Regret-Minimizers and Convergence to Price-Taking

Rachel C. Shafer *

October 31, 2015

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

*University of Illinois, shafer1@illinois.edu.

1 Introduction

Perfectly competitive markets are efficient only if traders act as price takers[1], 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 paper 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 paper examines three separate versions of 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 paper'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 paper, 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 paper. Taken together, the paper'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

[1]By "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.

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 paper considers, is well-suited to situations of complete ignorance because it essentially accommodates all priors at once. This paper finds that minimax regret traders do not converge to price-taking behavior.

The second reason against 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 paper 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 reason against 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 paper 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 paper'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 priors" 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), followed by a further exploration of the information structure (section 3). 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 7 concludes.

2 Traders and Auction Rules

This section describes the models traders, institution, and traders' profit functions, all of which follow the standard approach to a private value sealed-bid double auction.

2.1 Traders with Private Values

There are m buyers and n sellers. Each buyer i has a valuation $v^i \in [\underline{v}, \overline{v}] \subset \mathbb{R}_+$, which is the buyer's maximum willingness to pay for a single unit of the good. Each seller j has a cost $c^j \in [\underline{c}, \overline{c}] \subset \mathbb{R}_+$ 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.

2.2 k-Double Auction Rules

Agents submit sealed bids (denoted b^i) and asks (a^j) to the auctioneer. We assume that the submitted bids $(b^1, b^2, ..., b^m)$ are positive real numbers that cannot exceed the maximum valuation $\overline{v}$. We assume that sellers submit asks $(a^1, a^2, ..., a^n)$ 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 $Z = [\underline{c}, \overline{v}]$, where $\underline{c} < \overline{v}$.

These bids and asks determine a single price at which all units are traded and identifies 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 $k \in [0, 1]$:

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

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* $m+1^{st}$*-lowest bid or ask. Let* $(z_{(1)}, z_{(2)}, ..., z_{(m+n)}) \in Z^{m+n}$ *be the ordered set of bids and asks, with* $z_{(1)} < z_{(2)} < ... < z_{(m+n)}$. *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).

2.3 Profit Functions

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 ζ 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

$$\Pi^B(b, v, \zeta) = \begin{cases} v - [(1-k)\zeta_{(m)} + k\zeta_{(m+1)}] & \text{if } \zeta_{(m+1)} < b \\ v - [(1-k)\zeta_{(m)} + kb] & \text{if } \zeta_{(m)} < b < \zeta_{(m+1)} \\ 0 & \text{if } b < \zeta_{(m)} \end{cases}. \tag{2}$$

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

$$\Pi^A(a, c, \zeta) = \begin{cases} [(1-k)\zeta_{(m)} + k\zeta_{(m+1)}] - c & \text{if } a < \zeta_{(m)} \\ [(1-k)a + k\zeta_{(m+1)}] - c & \text{if } \zeta_{(m)} < a < \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < a \end{cases}. \tag{3}$$

The levels of profit, as it relates to the rival bids and asks $\zeta_{(m)}$ and $\zeta_{(m+1)}$, are shown in figures 1 and 2. Note that $\zeta_{(m)}$ is always less than $\zeta_{(m+1)}$, so that the possible outcomes all lie above the 45 degree line. The bidder's profit depends on $\zeta_{(m)}$ and $\zeta_{(m+1)}$ relative to his bid b (and likewise the seller for his ask a).

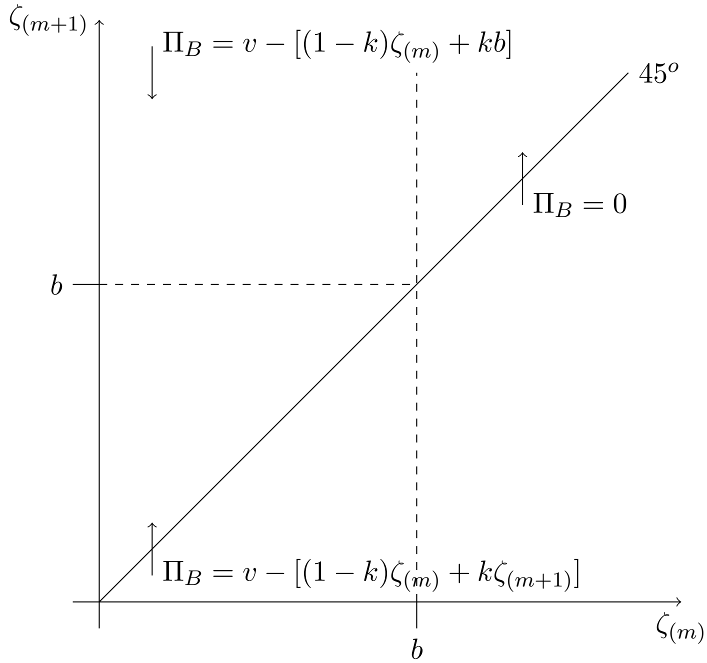

Figure 1: Bidder's Profit from bid b

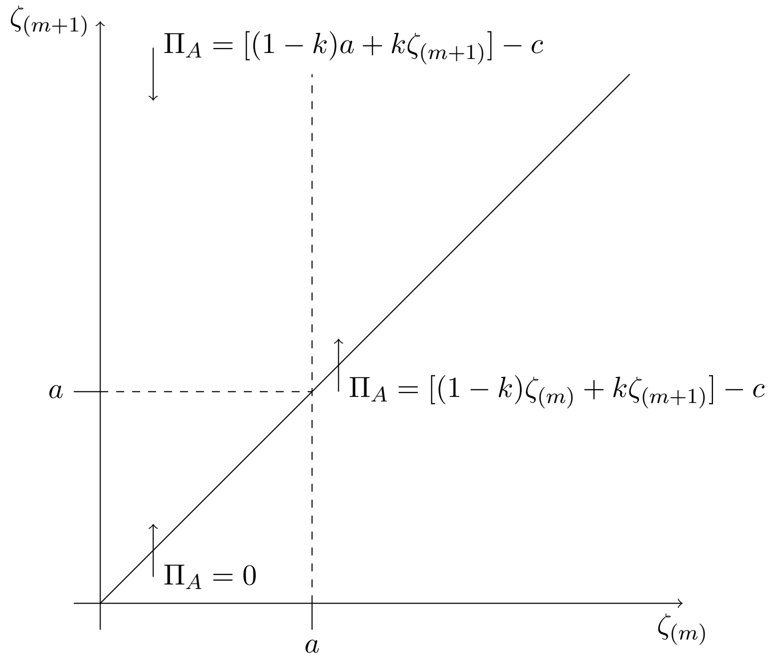

Figure 2: Sellers Profit from ask a

3 Information Structure

In analyzing the market described in section 2, it is typical to think of the auction as a game of incomplete information 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.

3.1 Defining a Decision Problem 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 $u : \mathcal{S} \times \mathcal{A} \to \mathbb{R}$.

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.

3.2 Contrasting Incomplete Information and Knightian Uncertainty

In an incomplete information game, the bidder does not know the costs and valuations of the other traders. However, the bidder does know the distribution from which these redemption values are drawn. This distribution, together with some equilibrium belief about the traders' strategies, is the bidder's source for his 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.

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 valuations and costs for the other traders. They know their own valuation/cost with certainty. However, they do not know the distribution of other traders' valuations and costs.

Two possible approaches to this decision problem under Knightian uncertainty are depicted in Figures 4 and 5. The crucial difference between these figures and Figure 3 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.

3.3 The Double Auction as a Decision Problem 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 4, 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 5, 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.

In analyzing the agent's problem as a decision problem under Knightian uncertainty, rather than as a game, it is important that this is a sealed-bid auction with private valuations. Because this is a situation in which traders act once (submitting sealed bids), the traders need not consider how their actions will reveal information to other traders, and thus alter their rivals' future behavior. Furthermore, private valuations mean that traders will not revise their own estimate of an item's true value based on the revelation of other traders' valuations. The only reason that the other traders' redemption values matter is that it affects how the other traders may be expected to bid.

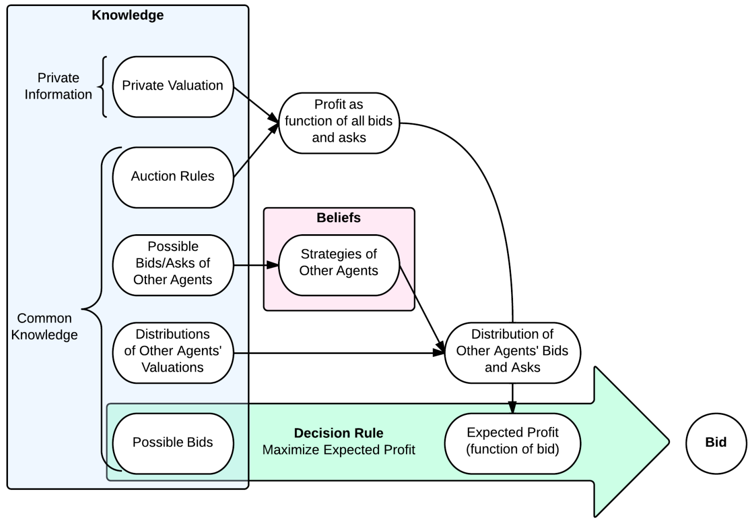

Figure 3: Bidder Decision in Private-Value k-Double Auction: Game of Incomplete Information

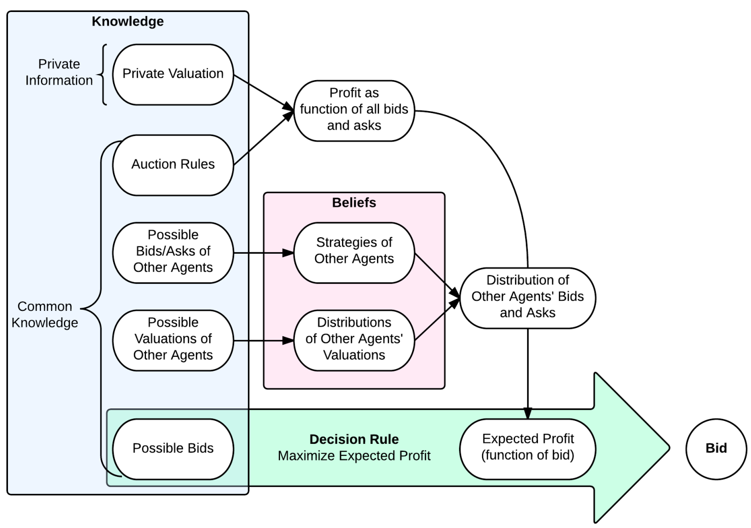

Figure 4: Bidder Decision in Private-Value k-Double Auction: Decision Problem Under Knightian Uncertainty (Bayesian Approach)

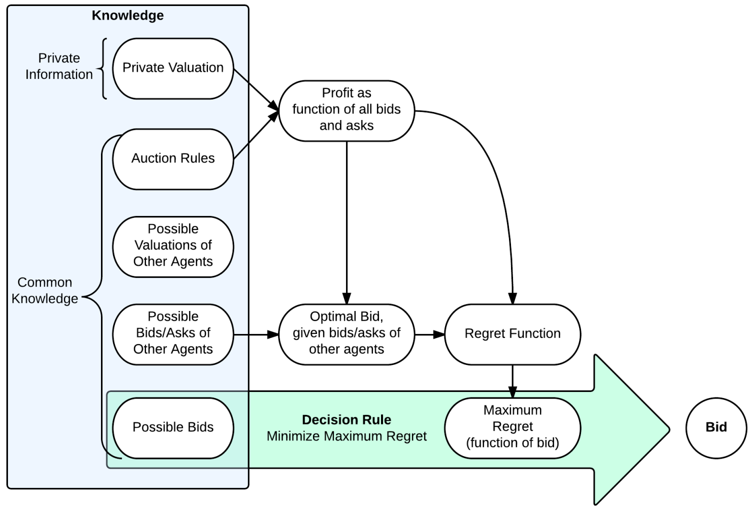

Figure 5: Bidder Decision in Private-Value k-Double Auction: Decision Problem Under Knightian Uncertainty (Minimax Regret))

4 First approach: minimizing maximum regret

This section formally defines minimax regret, and shows that minimax regret traders bid and ask functions do not reflect the traders true redemption values, and do not converge to truthful bidding no matter how large the market grows.

4.1 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 1 *An action a attains minimax regret if*

$$a \in \arg\min_{a \in \mathcal{A}} \max_{s \in \mathcal{S}} \{ \max_{a^* \in \mathcal{A}} u(a^*, s) - u(a, s) \} \} \tag{4}$$

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.

4.2 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.

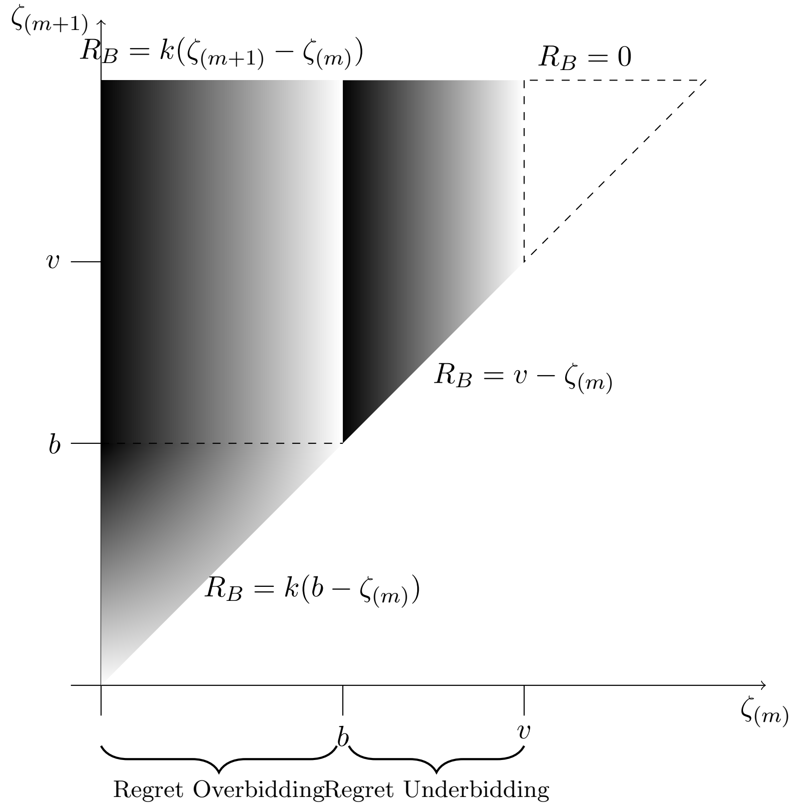

Figure 6: Bidder's Profit from bid b

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 using their first approach to the incomplete information bilateral case.

4.3 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 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 $n, m \to \infty$. The thin lines represent the truthful 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. 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.

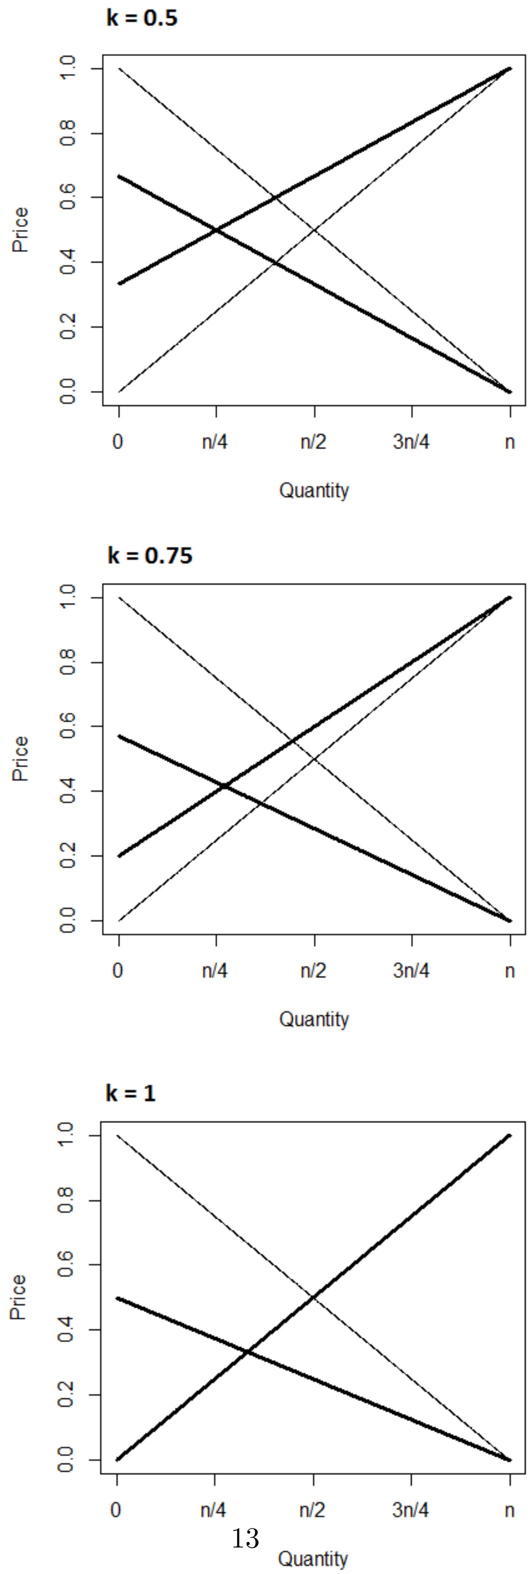

Figure 7: The Distribution of Bids and Asks Depends on k

5 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.

5.1 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 rationalizable[2] 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 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

[2]For a seller with cost c^i, any ask $a^i \in [c^i, \overline{c}]$ is rationalizable. For a buyer with cost v^i, any bid $b^i \in [\underline{b}, v^i]$ is rationalizable.

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 $\widetilde{F}$ 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*

$$\widetilde{F}(\tfrac{(k+1)b^i - v^i}{k}) = \tfrac{\widetilde{F}(b^i)}{1+k}$$

Such a bid b^i exists on the interval $\left[\frac{v^i}{1+k}, v^i\right]$.

Similarly, let $\widetilde{G}$ 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

$$\widetilde{G}(\tfrac{a^i(1+(1-k)) - c^i}{1-k}) = \tfrac{\widetilde{G}(a^i)+(1-k)}{1+(1-k)}$$

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".

5.2 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 to 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* $\lim_{n\to\infty} b(v;n) = \dfrac{v}{1+k}$.

Figure 8 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.

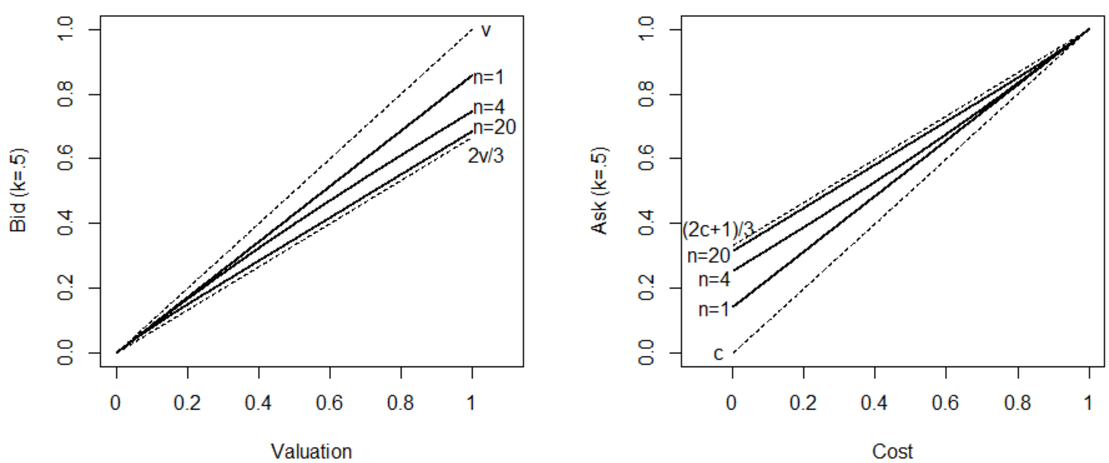

Figure 8: Bids and asks for traders that minimize expected maximum regret, at various market sizes

6 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.

6.1 Minimizing Maximum Expected Regret defined

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

Definition 2 *Let* Γ *denote a set of probability distriubtions on* $\mathcal{S}$. *An action* a *attains* Γ-*minimax regret if*

$$a \in \arg\min_{a \in \mathcal{A}} \{ \max_{\pi \in \Gamma} \{ \int \max_{a^* \in \mathcal{A}} u(a^*, s) - u(a, s) d\pi \} \} \tag{5}$$

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 $u(a,s) = \min_{s \in S} u(a,s)$. 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.

6.2 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. Converge 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* $i \geq N(b, G_\gamma)$*:*

$$\int u(v-\epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) > \int u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \quad (6)$$

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*

$$\overline{N}(\epsilon) = \max_{G_\gamma \in \Gamma} \{N(\epsilon, G_\gamma)\} \quad (7)$$

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* $m_i + 1^{th}$ *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 $i \to \infty$*.*

7 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.

Proofs

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 $\zeta_{(m+1)}$:

$$\Pi_B = \begin{cases} v^i - (k\zeta_{(m+1)} + (1-k)\zeta_{(m)}) & \text{if } \zeta_{(m+1)} < b^i \\ v^i - (kb^i + (1-k)\zeta_{(m)}) & \text{if } \zeta_{(m)} < b^i < \zeta_{(m+1)} \\ 0 & \text{if } b^i < \zeta_{(m)} \end{cases} \tag{8}$$

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

$$\Pi_B^* = \begin{cases} v^i - \zeta_{(m)} & \text{if } \zeta_{(m)} \leq v^i \\ 0 & \text{if } \zeta_{(m)} > v^i \end{cases} \tag{9}$$

Then the buyer's regret function is

$$R_B = \begin{cases} k(\zeta_{(m+1)} - \zeta_{(m)}) & \text{if } \zeta_{(m)} \leq v^i \text{ and } \zeta_{(m+1)} < b^i \\ k(b^i - \zeta_{(m)}) & \text{if } \zeta_{(m)} \leq v^i \text{ and } \zeta_{(m)} < b^i < \zeta_{(m+1)} \\ v^i - \zeta_{(m)} & \text{if } \zeta_{(m)} \leq v^i \text{ and } b^i < \zeta_{(m)} \\ 0 & \text{if } \zeta_{(m)} > v^i \text{ and } b^i < \zeta_{(m)} \end{cases} \tag{10}$$

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

$$\sup R_B = \max(kb^i, v^i - b^i, 0) \tag{11}$$

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.

$$\Pi_S = \begin{cases} (k\zeta_{(m+1)} + (1-k)\zeta_{(m)}) - c^i & \text{if } a^i < \zeta_{(m+1)} \\ (k\zeta_{(m+1)} + (1-k)a^i) - c^i & \text{if } \zeta_{(m)} < a^i < \zeta_{(m+1)} \\ 0 & \text{if } a^i > \zeta_{(m+1)} \end{cases} \tag{12}$$

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

$$\Pi_S^* = \begin{cases} \zeta_{(m+1)} - c^i & \text{if } c^i \leq \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < c^i \end{cases} \tag{13}$$

Then the seller's regret function is

$$R_S = \begin{cases} (1-k)(\zeta_{(m+1)} - \zeta_{(m)}) & \text{if } \zeta_{(m+1)} \geq c^i \text{ and } a^i < \zeta_{(m+1)} \\ (1-k)(\zeta_{(m+1)} - a^i) & \text{if } \zeta_{(m+1)} \geq c^i \text{ and } \zeta_{(m)} < a^i < \zeta_{(m+1)} \\ \zeta_{(m+1)} - c^i & \text{if } \zeta_{(m+1)} \geq c^i \text{ and } a^i > \zeta_{(m+1)} \\ 0 & \text{if } \zeta_{(m+1)} < c^i \text{ and } a^i > \zeta_{(m+1)} \end{cases} \tag{14}$$

The supremum of the seller's regret function is

$$\sup R_B = \max((1-k)(1-a^i), a^i - c^i, 0) \tag{15}$$

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 $\widetilde{F}$ 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*

$$\widetilde{F}\left(\tfrac{(k+1)b^i - v^i}{k}\right) = \tfrac{\widetilde{F}(b^i)}{1+k}$$

Such a bid b^i exists on the interval $\left[\frac{v^i}{1+k}, v^i\right]$.

Similarly, let $\widetilde{G}$ 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

$$\widetilde{G}\left(\tfrac{a^i(1+(1-k)) - c^i}{1-k}\right) = \tfrac{\widetilde{G}(a^i)+(1-k)}{1+(1-k)}$$

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 regret[3]. The maximum regret for a bid $b^i < v^i$ given the realization of costs c and values v is

$$R_B(b^i|v,c) = \begin{cases} 0 & \text{if } c_{(1)} > v^i \\ \max\{v^i - b^i, k(b^i - c_{(1)})\} & \text{if } v^i \geq b^i \geq c_{(1)} \\ v^i - c_{(1)} & \text{if } v^i \geq c_{(1)} > b^i \end{cases} \tag{16}$$

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

[3]The following discussion follows Linhardt and Radner (1989) closely.

necessary, depending on the value of $c_{(1)}$:

$$v^i - b^i < k(b^i - c_{(1)})$$
$$\Rightarrow c_{(1)} < \frac{(k+1)b^i - v^i}{k}$$

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

$$\overline{R}_B(b^i|v^i) = \int_0^{b^*} k(b^i - c)d\widetilde{F}(c) + \int_{b^*}^{b^i} (v^i - b^i)d\widetilde{F}(c) + \int_{b^i}^{v^i} (v^i - c)d\widetilde{F}(c) \tag{17}$$

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

$$\overline{R}_B(b^i|v^i) = \int_0^{b^i} (v^i - b^i)d\widetilde{F}(c) + \int_{b^i}^{v^i} (v^i - c)d\widetilde{F}(c) \tag{18}$$

Integrating by parts,

$$\overline{R}_B(b^i|v^i) = \begin{cases} k\int_0^{\frac{(k+1)b^i - v^i}{k}} \widetilde{F}(c)dc + \int_{b^i}^{v^i} \widetilde{F}(c)dc & \text{if } b^i > \frac{v^i}{1+k} \\ \int_{b^i}^{v^i} \widetilde{F}(c)dc & \text{if } b^i \leq \frac{v^i}{1+k} \end{cases} \tag{19}$$

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

$$\frac{d}{db^i}\overline{R}_B = \begin{cases} (k+1)\widetilde{F}(\frac{(k+1)b^i - v^i}{k}) - \widetilde{F}(b^i) & \text{if } \frac{v^i}{1+k} \leq b^i \leq v^i \\ -\widetilde{F}(b^i) & \text{if } 0 \leq b^i \leq \frac{v^i}{1+k} \end{cases} \tag{20}$$

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

$$R_S(a^i|v, c) = \begin{cases} 0 & \text{if } v_{(m)} < c^i \\ \max\{a^i - c^i, (1-k)(v_{(m)} - a^i)\} & \text{if } a^i \geq v_{(m)} \\ v_{(m)} - c^i & \text{if } c^i \geq v_{(m)} < a^i \end{cases} \tag{21}$$

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 $v_{(m)}$:

$$a^i - c^i < (1-k)(v_{(m)} - C_i)$$
$$\Rightarrow v_{(m)} > \frac{(1+(1-k))a^i - c^i}{1-k}$$

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

$$\overline{R}_S(a^i|c^i) = \int_{c^i}^{a^i} (v - c^i)d\widetilde{G}(v) + \int_{a^i}^{C^*} (a^i - c^i)d\widetilde{G}(v) + \int_{C^*}^{1} (1-k)(v - a^i)d\widetilde{G}(v) \quad (22)$$

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:

$$\overline{R}_S(a^i|c^i) = \int_{c^i}^{a^i} (v - c^i)d\widetilde{G}(v) + \int_{a^i}^{1} (a^i - c^i)d\widetilde{G}(c) \quad (23)$$

Integrating by parts,

$$\overline{R}_S(a^i|c^i) = \begin{cases} (1-k)(1-a^i) - \int_{c^i}^{a^i} \widetilde{G}(v)dv - \int_{C^*}^{1} (1-k)\widetilde{G}(v)dv & \text{if } a^i > \frac{c^i+(1-k)}{1+(1-k)} \\ (a^i - c^i) - \int_{c^i}^{a^i} \widetilde{G}(v)dv & \text{if } a^i < \frac{c^i+(1-k)}{1+(1-k)} \end{cases} \quad (24)$$

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

$$\frac{d\overline{R}_S(a^i|c^i)}{da^i} = \begin{cases} -(1-k) - \widetilde{G}(a^i) + (2-k)\widetilde{G}(\frac{(1+(1-k))a^i - c^i}{1-k}) & \text{if } a^i > \frac{c^i+(1-k)}{1+(1-k)} \\ 1 - \widetilde{G}(c^i) & \text{if } a^i < \frac{c^i+(1-k)}{1+(1-k)} \end{cases} \quad (25)$$

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. ■

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

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

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

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:

$$\lim_{n\to\infty} \left[(k+1)\widetilde{F}\left(\frac{(k+1)(\frac{v}{1+k}+\epsilon) - v}{k} \right) - \widetilde{F}\left(\frac{v}{1+k} + \epsilon \right) \right] < 0 \tag{26}$$

$$(k+1) \lim_{n\to\infty} \widetilde{F}\left(\frac{(k+1)\epsilon}{k} \right) - \lim_{n\to\infty} \widetilde{F}\left(\frac{(k+1)\epsilon}{k+1} \right) < 0 \tag{27}$$

$$(k+1)(1) - 1 < 0 \tag{28}$$

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* $i \geq N(b, G_\gamma)$*:*

$$\int u(v-\epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) > \int u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \tag{29}$$

   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*

$$\overline{N}(\epsilon) = \max_{G_\gamma \in \Gamma} \{N(\epsilon, G_\gamma)\} \tag{30}$$

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: (30) implies that for $i \geq \overline{N}(\epsilon)$,

$$\int \max_{b^*\in[0,v]} u(b^*, \zeta) - u(v-\epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) < \int \max_{b^*\in[0,v]} u(b^*, \zeta) - u(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v - \epsilon \tag{31}$$

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

$$\int R(v-\epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ < \int R(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dG_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v-\epsilon \quad (32)$$

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

$$\max_{G_{\gamma,i}\in\Gamma_i} \int R(v-\epsilon, \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ < \max_{G_{\gamma,i}\in\Gamma_i} \int R(b', \zeta_{(m_i)}, \zeta_{(m_i+1)}) dF_{\gamma,i}(\zeta_{(m_i)}, \zeta_{(m_i+1)}), \forall b' < v-\epsilon \quad (33)$$

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

$$\int R(v-\epsilon, \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ > \max_{G_{\gamma,j}\in\Gamma_j} \int R(b', \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{\gamma,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \quad (34)$$

implying

$$\int R(v-\epsilon, \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \\ > \int R(b', \zeta_{(m_j)}, \zeta_{(m_j+1)}) dG_{*,j}(\zeta_{(m_i)}, \zeta_{(m_i+1)}) \quad (35)$$

which contradicts (32).

We conclude that the maximum expected regret-minimizing bid can be arbitrarily close to bidding v, given a sufficient number of competitors. For (37) establishes that for a number of bidders greater than or equal to $\overline{N}(\epsilon)$, 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 $m_i + 1^{th}$ 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 $i \to \infty$.

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 $m_i^t h$ and $m_i + 1^t h$ order statistics is

$$f_{(m_i)(m_i+1)}(x, y) = \frac{(m_i + n_i - 1)!}{(m_i - 1)!n_i!}(x)f(x)f(y)F^{m_i-1}[1 - F(y)]^{n_i} \tag{36}$$

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:

$$\begin{aligned} E[R_B] = &\int_0^b \int_x^b k(y - x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1 - F(y)]^n dydx \\ &+ \int_0^b \int_b^1 k(b - x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1 - F(y)]^n dydx \\ &+ \int_b^v \int_x^1 (v - x)C_{m,m+1}f(x)f(y)F(x)^{m-1}[1 - F(y)]^n dydx \end{aligned} \tag{37}$$

$$\begin{aligned} \frac{dE[R_B]}{db} = &\int_0^b k(b - x)C_{m,m+1}f(x)f(b)F(x)^{m-1}[1 - F(b)]^n dx \\ &- \int_0^b k(b - x)C_{m,m+1}f(x)f(b)F(x)^{m-1}[1 - F(b)]^n dx \\ &+ \int_0^b \int_b^1 kC_{m,m+1}f(x)f(y)F(x)^{m-1}[1 - F(y)]^n dydx \\ &- \int_b^1 (v - b)C_{m,m+1}f(b)f(y)F(b)^{m-1}[1 - F(y)]^n dy \end{aligned} \tag{38}$$

$$\begin{aligned} \frac{dE[R_B]}{db} = &\int_0^b \int_b^1 kC_{m,m+1}f(x)f(y)F(x)^{m-1}[1 - F(y)]^n dydx \\ &- \int_b^1 (v - b)C_{m,m+1}f(b)f(y)F(b)^{m-1}[1 - F(y)]^n dy \end{aligned} \tag{39}$$

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

$$\int_0^b \int_b^1 kf(x)f(y)F^{m-1}(x)[1-F(y)]^n dydx = \int_b^1 (v-b)f(b)f(y)F^{m-1}(b)[1-F(y)]^n dy \tag{40}$$

$$k\int_0^b f(x)F(x)^{m-1}dx = (v-b)f(b)F(b)^{m-1} \tag{41}$$

$$k\frac{F(b)^m}{m} = (v-b)f(b)F(b)^{m-1} \tag{42}$$

$$\frac{kF(b)}{mf(b)} = (v-b) \tag{43}$$

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. ■

References

Anscombe, F. J., and R. J. Aumann. 1963. "A Definition of Subjective Probability." *The Annals of Mathematical Statistics*, 34(1): 199–205.

Arrow, K. J., and L. 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.

Chernoff, H. 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 Allen Zhu. 2014. "Knightian Robustness from Regret Minimization." *CoRR*, abs/1403.6409.

Cohen, Michele, and Jean-Yves Jaffray. 1980. "Rational Behavior under Complete Ignorance." *Econometrica*, 48(5): 1281–99.

Dempster, A.P. 1975. "A subjectivist look at robustness." *Bulletin of International Statistical Institute*, 46: 349–374.

Fishburn, P.C. 1970. *Utility Theory for Decision Making.* Wiley, New York.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin expected utility with non-unique prior." *Journal of mathematical economics*, 18(2): 141–153.

Harsanyi, John C. 1967. "Games with Incomplete Information Played by Bayesian Players. Part I: The Basic Model." *Management Science*, 14: 159–182.

Knight, Frank H. 1912. *Risk, uncertainty and profit.* Courier Dover Publications.

Kreps, D. 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.

Lubin, Benjamin, and David C. Parkes. 2012. "Approximate Strategyproofness." *Current Science*, 103(9): 10211032.

Luce, R. Duncan, and Howard Raiffa. 1957. *Games and Decisions: Introduction and Critical Survey.* Dover, NY:Dover Publications.

Marschak, Jacob, and Roy Radner. 1954. *Note on Some Proposed Decision Criteria.* John Wiley.

Milnor, J. 1954. "Games Against Nature." In *Decisions Processes.* , ed. R. Thrall, C. Coombs and R. Davis, 49–59. New York:Wiley.

Ray, Paramesh. 1973. "Independence of Irrelevant Alternatives." *Econometrica*, 41(5): 987–91.

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.

Savage, L. J. 1951. "The Theory of Statistical Decision." *Journal of the American Statistical Association*, 46(253): pp. 55–67.

Seidenfeld, Teddy. 1988. "Decision Theory Without "Indepdendence" or Without "Ordering"." *Economics and Philosophy*, 4: 267–290.

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, R. 1987. “On Equilibria of Bid-Ask Markets.” In *Arrow and the Ascent of Modern Economic Theory.* , ed. G. W. Feiwel. New York University Press.