zeta_{(m)} & \text{if } \zeta_{(m)} \leq v_i \\ 0 & \text{if } \zeta_{(m)} > v_i \end{cases} \quad (\text{A.2})$$

Then the buyer's regret function is

$$R_B = \begin{cases} k(\zeta_{(m+1)} - \zeta_{(m)}) & \text{if } \zeta_{(m)} \leq v_i \text{ and } \zeta_{(m+1)} < b_i \\ k(b_i - \zeta_{(m)}) & \text{if } \zeta_{(m)} \leq v_i \text{ and } \zeta_{(m)} < b_i < \zeta_{(m+1)} \\ v_i - \zeta_{(m)} & \text{if } \zeta_{(m)} \leq v_i \text{ and } b_i < \zeta_{(m)} \\ 0 & \text{if } \zeta_{(m)} > v_i \text{ and } b_i < \zeta_{(m)} \end{cases} \quad (\text{A.3})$$

In the first two cases, the buyer's regret is from winning at a higher price than necessary; in the third case, the buyer regrets failing to win a unit when the price is less than his valuation. The buyer's regret is zero if $\zeta_{(m)}$ is higher than his valuation. We omit the case that the buyer wins at a price higher than his own valuation, because the regret resulting from that action will always be at least as great as bidding v_i , and sometimes greater, so we eliminate the possibility of bidding more than v_i .

The supremum of the buyer's regret function is

$$\sup R_B = \max(kb_i, v_i - b_i, 0) \quad (\text{A.4})$$

The first term is increasing in buyer i 's bid b_i ; the second term is decreasing in b_i . (Each of the first two terms are greater than zero for all $b_i < v_i$.) The maximum regret is minimized when $kb_i = v_i - b_i$. Therefore, a buyer choosing his bid p_j to minimize this function will bid $\frac{v_i}{1+k}$.

The calculations for the seller's minimax regret ask are similar to the calculations

for the buyer.

$$\Pi_S = \begin{cases} (k\zeta_{(m+1)} + (1-k)\zeta_{(m)}) - c_i & \text{if } a_i < \zeta_{(m+1)} \\ (k\zeta_{(m+1)} + (1-k)a_i) - c_i & \text{if } \zeta_{(m)} < a_i < \zeta_{(m+1)} \\ 0 & \text{if } a_i > \zeta_{(m+1)} \end{cases} \quad (\text{A.5})$$

For any set of rival bids and offers ζ , the supremum of seller i 's possible profit is

$$\Pi_S^* = \begin{cases} \zeta_{(m+1)} - c_i & \text{if } c_i \leq \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < c_i \end{cases} \quad (\text{A.6})$$

Then the seller's regret function is

$$R_S = \begin{cases} (1-k)(\zeta_{(m+1)} - \zeta_{(m)}) & \text{if } \zeta_{(m+1)} \geq c_i \text{ and } a_i < \zeta_{(m+1)} \\ (1-k)(\zeta_{(m+1)} - a_i) & \text{if } \zeta_{(m+1)} \geq c_i \text{ and } \zeta_{(m)} < a_i < \zeta_{(m+1)} \\ \zeta_{(m+1)} - c_i & \text{if } \zeta_{(m+1)} \geq c_i \text{ and } a_i > \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < c_i \text{ and } a_i > \zeta_{(m+1)} \end{cases} \quad (\text{A.7})$$

The supremum of the seller's regret function is

$$\sup R_B = \max((1-k)(1-a_i), a_i - c_i, 0) \quad (\text{A.8})$$

The first term is decreasing in the seller's ask a_i ; the second term is increasing in a_i . (Each of the first two terms are greater than zero for all $a_i > c_i$.) The maximum

regret is minimized when $(1 - k)(1 - a_i) = a_i - c_i$. Therefore, a seller choosing his ask a_i to minimize this function will choose $a_i = \frac{c_i + (1-k)}{1 + (1-k)}$. ■

Theorem 2 *Let \tilde{F} denote the cumulative distribution function of the lowest cost among the n sellers in the market. The bid b_i that minimizes expected maximum regret for bidder i with valuation v_i satisfies*

$$\tilde{F}\left(\frac{(k+1)b_i - v_i}{k}\right) = \frac{\tilde{F}(b_i)}{1+k}$$

Such a bid b_i exists on the interval $\left[\frac{v_i}{1+k}, v_i\right]$.

Similarly, let \tilde{G} denote the cumulative distribution function of the highest valuation among the m sellers in the market. The ask a_i that minimizes expected maximum regret for seller i with cost c_i satisfies

$$\tilde{G}\left(\frac{a_i(1+(1-k)) - c_i}{1-k}\right) = \frac{\tilde{G}(a_i) + (1-k)}{1+(1-k)}$$

Such a bid a_i exists on the interval $\left[c_i, \frac{c_i + (1-k)}{1 + (1-k)}\right]$.

Proof: Each seller's ask is bounded below by his cost. Each bidder's bid is bounded above by his valuation. It follows that:

- The lower bound of $\zeta_{(m)}$ is the lowest realized cost, $c_{(1)}$. Since there are m buyers, the lowest m bids not including bidder i 's bid must include at least one seller. The $m - 1$ buyers can submit arbitrarily low bids. In every state of the world, it is possible for all of the rival bidders to submit bids of 0 - however "farfetched" that may seem. The lowest ask, however, cannot be less than $c_{(1)}$.
- The upper bound of $\zeta_{(m+1)}$ is 1. It is possible for $n - 1$ sellers to all submit asks of 1, since their asks are bounded above only by the highest price that any buyer could conceivably be willing to accept.

For this reason, we need only consider the distribution of the lowest realized cost, $c_{(1)}$, when we calculate the expected maximum regret¹. The maximum regret for a bid $b_i < v_i$ given the realization of costs c and values v is

$$R_B(b_i|v, c) = \begin{cases} 0 & \text{if } c_{(1)} > v_i \\ \max\{v_i - b_i, k(b_i - c_{(1)})\} & \text{if } v_i \geq b_i \geq c_{(1)} \\ v_i - c_{(1)} & \text{if } v_i \geq c_{(1)} > b_i \end{cases} \quad (\text{A.9})$$

In the second case, where the lowest realized cost $c_{(1)}$ is less than the bidder's bid b_i , the maximum regret could result from bidding more than necessary or from bidding less than necessary, depending on the value of $c_{(1)}$:

$$v_i - b_i < k(b_i - c_{(1)}) \quad (\text{A.10})$$

$$\Rightarrow c_{(1)} < \frac{(k+1)b_i - v_i}{k} \quad (\text{A.11})$$

Since costs of the sellers are independently and identically distributed with cdf F , the smallest cost realized by n sellers is has cdf $\tilde{F}(c) = 1 - (1 - F(c))^n$. Therefore,

$$\bar{R}_B(b_i|v_i) = \int_0^{b^*} k(b_i - c)d\tilde{F}(c) + \int_{b^*}^{b_i} (v_i - b_i)d\tilde{F}(c) + \int_{b_i}^{v_i} (v_i - c)d\tilde{F}(c) \quad (\text{A.12})$$

Where $b^* = \frac{(k+1)b_i - v_i}{k}$. If $b_i \leq \frac{v_i}{1+k}$, then the first term disappears:

$$\bar{R}_B(b_i|v_i) = \int_0^{b_i} (v_i - b_i)d\tilde{F}(c) + \int_{b_i}^{v_i} (v_i - c)d\tilde{F}(c) \quad (\text{A.13})$$

¹The following discussion follows Linhardt and Radner (1989) closely.

Integrating by parts,

$$\bar{R}_B(b_i|v_i) = \begin{cases} k \int_0^{\frac{(k+1)b_i - v_i}{k}} \tilde{F}(c)dc + \int_{b_i}^{v_i} \tilde{F}(c)dc & \text{if } b_i > \frac{v_i}{1+k} \\ \int_{b_i}^{v_i} \tilde{F}(c)dc & \text{if } b_i \leq \frac{v_i}{1+k} \end{cases} \quad (\text{A.14})$$

Differentiating the expected regret function with respect to bidder i 's bid b_i ,

$$\frac{d}{db_i} \bar{R}_B = \begin{cases} (k+1)\tilde{F}\left(\frac{(k+1)b_i - v_i}{k}\right) - \tilde{F}(b_i) & \text{if } \frac{v_i}{1+k} \leq b_i \leq v_i \\ -\tilde{F}(b_i) & \text{if } 0 \leq b_i \leq \frac{v_i}{1+k} \end{cases} \quad (\text{A.15})$$

This derivative is continuous, and it is non-positive for $b_i = \frac{v_i}{1+k}$ and non-negative for $b_i = v_i$. Then there exists a bid b_i such that the derivative of the expected regret function is zero, where the expected maximum regret is minimized.

The maximum regret for a bid $a_i > c_i$ given the realization of costs c and values v is

$$R_S(a_i|v, c) = \begin{cases} 0 & \text{if } v_{(m)} < c_i \\ \max\{a_i - c_i, (1-k)(v_{(m)} - a_i)\} & \text{if } a_i \geq v_{(m)} \\ v_{(m)} - c_i & \text{if } c_i \geq v_{(m)} < a_i \end{cases} \quad (\text{A.16})$$

In the second case, where the highest realized valuation $v_{(m)}$ is less than the seller's ask a_i , the maximum regret could result from asking more than necessary or from

asking less than necessary, depending on the value of $v_{(m)}$:

$$a_i - c_i < (1 - k)(v_{(m)} - C_i) \quad (\text{A.17})$$

$$\Rightarrow v_{(m)} > \frac{(1 + (1 - k))a_i - c_i}{1 - k} \quad (\text{A.18})$$

Since costs of the buyers are independently and identically distributed with cdf G , the highest valuation realized by m sellers is has cdf $\tilde{G}(v) = G(v)^n$. Therefore,

$$\bar{R}_S(a_i|c_i) = \int_{c_i}^{a_i} (v - c_i)d\tilde{G}(v) + \int_{a_i}^{C^*} (a_i - c_i)d\tilde{G}(v) + \int_{C^*}^1 (1 - k)(v - a_i)d\tilde{G}(v) \quad (\text{A.19})$$

Where $C^* = \frac{(1+(1-k))a_i - c_i}{1-k}$. If $a_i \geq \frac{a_i + (1-k)}{1+(1-k)}$, then the last term disappears:

$$\bar{R}_S(a_i|c_i) = \int_{c_i}^{a_i} (v - c_i)d\tilde{G}(v) + \int_{a_i}^1 (a_i - c_i)d\tilde{G}(v) \quad (\text{A.20})$$

Integrating by parts,

$$\bar{R}_S(a_i|c_i) = \begin{cases} (1 - k)(1 - a_i) - \int_{c_i}^{a_i} \tilde{G}(v)dv - \int_{C^*}^1 (1 - k)\tilde{G}(v)dv & \text{if } a_i > \frac{c_i + (1-k)}{1+(1-k)} \\ (a_i - c_i) - \int_{c_i}^{a_i} \tilde{G}(v)dv & \text{if } a_i < \frac{c_i + (1-k)}{1+(1-k)} \end{cases} \quad (\text{A.21})$$

Differentiating the expected regret function with respect to seller i 's asm a_i ,

$$\frac{d\bar{R}_S(a_i|c_i)}{da_i} = \begin{cases} -(1-k) - \tilde{G}(a_i) + (2-k)\tilde{G}\left(\frac{(1+(1-k))a_i - c_i}{1-k}\right) & \text{if } a_i > \frac{c_i + (1-k)}{1+(1-k)} \\ 1 - \tilde{G}(c_i) & \text{if } a_i < \frac{c_i + (1-k)}{1+(1-k)} \end{cases} \quad (\text{A.22})$$

This derivative is continuous, and it is non-negative for $a_i = \frac{c_i + (1-k)}{1+(1-k)}$ and non-positive for $a_i = c_i$. Then there exists a bid a_i such that the derivative of the expected regret function is zero, where the expected maximum regret is minimized. \blacksquare

Corollary 1 *Let $b(v; n)$ denote the bid that minimizes expected maximum regret in a k -double auction with n sellers. Then $\lim_{n \rightarrow \infty} b(v; n) = \frac{v}{1+k}$.*

Proof: Suppose not. Then there exists some valuation v and $\epsilon > 0$ such that

$$\lim_{n \rightarrow \infty} b(v; n) > \frac{v}{1+k} + \epsilon$$

Since the derivative of expected maximum regret at $b(v; n)$ is zero, and since the derivative of the expected maximum regret is strictly increasing, $\frac{v}{1+k} + \epsilon < b(v; n)$ implies that the derivative of expected maximum regret at $\frac{v}{1+k} + \epsilon$ is less than zero in the limit:

$$\lim_{n \rightarrow \infty} \left[(k+1)\tilde{F}\left(\frac{(k+1)\left(\frac{v}{1+k} + \epsilon\right) - v}{k}\right) - \tilde{F}\left(\frac{v}{1+k} + \epsilon\right) \right] < 0 \quad (\text{A.23})$$

$$(k+1) \lim_{n \rightarrow \infty} \tilde{F}\left(\frac{(k+1)\epsilon}{k}\right) - \lim_{n \rightarrow \infty} \tilde{F}\left(\frac{(k+1)\epsilon}{k+1}\right) < 0 \quad (\text{A.24})$$

$$(k+1)(1) - 1 < 0 \quad (\text{A.25})$$

This is a contradiction. ■

Theorem 3 *Suppose that the following conditions hold for $\Gamma = \{\Gamma_i\}_{i=1}^\infty$:*

1. *For each sequence of priors $\{G_{\gamma,i}\} \in \{\Gamma_i\}_{i=1}^\infty$, for every $\epsilon \in (0, v)$, there exists $N(\epsilon, G_\gamma) \in \mathbb{N}$ such that for all $i \geq N(b, G_\gamma)$:*

$$\begin{aligned} \int u(v - \epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ > \int u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \end{aligned} \quad (\text{A.26})$$

for all $b' < v - \epsilon$. That is, under each prior $\{G_\gamma\}_{i=1}^\infty \in \{\Gamma_i\}_{i=1}^\infty$, the utility-maximizing bid converges to v over the sequence of markets.

2. *There exists a well-defined function*

$$\bar{N}(\epsilon) = \max_{G_\gamma \in \Gamma} \{N(\epsilon, G_\gamma)\} \quad (\text{A.27})$$

Then the bid that minimizes maximum expected regret under $\{\Gamma_i\}_{i=1}^\infty$ converges to truthful bidding over the sequence of markets $\{(m_i, n_i)\}_{i=1}^\infty$.

Proof: (A.27) implies that for $i \geq \bar{N}(\epsilon)$,

$$\begin{aligned} \int \max_{b^* \in [0, v]} u(b^*, \zeta) - u(v - \epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ < \int \max_{b^* \in [0, v]} u(b^*, \zeta) - u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v - \epsilon \end{aligned} \quad (\text{A.28})$$

holds for each $G_\gamma \in \Gamma$. In other words, for markets subsequent to $(m_{\bar{N}(\epsilon)}, n_{\bar{N}(\epsilon)})$, expected regret is minimized at a bid within ϵ of truthful bidding, for each prior $G_\gamma \in \Gamma$.

$$\begin{aligned} & \int R(v - \epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ & < \int R(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v - \epsilon \quad (\text{A.29}) \end{aligned}$$

Therefore, for $i > \bar{N}(\epsilon)$,

$$\begin{aligned} & \max_{G_{\gamma,i} \in \Gamma_i} \int R(v - \epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ & < \max_{G_{\gamma,i} \in \Gamma_i} \int R(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v - \epsilon \quad (\text{A.30}) \end{aligned}$$

For suppose not. Then there would exist $G_* \in \Gamma$ such that for some $j \geq \bar{N}(\epsilon)$, for some $b' < v - \epsilon$,

$$\begin{aligned} & \int R(v - \epsilon, \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ & > \max_{G_{\gamma,j} \in \Gamma_j} \int R(b', \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{\gamma,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \quad (\text{A.31}) \end{aligned}$$

implying

$$\begin{aligned} & \int R(v - \epsilon, \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ & > \int R(b', \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \quad (\text{A.32}) \end{aligned}$$

which contradicts (A.29).

We conclude that the maximum expected regret-minimizing bid can be arbitrarily close to bidding v , given a sufficient number of competitors. As equation A.30 establishes, for a number of bidders greater than or equal to $\bar{N}(\epsilon)$, the optimal bid must be within ϵ of bidding one's true valuation. ■

Lemma 1 *Let $\{(m_i, n_i)\}_{i=1}^{\infty}$ be a sequence of markets in which the m_i buyers approaches infinity. Let each $G_\gamma = \{G_{\gamma,i}\}_{i=1}^{\infty}$ in Γ be a joint distribution of the m_i^{th} and $m_i + 1^{\text{th}}$ order statistics in which all bids and asks are treated as $(m_i + n_i - 1)$ iid draws from a distribution $f_\gamma(x)$, where $f_\gamma(x) > \epsilon > 0$.*

Then the bid that minimizes the maximum expected regret will approach truthful bidding as $i \rightarrow \infty$.

Proof:

If the rival bids and asks in a market of size (m_i, n_i) are independently and identically distributed according to some distribution with cdf F and pdf f , then the joint distribution of the m_i^{th} and $m_i + 1^{\text{th}}$ order statistics is

$$f_{(m_i)(m_i+1)}(x, y) = \frac{(m_i + n_i - 1)!}{(m_i - 1)!n_i!} (x)f(x)f(y)F^{m_i-1}[1 - F(y)]^{n_i} \quad (\text{A.33})$$

The subscripts i denoting the market will be omitted in the following proof. For brevity, let $\frac{(m+n-1)!}{(m-1)!n!} = C_{m,m+1}$.

We can calculate the expected regret of a bid b under the prior f :

$$\begin{aligned}
E[R_B] &= \int_0^b \int_x^b k(y-x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1-F(y)]^n dy dx \\
&\quad + \int_0^b \int_b^1 k(b-x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1-F(y)]^n dy dx \\
&\quad + \int_b^v \int_x^1 (v-x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1-F(y)]^n dy dx
\end{aligned} \tag{A.34}$$

$$\begin{aligned}
\frac{dE[R_B]}{db} &= \int_0^b k(b-x)C_{m,m+1}f(x)f(b)F(x)^{m-1}[1-F(b)]^n dx \\
&\quad - \int_0^b k(b-x)C_{m,m+1}f(x)f(b)F(x)^{m-1}[1-F(b)]^n dx \\
&\quad + \int_0^b \int_b^1 kC_{m,m+1}f(x)f(y)F(x)^{m-1}[1-F(y)]^n dy dx \\
&\quad - \int_b^1 (v-b)C_{m,m+1}f(b)f(y)F(b)^{m-1}[1-F(y)]^n dy
\end{aligned} \tag{A.35}$$

$$\begin{aligned}
\frac{dE[R_B]}{db} &= \int_0^b \int_b^1 kC_{m,m+1}f(x)f(y)F(x)^{m-1}[1-F(y)]^n dy dx \\
&\quad - \int_b^1 (v-b)C_{m,m+1}f(b)f(y)F(b)^{m-1}[1-F(y)]^n dy
\end{aligned} \tag{A.36}$$

Then the first-order condition to minimize expected regret is:

$$\int_0^b \int_b^1 k f(x) f(y) F^{m-1}(x) [1 - F(y)]^n dy dx = \int_b^1 (v - b) f(b) f(y) F^{m-1}(b) [1 - F(y)]^n dy \quad (\text{A.37})$$

$$k \int_0^b f(x) F(x)^{m-1} dx = (v - b) f(b) F(b)^{m-1} \quad (\text{A.38})$$

$$k \frac{F(b)^m}{m} = (v - b) f(b) F(b)^{m-1} \quad (\text{A.39})$$

$$\frac{k F(b)}{m f(b)} = (v - b) \quad (\text{A.40})$$

Therefore, **if** $f(b) > 0$, the amount that the bidder shades her bid will converge to zero as the number of other bidders grows large. ■

Appendix B

Proofs for Chapter 3

Theorem 4 *In a common value Vickrey auction, the minimax regret bid is $b_i = s_i$.*

Proof: The bidder's profit function in a Vickrey auction is:

$$\pi = \begin{cases} \omega - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.1})$$

The greatest profit that a bidder can achieve, given the state ω , is:

$$\pi^* = \begin{cases} \omega - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.2})$$

Therefore, the bidder's regret function is:

$$R(b, \omega) = \begin{cases} \omega - \max_{j \neq i} b_j - (\omega - \max_{j \neq i} b_j) & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{otherwise} \end{cases} \quad (\text{B.3})$$

Maximum regret for a given bid b_i is therefore

$$\sup R(b_i, s_i) = \max\{b_i - s_i + \epsilon, s_i - b_i + \epsilon, 0\} \quad (\text{B.4})$$

To minimize the maximum regret, therefore, the bid b_i should satisfy the condition

$$b_i - s_i + \epsilon = s_i - b_i + \epsilon \quad (\text{B.5})$$

This condition is met when $b_i = s_i$. ■

Lemma 2 *In a common value first price sealed bid auction, a bidder who observed signal s_i has the following maximum regret function:*

$$\sup R = \max\{b_i - \inf\{\max_{j \neq i} b_j\}, b_i - s_i + \epsilon, s_i + \epsilon - b_i\} \quad (\text{B.6})$$

Proof: In the first price auction, the bidder's profit function is

$$\pi = \begin{cases} \omega - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i \leq \max_{j \neq i} b_j \end{cases} \quad (\text{B.7})$$

The greatest profit that a bidder can achieve, given the state ω , is:

$$\pi^* = \begin{cases} \omega - \max_{j \neq i} b_j & \text{if } \omega > \max_{j \neq i} b_j \\ 0 & \text{if } \omega < \max_{j \neq i} b_j, b_i < \max_{j \neq i} b_j \end{cases} \quad (\text{B.8})$$

Therefore, the regret function is

$$R(b, \omega) = \begin{cases} b_i - \max_{j \neq i} b_j & \text{if } \omega > \max_{j \neq i} b_j, b_i > \max_{j \neq i} b_j \\ b_i - \omega & \text{if } \omega < \max_{j \neq i} b_j, b_i > \max_{j \neq i} b_j \\ \omega - \max_{j \neq i} b_j & \text{if } \omega > \max_{j \neq i} b_j, b_i < \max_{j \neq i} b_j \\ 0 & \text{if } \omega < \max_{j \neq i} b_j, b_i < \max_{j \neq i} b_j \end{cases} \quad (\text{B.9})$$

In the first case, there is opportunity to win the asset at a profit, and the bidder does win the asset; any amount that the winning bidder bids above the highest rival bid is regretted. This regret is greatest, then, when the highest rival bid turns out to be at the lowest possible value. Call this lowest possible value for the highest rival bid, $\inf\{\max_{j \neq i} b_j\}$. Note that if the decision maker's signal conveys information about other bidder's signals, then the range of possible rival bids that the bidder considers could depend on his own signal.

$$\sup\{b_i - \max_{j \neq i} b_j\} = b_i - \inf\{\max_{j \neq i} b_j\} \quad (\text{B.10})$$

In the second case, the true value of the asset is less than the bid required to win it; therefore, the bidder regrets winning at an unprofitable price. This regret is greatest when the asset ω is at its lowest possible value, given bidder's signal: $\omega = s_i - \epsilon$.

$$\sup\{b_i - \omega\} = b_i - (s_i - \epsilon) \quad (\text{B.11})$$

In the third case, there is opportunity to win the asset at a profit, but the bidder does not win the asset; the losing bidder regrets the foregone profit. This regret is greatest when the asset is at its highest possible value (given the bidder's signal, that is, $\omega = s_i + \epsilon$) and the rival bid has just barely edged out the bid b_i .

$$\sup\{\omega - \max_{j \neq i} b_j\} = s_i + \epsilon - \inf_{\max_{j \neq i} b_j > b_i} \{\max_{j \neq i} b_j\} \quad (\text{B.12})$$

$$= s_i + \epsilon - b_i \quad (\text{B.13})$$

Therefore, the supremum of the regret function given the bidder's signal s_i and the maximum possible error ϵ is

$$\sup R = \max\{b_i - \inf_{j \neq i} \{\max_{j \neq i} b_j\}, b_i - s_i + \epsilon, s_i + \epsilon - b_i\} \quad (\text{B.14})$$

■

Theorem 5 *If each bidder believes that the other bidders will reduce their bid from their own signal by the maximum possible error,*

$$b(s) = s - \epsilon, \quad (\text{B.15})$$

Then this same bidding function will minimize the bidder's own maximum regret. Thus, bidders who reduce their bids from their signals by the maximum error ϵ will mutually minimize maximum regret.

Proof: If all of the other bidders bid $b_j = s_j - \epsilon$, then

$$\inf\{\max_{j \neq i} b_j\} = \inf_{s_j \in (s_i - 2\epsilon, s_i + 2\epsilon)} \{s_j - \epsilon\} \quad (\text{B.16})$$

$$= s_i - 3\epsilon \quad (\text{B.17})$$

Therefore, the maximum regret for a bidder's bid b_i is

$$\sup R = \max\{b_i - s_i + 3\epsilon, b_i - s_i + \epsilon, s_i + \epsilon - b_i\} \quad (\text{B.18})$$

$$= \max\{b_i - s_i + 3\epsilon, s_i + \epsilon - b_i\} \quad (\text{B.19})$$

Note that maximum regret is increasing in $b_i - s_i + 3\epsilon$, and decreasing in $s_i + \epsilon - b_i$. Maximum regret will be minimized when each scenario obtains the same level of regret:

$$b_i - s_i + 3\epsilon = s_i + \epsilon - b_i \quad (\text{B.20})$$

$$2b_i = 2s_i - 2\epsilon \quad (\text{B.21})$$

$$b_i = s_i - \epsilon \quad (\text{B.22})$$

■

Lemma 4 Let $\check{\zeta}_{(m)}$ denote the lowest possible value of $\zeta_{(m)}$ given the signal s_i ob-

served by the bidder. Then the minimax regret bid is

$$b = \begin{cases} s & \text{if } s \leq \frac{\epsilon}{k} + \check{\zeta}_{(m)} \\ \frac{s + \epsilon + k\check{\zeta}_{(m)}}{1+k} & \text{if } s > \frac{\epsilon}{k} + \check{\zeta}_{(m)} \end{cases} \quad (\text{B.23})$$

Proof: The calculation for the profit function and the regret function in this model is similar to the calculations for those functions in the case of interdependent values. The buyer's regret function is

$$R_B = \begin{cases} k(\check{\zeta}_{(m+1)} - \check{\zeta}_{(m)}) & \text{if } \check{\zeta}_{(m)} \leq \omega, \check{\zeta}_{(m+1)} < b_i \\ k(b_i - \check{\zeta}_{(m)}) & \text{if } \check{\zeta}_{(m)} \leq \omega, \check{\zeta}_{(m)} < b_i < \check{\zeta}_{(m+1)} \\ \omega - \check{\zeta}_{(m)} & \text{if } \check{\zeta}_{(m)} \leq \omega, b_i < \check{\zeta}_{(m)} \\ (kb_i + (1-k)\check{\zeta}_{(m)}) - \omega & \text{if } \check{\zeta}_{(m)} > \omega, \check{\zeta}_{(m)} < b_i < \check{\zeta}_{(m+1)} \\ (k\check{\zeta}_{(m+1)} + (1-k)\check{\zeta}_{(m)}) - \omega & \text{if } \check{\zeta}_{(m)} > \omega, \check{\zeta}_{(m+1)} < b_i \\ 0 & \text{if } \check{\zeta}_{(m)} > \omega, b_i < \check{\zeta}_{(m)} \end{cases} \quad (\text{B.24})$$

In the first two cases, the bidder's regret comes from bidding more than necessary in order to win an asset, although the asset's value of ω is worth buying at the minimum price $\check{\zeta}_{(m)}$. Note that this regret is greatest when $\check{\zeta}_{(m+1)}$ is as great as possible, and $\check{\zeta}_{(m)}$ is as low as possible, ie, $\check{\zeta}_{(m)}$.

In the third case, the buyer regrets failing to purchase a unit of the asset when it would be profitable to do so at the price $\check{\zeta}_{(m)}$. This regret is maximized when the asset's true value is as high as possible (given the signal observed by the buyer), and

the price is as low as possible (given that the buyer was not able to secure one with his bid b_i).

In the fourth and fifth case, the bidder regrets paying more for the item than it is worth. This regret is maximized when the asset's true value is as low as possible (given the buyer's signal s_i).

In the last case, the bidder has no regret from not purchasing an asset, because the required bid is too high.

Thus, the maximum regret function is

$$\sup R_B = \max\{k(b_i - \check{\zeta}_{(m)}), s_i + \epsilon - b_i, b_i - s_i + \epsilon\} \quad (\text{B.25})$$

If $k(b_i - \check{\zeta}_{(m)}) > b_i - s_i + \epsilon$, then maximum regret is minimized at

$$k(b_i - \check{\zeta}_{(m)}) = s_i + \epsilon - b_i \quad (\text{B.26})$$

$$(1 + k)b_i = s_i + \epsilon + k\check{\zeta}_{(m)} \quad (\text{B.27})$$

$$b_i = \frac{s_i + \epsilon + k\check{\zeta}_{(m)}}{1 + k} \quad (\text{B.28})$$

If $b_i - s_i + \epsilon > k(b_i - \check{\zeta}_{(m)})$, then maximum regret is minimized at

$$b_i - s_i + \epsilon = s_i + \epsilon - b_i \quad (\text{B.29})$$

$$2b_i = 2s_i \quad (\text{B.30})$$

$$b_i = s_i \quad (\text{B.31})$$

■

Lemma 5 Let $\hat{\zeta}_{(m+1)}$ denote the highest possible value of $\zeta_{(m+1)}$ given the signal s_i observed by the seller. Then the minimax regret ask is

$$a = \begin{cases} s & \text{if } s \geq \hat{\zeta}_{(m+1)} - \frac{\epsilon}{1-k} \\ \frac{s - \epsilon + (1-k)\hat{\zeta}_{(m+1)}}{1+(1-k)} & \text{if } s < \hat{\zeta}_{(m+1)} - \frac{\epsilon}{1-k} + \check{\zeta}_{(m)} \end{cases} \quad (\text{B.32})$$

Proof: The seller's regret function is

$$R_S = \begin{cases} (1-k)(\check{\zeta}_{(m+1)} - \check{\zeta}_{(m)}) & \text{if } \check{\zeta}_{(m+1)} \geq \omega, a_i < \check{\zeta}_{(m)} \\ k(\check{\zeta}_{(m+1)} - a_i) & \text{if } \check{\zeta}_{(m+1)} \geq \omega, \check{\zeta}_{(m)} < a_i < \check{\zeta}_{(m+1)} \\ \check{\zeta}_{(m+1)} - \omega & \text{if } \check{\zeta}_{(m+1)} \geq \omega, \check{\zeta}_{(m+1)} < a_i \\ \omega + (ka_i + (1-k)\check{\zeta}_{(m)}) & \text{if } \check{\zeta}_{(m+1)} < \omega, \check{\zeta}_{(m)} < a_i < \check{\zeta}_{(m+1)} \\ \omega + (k\check{\zeta}_{(m+1)} + (1-k)a_i) & \text{if } \check{\zeta}_{(m+1)} < \omega, a_i < \check{\zeta}_{(m)} \\ 0 & \text{if } \check{\zeta}_{(m+1)} < \omega, \check{\zeta}_{(m+1)} < a_i \end{cases} \quad (\text{B.33})$$

The supremum of the regret function is

$$\sup R_S = \max\{(1-k)(\hat{\zeta}_{(m+1)} - a), a - s + \epsilon, s + \epsilon - a\} \quad (\text{B.34})$$

If $(1 - k)(\hat{\zeta}_{(m+1)} - s) > \epsilon$, then maximum regret is minimized at

$$(1 - k)(\hat{\zeta}_{(m+1)} - a) = a - s + \epsilon \quad (\text{B.35})$$

$$a = \frac{s - \epsilon + (1 - k)\hat{\zeta}_{(m+1)}}{1 + (1 - k)} \quad (\text{B.36})$$

If $(1 - k)(\hat{\zeta}_{(m+1)} - s) < \epsilon$, then maximum regret is minimized at

$$a - s + \epsilon = s + \epsilon - a \quad (\text{B.37})$$

$$a = s \quad (\text{B.38})$$

■

Theorem 6 *The following bid and asks functions mutually minimize maximum regret.*

1. If $k = 0$: $b = s$, and $a = s + \frac{\epsilon}{2}$.
2. If $k = 1$: $b = s - \frac{\epsilon}{2}$, and $a = s$.
3. If $k = \frac{1}{2}$: $b = s$, and $a = s$.

Proof: Whenever the bid function is less than or equal to the ask function at each signal s_i , the following statements hold:

- For each bidder i , the value of $\inf \zeta_{(m)}$ depends on the lowest possible ask from a *seller*. This lowest possible ask is determined by the ask function and the lowest possible signal that the seller could observe, given buyer i 's signal s_i :

since the lowest possible state is $s_i - \epsilon$, the lowest signal that a seller could observe is $s_i - 2\epsilon$. If the buyer believes that sellers will bid truthfully, then $\inf \zeta_{(m)} = s_i - 2\epsilon$.

- Similarly, for each seller i , the value of $\sup \zeta_{(m+1)}$ depends on the highest possible bid from a *buyer*. This highest possible bid is determined by the bid function and the lowest possible signal that the buyer could observe, given seller i 's signal s_i : since the highest possible state is $s_i + \epsilon$, the highest signal that a buyer could observe is $s_i + 2\epsilon$. If the seller believes that the buyers are bidding truthfully, then $\sup \zeta_{(m+1)} = s_i + 2\epsilon$
1. If $k = 0$, then buyers have no influence on the price. Since the bidder faces no regret from being a pivotal trader, he has no reason to shade his bid:

$$\sup R_B = \max\{(0)(b_i - \check{\zeta}_{(m)}), s_i + \epsilon - b_i, b_i - s_i + \epsilon\} \quad (\text{B.39})$$

The bid that minimizes this maximum regret function is $b_i = s_i$. On the other hand, the seller's maximum regret function is:

$$\sup R_S = \max\{s + 2\epsilon - a, a - s + \epsilon, s + \epsilon - a\} \quad (\text{B.40})$$

because the seller knows that the bidders will bid their signals. The ask that minimizes this maximum regret function is $a_i = s_i - \frac{\epsilon}{2}$.

2. The proof for $k = 1$ is similar, with the seller's and buyer's situations reversed.

3. If $k = \frac{1}{2}$, we can verify that truthful bids and asks satisfy the conditions in the two previous lemmas:

$$b = \frac{s + \epsilon + \frac{1}{2}(s - 2\epsilon)}{\frac{3}{2}} = s \quad (\text{B.41})$$

and

$$a = \frac{s - \epsilon + \frac{1}{2}(s + 2\epsilon)}{\frac{3}{2}} = s \quad (\text{B.42})$$

■

Appendix C

Proofs for Chapter 4

Lemma 6 For any $\zeta \in Z^{m+n-1}$, there exists $\xi \in Z^{m'+n'-1}$ such that $u(b, \zeta) = u(b, \xi)$ for any possible bid $b \in Z$.

Proof: $u(b, \zeta)$ is determined by the relationship between $\zeta_{(m)}, \zeta_{(m+1)}$, and b . There exists $\xi \in Z^{m'+n'-1}$ such that $\xi_{m'} = \zeta_{(m)}$ and $\xi_{(m'+1)} = \zeta_{(m+1)}$. For example, define ξ as

$$\xi_{(i)} = \begin{cases} \min\{Z\} & \text{for } i < m' \\ \zeta_{(m)} & \text{for } i = m' \\ \zeta_{(m+1)} & \text{for } i = m' + 1 \\ \max\{Z\} & \text{for } i > m' + 1 \end{cases} \quad (\text{C.1})$$

A conscientious reader may ask whether it will be possible for a double auction with $m' - 1$ rival bidders and n' sellers to submit *rationalizable* bids and asks that satisfy this criterion. There may be some realizations of the other traders' valuations and costs such that the above actions are not rationalizable. However, since it is within the realm of possibility for all buyers to have valuations equal to \bar{v} , and for all sellers

to have costs equal to \underline{c} , any set of bids and asks $\zeta \in Z^{m+n-1}$ or $\xi \in Z^{m'+n'-1}$ could be rationalizable.

For any bid b_i , the payoff $u_i(\xi)$ will be

$$u_i(\xi) = \begin{cases} v - [(1-k)\xi_{(m)} + k\xi_{(m+1)}] & \text{if } \zeta_{(m+1)} = \xi_{(m+1)} < b_i \\ v - [(1-k)\xi_{(m)} + kb_i] & \text{if } \zeta_{(m)} = \xi_{(m)} < b_i < \zeta_{(m+1)} = \xi_{(m+1)} \\ 0 & \text{if } b_i < \zeta_{(m)} = \xi_{(m)} \end{cases} \quad (\text{C.2})$$

exactly equal to the payoff $u_i(b_i, \zeta)$.

Another way of expressing this result is that the set U_i of possible payoffs of bid b_i is the same for all m and n . ■

Theorem 7 *Suppose that a bidder's preferences over possible bidding strategies \succsim^S satisfies the symmetry axiom. Consider a menu of possible bids \mathcal{A} containing at least two possible actions, bids b_1 and b_2 . If $b_1 \succsim_{\mathcal{A}}^S b_2$ when there are n sellers and $m-1$ rival bidders, then $b_1 \succsim_{\mathcal{A}}^S b_2$ when there are n' sellers and $m'-1$ rival bidders. In fact, all of the preferences over actions in \mathcal{A} remain unchanged regardless of the number of traders in the market.*

Proof: The idea is to show that there exists ψ such that swapping constant profit events can transform the profit function of bidding b_i against $m+n-1$ rivals into the profit function of bidding b_i against $m'+n'-1$ rivals. This function ψ can be defined for each state z in the following way:

- For state (ζ, ξ) , calculate the payoff from bids b_1 and b_2 , when there are n sellers and $m-1$ rival bidders:

$$u_1(\zeta, \xi) = \Pi(b_1, \zeta) \text{ and } u_2(\zeta, \xi) = \Pi(b_2, \zeta).$$

- Calculate also the payoff from bids b_1 and b_2 , when there are n' sellers and $m' - 1$ rival bidders:

$$u'_1(\zeta, \xi) = \Pi(b_1, \xi) \text{ and } u'_2(\zeta, \xi) = \Pi(b_2, \xi).$$

- Applying proposition 1, find a state of the world $(\tilde{\zeta}, \tilde{\xi})$ such that $u_i(\tilde{\zeta}, \tilde{\xi}) = u'_i(\zeta, \xi)$ and $u_i(\zeta, \xi) = u'_i(\tilde{\zeta}, \tilde{\xi})$ for every possible bid.

That is, $\Pi(b_i, \tilde{\zeta}) = \Pi(b_i, \xi)$ and $\Pi(b_i, \xi) = \Pi(b_i, \tilde{\zeta})$ for every permissible bid b_i .

- Set $\psi((\zeta, \xi)) = (\tilde{\zeta}, \tilde{\xi})$. Then $u'_i(\zeta, \xi) = u_i(\psi(\zeta, \xi)) = u_i(\tilde{\zeta}, \tilde{\xi})$, for all i . That is, $\Pi(b_i, \xi) = \Pi(b_i, \tilde{\zeta})$ for all i .

Note that this function ψ preserves the profile of outcomes. If $\psi((\zeta, \xi)) = \psi((\hat{\zeta}, \hat{\xi})) = (\tilde{\zeta}, \tilde{\xi})$, then $u_i(\tilde{\zeta}, \tilde{\xi}) = u'_i(\zeta, \xi) = u'_i(\hat{\zeta}, \hat{\xi})$.

From symmetry, we have that $b_1 \succ^S b_2$ in the original decision problem if and only if $b_1 \succ^S b_2$ in the transformed decision problem. ■

References

- Anscombe, Frank J., and Robert J. Aumann.** 1963. “A Definition of Subjective Probability.” *The Annals of Mathematical Statistics*, 34(1): 199–205.
- Arrow, Kenneth. J., and Leonid Hurwicz.** 1972. *An Optimality Criterion for Decision-making Under Complete Ignorance*. A.M. Kelley.
- Blin, Jean-Marie.** 1976. “How Relevant are ‘Irrelevant’ Alternatives?” *Theory and Decision*, 7: 95–105.
- Cason, Timothy N., and Daniel Friedman.** 1996. “Price formation in double auction markets.” *Journal of Economic Dynamics and Control*, 20(8): 1307–1337.
- Chamberlin, John.** 1974. “Provision of Collective Goods as a Function of Group Size.” *American Political Science Review*, 68: 707–716.
- Chatterjee, Kalyan, and William Samuelson.** 1983. “Bargaining Under Incomplete Information.” *Operations Research*, 31: 835–851.
- Chernoff, Herman.** 1954. *Rational Selection of Decision Functions*. Vol. 22 of *Cowles Commission Papers: New Series*, Cowles Commission for Research in Economics.
- Chiesa, Alessandro, Silvio Micali, and Zeyuan A. Zhu.** 2014. “Knightian self-uncertainty in the VCG mechanism for unrestricted combinatorial auctions.”
- Cohen, Michele, and Jean-Yves Jaffray.** 1980. “Rational Behavior under Complete Ignorance.” *Econometrica*, 48(5): 1281–99.
- Cripps, Martin W., and Jeroen M. Swinkels.** 2006. “Efficiency of Large Double Auctions.” *Econometrica*, 74(1): 47–92.
- Dempster, Arthur P.** 1975. “A subjectivist look at robustness.” *Bulletin of International Statistical Institute*, 46: 349–374.

- Fishburn, Peter C.** 1970. *Utility Theory for Decision Making*. Wiley, New York.
- Friedman, Daniel.** 1991. “A simple testable model of double auction markets.” *Journal of Economic Behavior & Organization*, 15(1): 47 – 70.
- Friedman, Daniel.** 2007. “Market theories evolve, and so do markets.” *Journal of Economic Behavior & Organization*, 63(2): 247–255.
- Gilboa, Itzhak, and David Schmeidler.** 1989. “Maxmin expected utility with non-unique prior.” *Journal of mathematical economics*, 18(2): 141–153.
- Gjerstad, Steven.** 2003. “The Impact of Pace in Double Auction Bargaining.” UCLA Department of Economics Levine’s Bibliography 66615600000000192.
- Gjerstad, Steven, and John Dickhaut.** 1998. “Price formation in double auctions.” *Games and economic behavior*, 22(1): 1–29.
- Guarnaschelli, Serena, Anthony M. Kwasnica, and Charles R. Plott.** 2003. “Information Aggregation in Double Auctions: Rational Expectations and the Winner’s Curse.” *Information Systems Frontiers*, 5: 63–77.
- Harsanyi, John C.** 1967. “Games with Incomplete Information Played by Bayesian Players. Part I: The Basic Model.” *Management Science*, 14: 159–182.
- Harstad, Ronald M., Aleksandar Sasa Pekec, and Ilia Tsetlin.** 2008. “Information Aggregation in Auctions with an Unknown Number of Bidders.” *Games and Economic Behavior*, 62: 467–508.
- Hayashi, Takashi.** 2009. “Context dependence and consistency in dynamic choice under uncertainty: the case of anticipated regret.” *Theory and Decision*, 70(4): 399–430.
- Hayashi, Takashi, and Hisayuki Yoshimoto.** 2012. “Risk- & Regret-Averse Bidders in Sealed-Bid Auctions.” Working Paper.
- Hayek, Friedrich A.** 1945. “The Use of Knowledge in Society.” *American Economic Review*, 35: 519–530.
- He, Minghua, Ho fung Leung, and Nick R. Jennings.** 2003. “A fuzzy-logic based bidding strategy for autonomous agents in continuous double auctions.” *Knowledge and Data Engineering, IEEE Transactions on*, 15(6): 1345 – 1363.

- Hong, Han, and Matthew Shum.** 2004. “Rates of Information Aggregation in Common Value Auctions.” *Journal of Economic Theory*, 116: 1–40.
- Hsu, Ming, Meghana Bhatt, Ralph Adolph, Daniel Tranel, and Colin F. Camerer.** 2005. “Neural Systems Responding to Degrees of Uncertainty in Human Decision-Making.” *Science*, 310: 1680–1683.
- Isaac, R. Mark, and James M. Walker.** 1988. “Group Size Effects in Public Good Provision: The Voluntary Contributions Mechanism.” *The quarterly Journal of Economics*, 103(1): 179–199.
- Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji.** 2005. “A Smooth Model of Decision Making Under Ambiguity.” *Econometrica*, 73: 1849–1892.
- Knight, Frank H.** 1912. *Risk, uncertainty and profit*. Courier Dover Publications.
- Kremer, Ilan.** 2002. “Information Aggregation in Common Value Auctions.” *Econometrica*, 70(4): 1675–1682.
- Kreps, David.** 1988. *Notes On The Theory Of Choice. Underground Classics in Economics*, Westview Press.
- Linhart, Peter B, and Roy Radner.** 1989. “Minimax-regret strategies for bargaining over several variables.” *Journal of Economic Theory*, 48(1): 152–178.
- Loertscher, Simon, and Leslie M. Marx.** 2015. “Prior-Free Bayesian Optimal Double-Clock Auctions.” Working Paper, University of Melbourne.
- Lubin, Benjamin, and David C. Parkes.** 2012. “Approximate Strategyproofness.” *Current Science*, 103(9): 10211032.
- Lucas, Robert.** 1978. “Asset Prices in an Exchange Economy.” *Econometrica*, 46: 103–124.
- Luce, R. Duncan, and Howard Raiffa.** 1957. *Games and Decisions: Introduction and Critical Survey*. Dover, NY:Dover Publications.
- Machina, Mark J.** 1982. “‘Expected Utility’ Analysis Without the Independence Axiom.” *Econometrica*, 50(2): 277–323.
- Machina, Mark J.** 1989. “Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty.” *Journal of Economic Literature*, 27: 1622–1668.

- Mackie, Gerry.** 2003. *Democracy Defended*. Cambridge University Press.
- Ma, Huiye, and Ho-Fung Leung.** 2007. “An adaptive attitude bidding strategy for agents in continuous double auctions.” *Electronic Commerce Research and Applications*, 6(4): 383 – 398. |ce:title;Intelligent agents in e-services|/ce:title;.
- Malinvaud, Edmond.** 1952. “Note on Von Neumann-Morgenstern’s Strong Independence Axiom.” *Econometrica*, 20: 679.
- Manne, Alan S.** 1952. “The Strong Independence Assumption – Gasoline Blends and Probability Mixtures.” *Econometrica*, 20: 665–669.
- Marschak, Jacob, and Roy Radner.** 1954. *Note on Some Proposed Decision Criteria*. John Wiley.
- Milgrom, Paul.** 1979. “A Convergence Theorem fo Competitive Bidding with Differential Information.” *Econometrica*, 47: 670–688.
- Milnor, John.** 1954. “Games Against Nature.” In *Decisions Processes*. , ed. R. Thrall, C. Coombs and R. Davis, 49–59. New York:Wiley.
- Myerson, Roger B., and Mark A. Satterthwaite.** 1983. “Efficient Mechanisms for Bilateral Trading.” *Journal of Economic Theory*, 29: 265–281.
- Pesendorfer, Wolfgang, and Jeroen M. Swinkels.** 1997. “The Loser’s Curse and Information Aggregation in Common Value Auctions.” *Econometrica*, 65(6): 1247–1281.
- Radner, Roy.** 1979. “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices.” *Econometrica*, 47(3): 655–678.
- Rajan, Vijay, James R. Slagle, John Dickhaut, and Arijit Mukherji.** 1997. “Decentralized Problem Solving Using the Double Auction Market Institution.” *Expert Systems with Applications*, 12(1): 1–10.
- Ray, Paramesh.** 1973. “Independence of Irrelevant Alternatives.” *Econometrica*, 41(5): 987–91.
- Renou, Ludovic, and Karl H. Schlag.** 2008. “Minimax Regret and Strategic Uncertainty.”
- Reny, Philip J., and Motty Perry.** 2006. “Toward a Strategic Foundation for Rational Expectations Equilibrium.” *Econometrica*, 74(5): 1231–1269.

- Rustichini, Aldo, Mark A. Satterthwaite, and Steven R. Williams.** 1994. "Convergence to Efficiency in a Simple Market with Incomplete Information." *Econometrica*, 62(5): pp. 1041–1063.
- Samuelson, Paul A.** 1952. "Probability, Utility, and the Independence Axiom." *Econometrica*, 20: 670–678.
- Satterthwaite, Mark A., and Steven R. Williams.** 1993. "The Bayesian Theory of the k-Double Auction." In *The Double Auction Market*, ed. Daniel Friedman and John Rust, 99–123. Addison-Wesley.
- Savage, Leonard J.** 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association*, 46(253): pp. 55–67.
- Seidenfeld, Teddy.** 1988. "Decision Theory Without "Independence" or Without "Ordering"." *Economics and Philosophy*, 4: 267–290.
- Sosic, Greys.** 2007. "Collusion in Second-Price Auctions under Minimax Regret Criterion." *Production and Operations Management*, 16(4): 471–482.
- Stoye, Joerg.** 2011*a*. "Axioms for minimax regret choice correspondences." *Journal of Economic Theory*, 146(6): 2226–2251.
- Stoye, Joerg.** 2011*b*. "Statistical decisions under ambiguity." *Theory and Decision*, 70(2): 129–148.
- Stoye, Joerg.** 2012. "New Perspectives on Statistical Decisions Under Ambiguity." *Annual Review of Economics*, 4(1): 257–282.
- Wald, Abraham.** 1945. "Statistical Decision Functions Which Minimize the Maximum Risk." *Annals of Mathematics*, 46(2): pp. 265–280.
- Wilson, Robert.** 1977. "A Bidding Model of Perfect Competition." *Review of Economic Studies*, 44: 511–518.
- Wilson, Robert.** 1987. "On Equilibria of Bid-Ask Markets." In *Arrow and the Ascent of Modern Economic Theory*, ed. G. W. Feiwel. New York University Press.
- Wold, Herman.** 1952. "Ordinal Preferences or Cardinal Utility?" *Econometrica*, 20: 661–664.