

The Evolution of Bidding Behavior in Private-Values Auctions and Double Auctions

Rene Saran*

Brown University, Providence, U.S.A.

Roberto Serrano†

Brown University, Providence, U.S.A.

IMDEA, Madrid, Spain

Abstract

We apply stochastic stability to study the evolution of bidding behavior in private-values second-price, first-price and k -double auctions. The learning process has a strong component of inertia but with a small probability, the bids are modified in the direction of ex-post regrets. We identify essentially a unique bid that will be used by each type in the long run. In the second-price auction, this is the truthful bid. In the first-price auction, bidding half of one's valuation is stable. The stable bid in the k -double auction is a toughening of the Chatterjee-Samuels linear equilibrium strategy. If we add a friction in changing one's bid, then truth-telling behavior is also obtained in the first-price and k -double auctions. Intuitively, the stochastically stable bid minimizes the maximal regret.

Keywords: Stochastic Stability; Ex-post Regret; Second-Price Auction; First-Price Auction; k -Double Auction

JEL Classification: C73; C78; D44; D83

**Email address:* rene_saran@brown.edu

†*Email address:* roberto_serrano@brown.edu

1 Introduction

The standard analysis of most auctions and double auctions with incomplete information is based on the notion of Bayesian equilibrium, and has led to beautiful theoretical constructions. Moreover, equilibrium insights can sometimes be applied with success to real-world problems. However, the use of equilibrium also has important limitations, as it relies on the existence of a common-knowledge type space to describe the underlying uncertainty. When one talks to game theorists that have provided advice on how to bid in real auctions, this is often a major stumbling block: we may calculate an equilibrium of the given auction, but in the absence of common-knowledge of type spaces and prior beliefs, it is implausible to expect the equilibrium to be played. One common way out in recent literature is to shift the analysis from the interim stage to the ex-post stage: by requiring that the equilibrium be robust for any type space, one ends up at the concept of ex-post equilibrium.¹

The approach we shall follow in this paper uses stochastic stability, one of the important tools introduced by evolutionary game theory (see Foster and Young (1990), Kandori, Mailath and Rob (1993), Young (1993)). Most applications of this methodology in non-cooperative game theory have been confined to games with complete information.² An exception is Jensen, Sloth and Witta-Jacobson (2005), which extends the model in Young (2003) to finite Bayesian games. They make three assumptions that are necessary for this extension. First, the players know the true distribution of types in the population; second, the types of the matched players are truthfully revealed to everyone after the end of the interaction; and third, for each type of each player, there is a record of the action taken by that type during some past periods in which that type was selected. Under these assumptions, the unperturbed best-response dynamic process, appropriately redefined, converges with probability one to a convention, which is a state that is “equivalent” to a strict Bayesian equilibrium of the game – if the latter exists. The perturbations then select among the different strict Bayesian equilibria. We have already noted that the scale of information needed to satisfy the first assumption is rarely available. It is also not clear how the second assumption would be satisfied. The only information that is revealed at the end of an interaction is the outcome and the consequent payoffs, which is not sufficient

¹One well-known difficulty with this concept is that it rarely exists.

²See Young (1998) for an account of different applications. See also Agastya (2004) for an application to double auctions with complete information.

to reveal the types – specially when players have private values.

Agreeing with the recent trends in the robust analysis of game theory with incomplete information, we shall turn to ex-post considerations. But, rather than working with the concept of ex-post equilibrium, we shall analyze an evolutionary model of behavior based on ex-post regrets. Hyafil and Boutilier (2004) propose minimizing the maximal ex-post regret as a decision criterion in such environments, which leads them to define a minimax-regret equilibrium. We will later show that the profile of stable bids in our model is a minimax-regret equilibrium. A description of our model follows.

There is a large number of traders, sellers and potential buyers of an indivisible good. We study three kinds of trading institutions, all under private-values asymmetric information (second-price auctions, first-price auctions and double auctions).

In our study of auctions, each seller has a zero valuation for the object and sets a reservation price of zero. Each bidder is privately informed about her valuation for the good and has no information about the valuations of the others. No assumptions on uncertainty are made, except that it is diffuse, in the sense that each bidder believes that there is always a positive probability of playing against any bid for the good.³ Each seller will be selling one good in each period, and is matched at random to a number of bidders. In each match, a one-shot auction procedure is used. We first analyze a second-price auction in Section 2, while a first-price auction is studied in Section 3. Each bidder's behavior is described as follows. First, as often found in much of the evolutionary literature, there is an important component of inertia. That is, with very high probability, a bidder, if she participates in an auction in period $t + 1$, will repeat the bid that she used in period t (she is convinced by a convention learned in the past, or by the game theorist that gives her expert advice). However, with small but positive probability, she will adjust her bid in the direction of ex-post regret. That is, she looks at the outcome of the last auction where she participated, and wonders what would have been her payoff had she bid differently against the specific bids used by the other bidders she met in the last match. If by changing to a different bid, her ex-post payoff would have increased, she regrets having used her bid instead of the alternative bid. The higher this regret the higher her probability of

³This is accomplished through the introduction of a rich set of obstinate types, agents who never change their behavior (see, for example, Ellingsen (1997), Abreu and Gul (2000)).

switching to such an alternative bid.

Thus, within the useful classification provided in Hart (2005), our models lie somewhere between “evolution” and “adaptive heuristics.” The evolutionary aspect of our work stems from the inertial component: with very high probability, our players are “programmed” to use a specific bid. However, our adaptive learning rule based on ex-post regret has an important component of heuristics. Quoting Hart (2005, p. 1403): “We use the term *heuristics* for rules of behavior that are simple, unsophisticated, simplistic and myopic. These are “rules of thumb” that the players use to make their decisions. We call them *adaptive* if they induce behavior that reacts to what happens in the play of the game, in directions that, loosely speaking, seem “better.”” Our process is related to the regret matching algorithms (e.g., Foster and Vohra (1998), Hart and Mas-Colell (2000)) with two important differences. First, we define regret with respect to the last one-shot interaction of a player, instead of the average regret entailed in regret matching, where a player carries over her time average payoff every period and wonders how that average payoff would have been affected had she played a different action. And second, in regret matching, the probability of switching is proportional to regret, so such switches are typically likely, whereas in our case they are infinitesimals, although the rate at which these probabilities go to zero is determined by the amount of regret. Our assumptions lead to an ergodic process and we characterize the stochastically stable states: that is, we identify the bids that will be used in the very long run most of the time, when behavior is affected by a negligible amount of ex-post regret.

In the second-price auction, our result identifies the truthful bid as the only one that is stochastically stable. This is understood as a consequence of this bid being dominant – in fact, the same result obtains in the double auction model, whenever it is dominant for a player to bid truthfully. On the other hand, the result for the first-price auction is sensitive to the presence of a friction to change one’s bid.⁴ If no such frictions exist, the unique stochastically stable bid is half of one’s valuation, independent of the number of competing bidders – as will be explained below, ex-post versus interim considerations explain this, since all that matters is the realization of the highest competitor bid. But if changing bids by much is not feasible,

⁴The presence of such frictions may be justified by “real-world” considerations. Often, for a bid to be accepted at an auction, complex legal work is required to back the plausibility of the bid in question.

stochastic stability selects essentially truth-telling behavior. This provides a different rationale to truth-telling or “trustworthy” behavior from others found in the literature.⁵

In Section 4 the trading institution is a k -double auction. That is, privately informed buyers and sellers of an indivisible good are matched in pairs, and when they are matched, they make simultaneous price announcements. Trade takes place if and only if the buyer’s bid exceeds the seller’s, at the k -weighted average of the two prices. Inertia and the adaptive learning rule based on ex-post regret as explained above determine traders’ behavior. Our result shows that, when frictions to change one’s bid do not exist, the unique bid selected by stochastic stability is a toughening of the linear equilibrium strategy of the uniform distribution model (Chatterjee and Samuelson (1983)). As in the first-price auction, if changing one’s bid by large amounts is not possible, we find that stochastic stability selects bids close to the truthful valuations.

Section 5 provides a detailed discussion of the intuition behind our results. Basically, the bids selected by stochastic stability in our adaptive learning model can be understood as those that equalize the maximum gain from increasing and decreasing one’s bid. Equivalently, they minimize the maximal regret, and they can also be seen as the best reply to a “veil of ignorance,” i.e., a uniform belief over other bids. In the absence of frictions on the amount of change in one’s bid, our agents behave as if they held uniform beliefs, and correspondingly, act tougher than if the corresponding equilibrium of a uniform game were being played. To enhance the comparison with that equilibrium, our agents’ behavior leads to not succeeding in trade more often, but when one does succeed, one obtains a substantially larger payoff. The paper closes with our brief conclusions, and a section that collects the proofs of the major results.

2 Second-Price Auction

There is a single indivisible object. Any seller of that indivisible object values it at 0 and this is known to the potential buyers or bidders. A bidder for that indivisible object values it at $v \in V \equiv \{0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\}$. A bidder’s

⁵It has been argued that trustworthy behavior arises due to self-interest in maintaining a relationship, other-regarding preferences, moral commitment etc. See Ben-Ner and Putterman (2001) for a brief discussion.

valuation is her private information.

For every $v \in V$, there exist a large number of bidders with their valuations equal to v . Let \mathcal{B}_v denote the population of bidders with their valuations equal to v . Let \mathcal{S}_0 denote the population of sellers.

In each period t , a finite number M_t of sellers are selected randomly from \mathcal{S}_0 . Each selected seller is then matched with $M' \geq 2$ bidders that are selected randomly and independently from $\mathcal{B} = \bigcup_{v \in V} \mathcal{B}_v$. Note that M_t can be a random variable so that the number of auctions can be different in different periods.⁶ For simplicity, we shall fix M' , the size of any auction.⁷

The trading mechanism used by all the sellers is a second-price sealed-bid auction. That is, all the bidders matched to a seller submit sealed bids to the seller. The bidder with the highest bid wins the object. If more than one bidder bids the highest, then a random tie-breaking rule is used by the seller. The bidder who wins the object pays a price equal to the second-highest bid.⁸ If the object is traded at price p , the winning bidder i 's payoff is $v_i - p$ and all the other bidders matched to the seller get a payoff of 0. The seller's payoff is p .

A bidder can bid any number $\sigma \in \Sigma = \{0, \delta, 2\delta, \dots, 1\}$, where δ is the indivisible money unit. Assume that δ is small enough so that there exists an integer N such that $N\delta = \frac{1}{n}$. Otherwise, some bidders will not have the option to bid truthfully.

Each bidder in this market can be identified by her bidding rule, which specifies her unique bid in the event she is matched with a seller. So when we say that a bidder's bid in period t is σ , we mean that she will bid σ in case she is matched with a seller in period t .

We assume that the bidders in this market have complete inertia with respect to their bids, that is, in each period, every bidder uses the same bid as in the previous period. Thus, the distribution of bids in the market during any period is exactly the initial distribution. Therefore, without adding more structure, the model has no predictive power. To answer which of the large

⁶In that case, we assume that M_t and $M_{t'}$ have the same distribution to ensure that the dynamic process is time homogeneous.

⁷It is straightforward to make the size of each auction random. In this case, it will be natural to assume that the bidding rule of a bidder specifies her bids for every possible size of the auction in which she can play. The result will not change in both the second-price auction and first-price auction models if one adopts this modification.

⁸To be precise, if more than one bidder bid the highest amount, then the price is equal to this bid.

number of possible distributions of bids will be observed in the long run, we add a perturbation to the model using a random adaptive learning rule defined in the next subsection.

2.1 Random Adaptive Learning Rule

A subset of bidders, \mathcal{P} , adapt their bid from period t to $t + 1$ using the following random adaptive learning rule: suppose in period t a bidder in \mathcal{P} with valuation v_i is using the bid σ_i .

- If this bidder is not matched in period t , then she does not change her bid if she were matched in period $t + 1$.
- If this bidder is matched in period t , then let $\pi(v_i, \sigma_i, \sigma_{-i})$ be her payoff in that auction, where σ_{-i} is the profile of bids of all other bidders in that auction. Pick any $\sigma'_i \in \Sigma(\eta, \sigma_i) \equiv \{\sigma : 0 < |\sigma - \sigma_i| \leq \eta\delta\}$, where $\eta \geq 1$. Had she bid σ'_i instead of σ_i in that auction, *ceteris paribus*, her payoff would have been $\pi(v_i, \sigma'_i, \sigma_{-i})$. Let σ_{-i}^* denote the highest bid among σ_{-i} . We assume that if $\sigma'_i = \sigma_{-i}^*$, then $\pi(v_i, \sigma'_i, \sigma_{-i}) = 0$. That is, had her bid of σ'_i made her one of the more-than-one highest bidders, then she believes that she would have lost the random tie-breaking rule of the seller.⁹

Define $\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) = \pi(v_i, \sigma'_i, \sigma_{-i}) - \pi(v_i, \sigma_i, \sigma_{-i})$.

Let $\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) > 0$. Then, we refer to this number as valuation v_i 's *regret* from bidding σ_i instead of σ'_i against σ_{-i} . In this case, the bidder changes her bid to σ'_i in period $t+1$ with probability $\epsilon^{\frac{1}{\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i})}}$. That is, while inertia continues to determine her bid with high probability, switches to other bids are an increasing function of the corresponding ex-post regret. However, if $\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) \leq 0$, then she changes her bid to σ'_i in period $t + 1$ with probability $\epsilon^{\frac{1}{\gamma}}$, where γ is positive but

⁹The assumption of pessimism in the case of tied winning bids is not important. We could instead assume that they use the tie-breaking rule and expected utility to evaluate their payoffs over lotteries. In this case, had the bid of σ'_i made a buyer one of the M^* multiple highest bidders, then she believes that she would have won the auction with probability $\frac{1}{M^*}$. Then her payoff in the auction would have been $\pi(v_i, \sigma'_i, \sigma_{-i}) = \frac{1}{M^*}(v_i - \sigma'_i)$. The result for the second-price auction would not change with this alternative assumption. The result for the first-price auction studied in the next section would not change either for $v \neq \delta$.

smaller than any regret.¹⁰ That is, switches to bids for which there is no ex-post regret are least likely. All these events entailing switches in bids are independent across bidders and time.¹¹

The parameter η reflects certain frictions in the environment that prevent a trader from changing her bid by large amounts: a trader will not be allowed to change her bid by more than $\eta\delta$ in one period. A small η implies a high friction, whereas $\eta \geq \frac{1}{\delta}$ implies no friction at all.

We shall assume that a very small subset of bidders are obstinate. Call the set of obstinate bidders \mathcal{O} . An obstinate bidder never changes her bid, that is, she bids the same any time she is matched with a seller.

Let P_v denote the size of the population $\mathcal{P} \cap \mathcal{B}_v$. We assume that the size of any auction is small compared to these populations, that is, $M' \leq \min\{P_0, \dots, P_1\}$.

A state of the market in a period, denoted by ω , lists the number of bidders in \mathcal{P} with each valuation v who are using each bid σ in that period. Let $N_B(v, \sigma)$ denote the number of bidders in $\mathcal{P} \cap \mathcal{B}_v$ who use the bid σ . Therefore, $\omega = (N_B(v, \sigma))_{v \in V \& \sigma \in \Sigma}$. Let Ω denote the set of states of the market.

We do not include the distribution of bids among the obstinate bidders when defining the state of the market. This is because these bidders never change their bids whereas we are interested in analyzing the dynamic process of the bidders changing their bids using the learning rule described above. The only assumption we make about the distribution of bids among the obstinate bidders is that it is “interior,” that is, for every $\sigma \in \Sigma$, there exists at least one bidder in \mathcal{O} who uses the bid σ . Although the number of obstinate bidders required to satisfy this assumption is large for small values of δ , this number should be thought of as arbitrarily small relative to the number of buyers in \mathcal{P} .

Without the random adaptive learning rule, the state of the market does not change. In that case, every bidder in \mathcal{P} uses the same bid as the one

¹⁰This minimum (positive) regret exists since the sets of valuations and bids are finite; thus, γ is well defined.

¹¹The results are robust to any transformation λ of the probability of switching as follows: for a fixed v , let $\lambda_v : \mathfrak{R} \rightarrow \mathfrak{R}_{++}$ be a positive-monotonic transformation of $\Delta(v, \cdot, \cdot, \cdot)$. The result is unaffected if we instead assume that a bidder changes her bid from σ_i to $\sigma'_i \in \Sigma(\eta, \sigma_i)$ with probability $\epsilon^{\frac{1}{\lambda_{v_i}(\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i}))}}$. Note that the transformation λ_v need not be the same as $\lambda_{v'}$.

she used in the previous period. This trivial dynamics can thus be described as a Markov process on the state space Ω with the transition matrix equal to the identity matrix. Hence, all states of the market are absorbing and, as already mentioned, any prediction regarding which states will be more frequently visited in the long run depends on the initial state, that is, the initial distribution of bids in the population of bidders \mathcal{P} . By adding the random adaptive learning rule with $\epsilon > 0$, we have perturbed this dynamics, to get an irreducible and aperiodic Markov process, \mathcal{M}^ϵ , on the state space Ω . Hence there exists a unique stationary distribution μ^ϵ that approximates both the frequency with which a state is visited over a long horizon and the probability of being in a particular state at a point in time. Following Kandori, Mailath and Rob (1993) and Young (1993), the states in the support of $\lim_{\epsilon \rightarrow 0} \mu^\epsilon$ are called *stochastically stable* states. These are to be interpreted as our long run prediction, i.e., the states on which the system will spend a positive proportion of time in the very long run when the switches of bids are possible, but very unlikely events. Our attempt is to identify such states.

Before we do that, we list some facts that will be used repeatedly to identify the stochastically stable states in this section.

Facts 2.1. *Consider a second-price sealed-bid auction. The following statements are true for any η :*

- *A bidder with valuation v_i would have gained a positive amount had she increased her bid from σ_i to σ'_i in that auction, ceteris paribus,*
 - *only if such a bidder did not win the object in that auction using σ_i (i.e., $\sigma_i \leq \sigma_{-i}^*$),*
 - *and only if she would have been the unique highest bidder had she increased her bid to σ'_i (i.e., $\sigma'_i > \sigma_{-i}^*$).*
 - *Furthermore, such a gain is possible if and only if $\sigma_i < v_i$, and the maximum such possible gain for a bidder of valuation v_i is $v_i - \sigma_i$, which would happen when $\sigma_{-i}^* = \sigma_i$.*
- *A bidder with valuation v_i would have gained a positive amount had she decreased her bid from σ_i to σ'_i in that auction, ceteris paribus,*
 - *only if she won the object in that auction using σ_i (i.e., $\sigma_i \geq \sigma_{-i}^*$),*
 - *and only if she would not have been the unique highest bidder after reducing her bid to σ'_i (i.e., $\sigma'_i \leq \sigma_{-i}^*$).*

- Furthermore, such a gain is possible if and only if $\sigma_i > v_i$, and the maximum such gain for a bidder of valuation v_i is $\sigma_i - v_i$, which may happen when $\sigma_{-i}^* = \sigma_i$.

It follows from Facts 2.1 that a bidder with valuation v_i who bids $\sigma_i = v_i$ would not have gained a positive amount had she increased or decreased her bid in that auction. This gives a clue to the main finding of this section, which we proceed to establish formally.

Let $r(\omega, \omega')$ be the minimum resistance of going from state ω to state ω' . Let $N_B(\omega, v, \sigma)$ denote the number of bidders in $\mathcal{P} \cap \mathcal{B}_v$ who use the bid σ in state ω . Also, for every $v \in V$ and every $\sigma \in \Sigma$, define

$$\Delta^{**}(v, \sigma) = \max_{\sigma'_i \in \Sigma(\eta, \sigma), \sigma_{-i} \in \Sigma^{M'-1}} \{\Delta(v, \sigma, \sigma'_i, \sigma_{-i}), \gamma\}$$

The following is an auxiliary result used later in this section.

Lemma 2.2. *Suppose ω and ω' are such that $N_B(\omega', v, \sigma) < N_B(\omega, v, \sigma)$ for some $\sigma \in \Sigma$. Then $r(\omega, \omega') \geq \frac{1}{\Delta^{**}(v, \sigma)}$.*

Proof: In any transition from ω to ω' , it must be that at least one bidder from $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from σ , since $N_B(\omega', v, \sigma) < N_B(\omega, v, \sigma)$. Call this bidder i . The maximum this bidder i could have gained had she changed her bid from σ in any auction equals $\Delta^{**}(v, \sigma)$. Thus, $r(\omega, \omega') \geq \frac{1}{\Delta^{**}(v, \sigma)}$. ■

The main result of this section now follows:

Proposition 2.3. *If ω is stochastically stable, then $N_B(\omega, v, \sigma) = 0$ for all $\sigma \neq v$.*

Thus, our model predicts that if the sellers use a second-price sealed-bid auction, then only truthful bids are stable in the long run. Recall that it is a dominant strategy to bid truthfully in a second-price sealed-bid auction. Interestingly, even though the buyers in our model do not perform such rational calculations, over the long run, they converge to bids that are truthful. We will have more to say about this in section 5.

3 First-Price Auction

Consider now exactly the same setup as in the previous section, except that all the sellers in \mathcal{S}_0 use a first-price sealed-bid auction as the mechanism to

allocate goods. That is, all the bidders matched to a seller submit sealed bids to the seller. The bidder with the highest bid wins the object. If more than one bidder bid the highest, then a random tie-breaking rule is used by the seller. The bidder who wins the object pays a price equal to her bid.

We begin by stating some key facts for the analysis in this section:

Facts 3.1. *Consider a first-price sealed-bid auction. The following statements are true for any η :*

- *A bidder with valuation v_i would have gained a positive amount had she increased her bid from σ_i to σ'_i in that auction, ceteris paribus,*
 - *only if she did not win the object in that auction by using σ_i (i.e., $\sigma_i \leq \sigma_{-i}^*$),*
 - *only if she would have been the unique highest bidder had she increased her bid to σ'_i (i.e., $\sigma_{-i}^* < \sigma'_i$),*
 - *and only if $\sigma'_i < v_i$.*
 - *Furthermore, such a gain would be possible if and only if $\sigma_i < v_i - \delta$, and the maximum amount that the bidder with valuation v_i could have gained by such a bid increase equals $v_i - (\sigma_i + \delta)$, which happens when $\sigma_{-i}^* = \sigma_i$.*
- *A bidder with valuation v_i would have gained a positive amount had she decreased her bid from σ_i to σ'_i in that auction, ceteris paribus,*
 - *only if she won the object when she bid σ_i in that auction (i.e., $\sigma_i \geq \sigma_{-i}^*$).*
 - *If $\sigma_i \leq v_i$, then she would have been the unique highest bidder after reducing her bid to σ'_i (i.e., $\sigma'_i > \sigma_{-i}^*$), which implies that $\sigma_i > \delta$. In such a case ($\delta < \sigma_i \leq v_i$), the maximum amount that the bidder could have gained by decreasing her bid from σ_i equals $\min\{\eta\delta, \sigma_i - \delta\}$, which happens when $\sigma_{-i}^* = 0$.*
 - *If $\sigma_i > v_i$, then she would have gained irrespective of whether she would have won or lost the auction after reducing her bid to σ'_i . In this case, the maximum amount that the bidder could have gained by decreasing her bid from σ_i equals $\max\{\min\{\eta\delta, \sigma_i - \delta\}, \sigma_i - v_i\}$. Here $\min\{\eta\delta, \sigma_i - \delta\}$ is the highest possible price fall, which happens when $\sigma_{-i}^* = 0$ and $\sigma_i - v_i$ is the highest loss that she can avoid, which may happen when $\sigma_{-i}^* = \sigma_i$.*

For every $v \in V$, every $\sigma \in \Sigma$, and every $\sigma' \in \Sigma(\eta, \sigma)$, define

$$\Delta^*(v, \sigma, \sigma') = \max_{\sigma_{-i} \in \Sigma^{M'-1}} \{\Delta(v, \sigma, \sigma', \sigma_{-i}), \gamma\}.$$

What follows is another auxiliary lemma, used later to prove the main result of the section:

Lemma 3.2. *Suppose ω and ω' are such that a bidder in $\mathcal{P} \cap \mathcal{B}_v$ must change her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in any transition from ω to ω' . Then $r(\omega, \omega') \geq \frac{1}{\Delta^*(v, \check{\sigma}, \tilde{\sigma})}$.*

Proof: The maximum a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could have gained had she changed her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in any auction equals $\Delta^*(v, \check{\sigma}, \tilde{\sigma})$. Thus, $r(\omega, \omega') \geq \frac{1}{\Delta^*(v, \check{\sigma}, \tilde{\sigma})}$. ■

We now state the main result of the current section:

Proposition 3.3. *If ω is stochastically stable, then $N_B(\omega, v, \sigma) = 0$ for all σ satisfying any of the following inequalities:*

- $\sigma > \max\{v - \eta\delta, \frac{v+\delta}{2}\}$.
- $\eta\delta \leq \sigma < v - \eta\delta - \delta$.
- $\sigma < \min\{\eta\delta, \frac{v-\delta}{2}\}$.

The content of this proposition is best expressed by the following corollary:

Corollary 3.4. *If ω is stochastically stable, then $N_B(\omega, v, \sigma) > 0$ implies that:*

- if $\eta < \frac{v-\delta}{2\delta}$, $\sigma \in [v - \eta\delta - \delta, v - \eta\delta]$;
- and if $\eta \geq \frac{v-\delta}{2\delta}$, $\sigma \in [\frac{v-\delta}{2}, \frac{v+\delta}{2}]$.

Proof: It follows from rearranging the inequalities in the statement of Proposition 3.3. ■

The above corollary tells us that the bids outside the identified intervals will not be used by the bidders in $\mathcal{P} \cap \mathcal{B}_v$ in any stochastically stable state. Note that the answer obtained depends on the friction parameter η . If η is very small (e.g., $\eta = 1$), there is a strong rigidity to change one's bid. Then, the bid that every bidder ends up using almost all the time hovers around one's true valuation, i.e., bid shading is not a significant long run prediction

in this case. On the other hand, if the friction to change bids is insignificant (corresponding to large values of η , say $\eta = \frac{1}{\delta}$), a unique form of bid shading is identified as the long run prediction. Indeed, every bidder of valuation v is very likely to use a bid that approximates as close as possible exactly half of his valuation, i.e., $\frac{v}{2}$. For intermediate values of η , the highest valuations almost all the time shade their bid as much as the friction allows them, while the low valuations' most often used bid is again half of their valuation. To illustrate the intermediate case, if $\delta = 0.01$ is the monetary unit, $n = 50$ is the grid width of valuations, and $\eta = 30$ is the friction parameter, then those bidders whose valuation exceeds 0.60 are very likely to shade their bid by 30 or 31 cents with respect to their valuation, while those with valuations no greater than 0.60 bid half of their valuation most of the time in the very long run.

Thus, the model predicts that if the bidders face an auction environment with high frictions to change their bids, then they are likely to bid very close to their valuations in the long run. The amount of shading by the bidders increases as the frictions diminish. However, there is an upper bound on the shading that bidders will undertake in the long run: even if bidders face a frictionless auction on the bid updating, they are unlikely to bid below $\frac{v-\delta}{2}$.

4 k -Double Auction

There is a single indivisible object. A seller of that indivisible object values it at $v_s \in V = \{0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1\}$, and it is her private information. A buyer of that indivisible object values it at $v_b \in V$, and it is also her private information.

For every $v \in V$, there exist a large number of buyers and a large number of sellers with their valuations equal to v . Let \mathcal{B}_v denote the population of buyers with their valuations equal to v . Let \mathcal{S}_v denote the population of sellers with their valuations equal to v .

In each period t , a finite number M_t of buyers and the same number of sellers are selected randomly and independently from $\mathcal{B} = \bigcup_{v \in V} \mathcal{B}_v$ and $\mathcal{S} = \bigcup_{v \in V} \mathcal{S}_v$, respectively. These selected buyers and sellers are then randomly paired to transact the indivisible object. Note that M_t itself can be a random variable so that the number of matches can be different in different periods.¹²

¹²In that case, we assume that M_t and $M_{t'}$ have the same distribution to ensure that the dynamic process is time homogeneous.

The trading mechanism is a k -double auction, $k \in [0, 1]$, in which both players submit sealed bids. Namely, if the buyer bids p_b and the seller bids p_s , then the object is traded if and only if $p_b \geq p_s$ at a price equal to the weighted average of the two bids, $kp_b + (1 - k)p_s$. If the object is traded at price p , the buyer's payoff is $v_b - p$ while the seller's payoff is $p - v_s$; otherwise, they each get a payoff of 0.

A player can bid any number $\sigma \in \Sigma = \{0, \delta, 2\delta, \dots, 1\}$. Again, assume that δ is small enough so that there exists an integer N such that $N\delta = \frac{1}{n}$.

Each player in this market can be identified by her bidding rule, which specifies her unique bid in the event she is matched. So when we say that a player's bid in period t is σ , we mean that she will bid σ in case she is matched in period t .

As in the previous two models, we first assume that the players in this market have complete inertia with respect to their bids, that is, in each period, every player uses the same bid as in the previous period. Thus the distribution of bids in the market during any period is exactly the initial distribution. Therefore, without adding more structure, the model has no predictive power. To answer which of the large number of possible distributions of bids will be observed in the long run, we add a perturbation to the model using a random adaptive learning rule similar to the one defined earlier.

4.1 Random Adaptive Learning Rule

Let the subscript i denote a buyer or a seller, that is, $i = b, s$. A subset of players, \mathcal{P} , adapt their bid from period t to $t + 1$ using the following random adaptive learning rule: suppose in period t a player in \mathcal{P} with valuation v_i is using the bid σ_i .

- If this player is not matched in period t , then she does not change her bid if matched in period $t + 1$.
- If this player is matched in period t , then let $\pi_i(v_i, \sigma_i, \sigma_{-i})$ be her payoff in that double auction, where σ_{-i} is the other player's bid. Pick any $\sigma'_i \in \Sigma(\eta, \sigma_i) \equiv \{\sigma : 0 < |\sigma - \sigma_i| \leq \eta\delta\}$, where $\eta \geq 1$. Had she bid σ'_i instead of σ_i in that double auction, *ceteris paribus*, her payoff would have been $\pi_i(v_i, \sigma'_i, \sigma_{-i})$.

Define $\Delta_i(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) = \pi_i(v_i, \sigma'_i, \sigma_{-i}) - \pi_i(v_i, \sigma_i, \sigma_{-i})$.

As in previous sections, when $\Delta_i(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) > 0$, we refer to this magnitude as valuation v_i 's *regret* from bidding σ_i instead of σ'_i against σ_{-i} . Thus, if $\Delta_i(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) > 0$, we shall assume that this player changes her bid to σ'_i in period $t + 1$ with probability $\epsilon^{\frac{1}{\Delta_i(v_i, \sigma_i, \sigma'_i, \sigma_{-i})}}$. That is, although inertia continues to be the main driving force of her behavior, in the unlikely event that her bid is adjusted, the probability of switching to another bid is increasing in the corresponding regret. However, if $\Delta_i(v_i, \sigma_i, \sigma'_i, \sigma_{-i}) \leq 0$, then she changes her bid to σ'_i in period $t + 1$ with probability $\epsilon^{\frac{1}{\gamma}}$, where γ is positive but smaller than any regret. These switches are not being justified by any regret, and we assume that the probability of such switches is smaller. All these events entailing switches in bids are independent across players and time.¹³

Also as in the previous models, we assume that there is a very small subset of players who are of the obstinate type, denoted by \mathcal{O} . An obstinate player never changes her bid, that is, she bids the same value any time she is matched. We also assume that the distribution of bids among the obstinate buyers and sellers is “interior,” that is, for every $\sigma \in \Sigma$, there exist at least one buyer and one seller in \mathcal{O} who use the bid σ .

A state of the market in a period, denoted by ω , lists the number of buyers and sellers in \mathcal{P} with each valuation $v \in V$ who are using each bid $\sigma \in \Sigma$ in that period. Let $N_B(v_b, \sigma_b)$ denote the number of buyers in $\mathcal{P} \cap \mathcal{B}_{v_b}$ who use the bid σ_b . Let $N_S(v_s, \sigma_s)$ denote the number of sellers in $\mathcal{P} \cap \mathcal{S}_{v_s}$ who use the bid σ_s . Therefore, $\omega = (N_B(v_b, \sigma_b), N_S(v_s, \sigma_s))_{v_b, v_s \in V \& \sigma_b, \sigma_s \in \Sigma}$. As in the previous models, we do not include the distribution of bids among the obstinate players while defining the state of the market. Let Ω denote the set of states of the market.

As before, without the random adaptive learning rule, the state of the market does not change. After we introduce it, we have an irreducible and aperiodic perturbed Markov process, and we proceed to identify its stochastically stable states.

Facts 4.1. *Consider a k -double auction. The following statements are true for any η .¹⁴*

¹³As in the previous models, the result is robust to any positive-monotonic transformation of the probability of switching.

¹⁴We list these facts only for a seller since we write down the proof of the main result

- A seller with valuation v_s would have gained a positive amount had she decreased her bid from σ_s to σ'_s in that double auction, *ceteris paribus*,
 - only if she did not trade with the buyer when she bid σ_s (i.e., $\sigma_s > \sigma_b$),
 - only if she would have traded with the buyer had she decreased her bid (i.e., $\sigma'_s \leq \sigma_b$),
 - only if $k\sigma_b + (1 - k)\sigma'_s > v_s$.
 - Furthermore, such a gain would be possible if and only if $\sigma_s > v_s + \delta$, and the maximum amount that the seller with valuation v_s could have gained by decreasing her bid from σ_s in that double auction equals $\sigma_s - \delta - v_s$, which happens if and only if $\sigma_b = \sigma_s - \delta$.
- If a seller with valuation v_s would have gained a positive amount had she increased her bid from σ_s to σ'_s in that double auction, *ceteris paribus*,
 - then it must be that she traded with the buyer when she bid σ_s (i.e., $\sigma_s \leq \sigma_b$).
 - If $k\sigma_b + (1 - k)\sigma_s \geq v_s$, then it must be that she would have also traded with the buyer after increasing her bid (i.e., $\sigma'_s \leq \sigma_b$). In this case, the maximum such gain for this seller from increasing her bid from σ_s equals $(1 - k) \min\{\eta\delta, 1 - \sigma_s\}$, which happens when $\sigma_b = 1$.
 - If $k\sigma_b + (1 - k)\sigma_s < v_s$, then it must be that $\sigma_s < v_s$. In this case, she would have gained irrespective of whether she would have traded or not after increasing her bid. Moreover, the maximum such gain for this seller from increasing her bid from σ_s equals $\max\{(1 - k) \min\{\eta\delta, 1 - \sigma_s\}, v_s - \sigma_s\}$. Here, $(1 - k) \min\{\eta\delta, 1 - \sigma_s\}$ is the highest possible price increase, which happens when $\sigma_b = 1$, and $v_s - \sigma_s$ is the highest loss that she can avoid, which happens if and only if $\sigma_b = \sigma_s$.

Let $N_B(\omega, v, \sigma)$ denote the number of buyers in $\mathcal{P} \cap \mathcal{B}_v$ with valuation v who bid σ in state ω . Similarly, define $N_S(\omega, v, \sigma)$. Also, for $i = b, s$, define

only for the sellers.

- for every $v \in V$, every $\sigma \in \Sigma$, and every $\sigma' \in \Sigma(\eta, \sigma)$,

$$\Delta_i^*(v, \sigma, \sigma') = \max_{\sigma_{-i} \in \Sigma} \{\Delta_i(v, \sigma, \sigma', \sigma_{-i}), \gamma\};$$

- for every $v \in V$, and every $\sigma \in \Sigma$,

$$\Delta_i^{**}(v, \sigma) = \max_{\sigma'_i \in \Sigma(\eta, \sigma), \sigma_{-i} \in \Sigma} \{\Delta_i(v, \sigma, \sigma'_i, \sigma_{-i}), \gamma\}$$

The next two lemmas provide the lower bound on the resistance of a transition between two states.

Lemma 4.2. *Suppose ω and ω' are such that $N_B(\omega', v, \sigma) < N_B(\omega, v, \sigma)$ for some $\sigma \in \Sigma$. Then $r(\omega, \omega') \geq \frac{1}{\Delta_b^{**}(v, \sigma)}$. Similarly, if ω and ω' are such that $N_S(\omega', v, \sigma) < N_S(\omega, v, \sigma)$ for some $\sigma \in \Sigma$, then $r(\omega, \omega') \geq \frac{1}{\Delta_s^{**}(v, \sigma)}$.*

Proof: It must be that at least one buyer in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from σ in the transition from ω to ω' since $N_B(\omega', v, \sigma) < N_B(\omega, v, \sigma)$. The maximum this buyer could have gained had she changed her bid from σ in any double auction equals $\Delta_b^{**}(v, \sigma)$. Thus, $r(\omega, \omega') \geq \frac{1}{\Delta_b^{**}(v, \sigma)}$. We can similarly prove the second part of the lemma. ■

Lemma 4.3. *Suppose ω and ω' are such that a buyer in $\mathcal{P} \cap \mathcal{B}_v$ must change her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in any transition from ω to ω' . Then $r(\omega, \omega') \geq \frac{1}{\Delta_b^*(v, \check{\sigma}, \tilde{\sigma})}$. Similarly, if ω and ω' are such that a seller in $\mathcal{P} \cap \mathcal{S}_v$ must change her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in any transition from ω to ω' , then $r(\omega, \omega') \geq \frac{1}{\Delta_s^*(v, \check{\sigma}, \tilde{\sigma})}$.*

Proof: The maximum a buyer in $\mathcal{P} \cap \mathcal{B}_v$ could have gained had she changed her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in any double auction equals $\Delta_b^*(v, \check{\sigma}, \tilde{\sigma})$. Thus, $r(\omega, \omega') \geq \frac{1}{\Delta_b^*(v, \check{\sigma}, \tilde{\sigma})}$. The other part of the lemma can be proved similarly. ■

We can now state the main result of this section:

Proposition 4.4. *If ω is stochastically stable, then*

- $N_B(\omega, v_b, \sigma_b) = 0$ for all σ_b satisfying any of the following inequalities:

- $\sigma_b > \max\{v_b - k\eta\delta, \frac{1}{1+k}v_b\}$.
- $\eta\delta - \delta \leq \sigma_b < v_b - k\eta\delta - \delta$.
- $\sigma_b < \min\{\eta\delta - \delta, \frac{1}{1+k}v_b - \delta\}$.

- $N_S(\omega, v_s, \sigma_s) = 0$ for all σ_s satisfying any of the following inequalities:
 - $\sigma_s < \min\{v_s + (1 - k)\eta\delta, \frac{1}{2-k}v_s + \frac{1-k}{2-k}\}$.
 - $v_s + (1 - k)\eta\delta + \delta < \sigma_s \leq 1 + \delta - \eta\delta$.
 - $\sigma_s > \max\{1 + \delta - \eta\delta, \frac{1}{2-k}v_s + \frac{1-k}{2-k} + \delta\}$.

Again, the message of the former proposition is best brought out by the following corollary:

Corollary 4.5. *If ω is stochastically stable, then*

- $N_B(\omega, v_b, \sigma_b) > 0$ implies that:
 - if $\eta < \frac{v_b}{\delta(1+k)}$, $\sigma_b \in [v_b - k\eta\delta - \delta, v_b - k\eta\delta]$,
 - and if $\eta \geq \frac{v_b}{\delta(1+k)}$, $\sigma_b \in [\frac{1}{1+k}v_b - \delta, \frac{1}{1+k}v_b]$.
- $N_S(\omega, v_s, \sigma_s) > 0$ implies that:
 - if $\eta < \frac{1-v_s}{\delta(2-k)}$, $\sigma_s \in [v_s + (1 - k)\eta\delta, v_s + (1 - k)\eta\delta + \delta]$,
 - and if $\eta \geq \frac{1-v_s}{\delta(2-k)}$, $\sigma_s \in [\frac{1}{2-k}v_s + \frac{1-k}{2-k}, \frac{1}{2-k}v_s + \frac{1-k}{2-k} + \delta]$.

Proof: It follows from rearranging the inequalities in the statement of Proposition 4.4. ■

The above corollary tells us that the bids outside the identified intervals will not be used by the traders in the non-obstinate populations. Again, note how the answer depends on the friction parameter η . If frictions to change one's bid are very large (e.g., as in $\eta = 1$), a buyer of valuation v_b is extremely likely to shade her bid between $(k+1)$ and k times δ , and a seller of valuation v_s is extremely likely to bump her ask between $(1-k)$ and $(2-k)$ times δ . If, on the other hand, frictions are small (corresponding to large values of η , say $\eta = \frac{1}{\delta}$), the long run behavior of each of these populations of traders is also well identified. Indeed, while a buyer of valuation v_b is then very likely to bid $\frac{v_b}{1+k}$, each seller of valuation v_s would ask most of the time an amount $\frac{v_s+1-k}{2-k}$. For intermediate values of η , the high surplus traders (high valuation buyers and low valuation sellers) bid according to the first items in the corollary, and they misrepresent their valuations with their bids and asks in a lower percentage than do low surplus traders, who bid according to the second items in it. To illustrate this intermediate case, suppose again that $\delta = 0.01$ is the indivisible money unit, $n = 50$ is the grid width of valuations, and

$\eta = 30$ is the friction parameter, and consider the case $k = \frac{1}{2}$. Then, buyers with valuation $v_b \leq 0.44$ are very likely to offer the bid closest to $\frac{2}{3}v_b$, while buyers with valuation $v_b \geq 0.46$ will bid shading their valuation by either 15 or 16 cents. On the other hand, sellers with valuation $v_s \geq 0.56$ will ask for the amount closest to $\frac{1}{3} + \frac{2}{3}v_s$, while sellers with valuation $v_s \leq 0.54$ will ask for 15 or 16 cents more than their valuation in any stochastically stable state.

It also follows from Corollary 4.5 that if $k = 0$, then the buyers in $\mathcal{P} \cap \mathcal{B}_{v_b}$ are likely to use only the bids in $[v_b - \delta, v_b]$ in the long run irrespective of η . Thus, when the buyer's bid does not determine the price of the good, all the buyers will eventually bid within δ of their valuation. It is interesting to note that it is a dominant strategy for a buyer to bid equal to her valuation in a k -double auction when $k = 0$. The buyers in our model will converge to almost truthful bids over the long run when $k = 0$ but without using similar rational calculations. Similarly, if $k = 1$, then the sellers in $\mathcal{P} \cap \mathcal{S}_{v_s}$ are likely to use only the bids in $[v_s, v_s + \delta]$ in the long run irrespective of the friction η . That is, when the seller's bid does not determine the price of the good, all the sellers will eventually bid within δ of their valuation. Again, it is a dominant strategy for a seller to bid equal to her valuation in a k -double auction when $k = 1$.

On the other hand, we see that when $k > 0$, the set of bids that are likely to be used by the buyers in the long run will depend on η . Starting from $[v_b - (1 + k)\delta, v_b - k\delta]$ when $\eta = 1$, the set of bids that are likely to be used by the buyers in the long run keeps "shifting down" as the friction to switch one's bid becomes less severe. Once $\eta \geq \frac{v_b}{(1+k)\delta}$, then only the bids in $[\frac{1}{1+k}v_b - \delta, \frac{1}{1+k}v_b]$ are likely to be used by the buyers in the long run irrespective of η . Similarly, if $k < 1$, then the bids that are likely to be used in the long run by the sellers depend on the friction as well. When $\eta = 1$, then they are likely to use only the bids in $[v_s + (1 - k)\delta, v_s + (2 - k)\delta]$ in the long run. As the friction is removed when one increases η , this set of bids keeps "shifting up". However, there is a limit to how much this set will "shift up". Once $\eta \geq \frac{1-v}{(2-k)\delta}$, then only the bids in $[\frac{1}{2-k}v_s + \frac{1-k}{2-k}, \frac{1}{2-k}v_s + \frac{1-k}{2-k} + \delta]$ are likely to be used by the sellers in the long run regardless of the value of η .

5 Discussion

To discuss our results, we shall begin by explaining the intuition of how they are obtained. The bids that will be used in the stochastically stable states (approximately) satisfy the following *balance condition*: *at these bids, a player's maximal gain from increasing her bid equals her maximal gain from decreasing it.*

In the second-price auction, if a bidder with valuation v uses a bid $\sigma > v$, there is no gain at all from increasing her bid over σ , while the maximal gain from decreasing it is $\sigma - v > 0$ (see Facts 2.1). In this case, such a bid will be adjusted downwards to find the path of least resistance. On the other hand, if she uses a bid $\sigma < v$, the maximal gain from increasing her bid is $v - \sigma > 0$ while there is no gain at all from decreasing it (see again Facts 2.1). Therefore, such a bid will be adjusted upwards. From the combination of both arguments, it follows that in the long run one will see almost all the time the bid $\sigma = v$. Furthermore, note how the two maximal gains are equal only at $\sigma = v$, where both vanish.

In the first-price auction, if a bid $\sigma \geq v$ is being used, there is no gain whatsoever from increasing it, while the maximum gain from decreasing it, consisting of avoiding a loss, is $\max\{\min\{\eta\delta, \sigma - \delta\}, \sigma - v\}$ (see Facts 3.1). This means that no bids $\sigma > v$ will be observed a positive proportion of time in the long run, since in this case the gain from decreasing one's bid is strictly positive. Further, the same is true for $\sigma = v$ whenever $v > \delta$. (If $v = 0$ or $v = \delta$, bidding one's valuation will still be stochastically stable.) What remains now is to study the dynamic forces for bids below one's valuation. If a bidder uses a bid $\sigma < v$, the maximal gain from increasing it over σ , which will happen when she turns from losing to winning the auction, is $v - \sigma - \delta$ (see Facts 3.1). On the other hand, the maximal gain from decreasing her bid under σ , which results in the highest possible price savings, is $\min\{\eta\delta, \sigma - \delta\}$ (see again Facts 3.1). Thus, for $v > \delta$, the two are equal only at $\sigma = \max\{v - \eta\delta - \delta, \frac{v}{2}\}$.

Consider now the k -double auction. Suppose a seller uses an asking bid $\sigma_s < v_s$. Then, there is no possible gain associated with decreasing it, while the maximal gain from increasing it, resulting in the highest possible price increase or the avoidance of a loss, is $\max\{(1 - k) \min\{\eta\delta, 1 - \sigma_s\}, v_s - \sigma_s\}$ (see Facts 4.1). It follows that no asking bid $\sigma_s < v_s$ will be ever part of the long run prediction, as the dynamic forces will push to increase it. This already pins down the long run behavior of a seller with valuation $v_s = 1$.

In addition, for bids $\sigma_s \geq v_s$, the maximal gain a seller gets from increasing her asking bid over σ_s , now simply consisting of the highest price increase, is $(1-k) \min\{\eta\delta, 1-\sigma_s\}$. On the other hand, the maximal gain of a seller from decreasing her ask below σ_s , turning a situation of no-trade into trade, is $\sigma_s - \delta - v_s$ if $\sigma_s > v_s$, while no gain is possible from such a decrease if $\sigma_s = v_s$ (see Facts 4.1). Thus, for $v_s < 1$, the two maximal payoff gains are equal only at $\sigma_s = \min\{v_s + (1-k)\eta\delta + \delta, \frac{1}{2-k}v_s + \frac{1-k}{2-k} + \frac{\delta}{2-k}\}$. We can similarly argue that the balance condition for the buyers is never possible for $\sigma_b > v_b$ since there is no possible gain associated with increasing one's bid, while a positive gain exists if one decreases it. Further, when $v_b > 0$, the balance condition only holds at $\sigma_b = \max\{v_b - k\eta\delta - \delta, \frac{1}{1+k}v_b - \frac{\delta}{1+k}\}$.

To sum up, the stability of the bids that are likely to be used in stochastically stable states can thus be seen as stemming from a balance between two opposing forces. Indeed, it is not difficult to see that the proofs we have constructed implicitly use this balance condition. Basically, we argue that the gain that a player in some population could have got in a particular transition along the added branch to the rooted tree is greater than the gain that a player in that population could have got on any transition along the deleted branch. The gain along the added branch is always at least the maximal gain from increasing (decreasing) the bid, whereas the gain on the deleted branch is at most approximately the maximal gain from decreasing (increasing) the bid. The last interaction of a player makes her ponder whether she has an incentive to increase, decrease or not change her bid. The strength of this incentive is a function of the two opposing forces: the maximal gain she could have got by increasing her bid and the maximal gain she could have got by decreasing it. If the former force dominates, then it is much more likely to see the player increasing her bid, whereas if the latter dominates, the system will move much more often toward states in which she decreases her bid. Only at those bids that are used in stable states do the two opposing forces balance.

Another way to see our results is that stochastically stable bids minimize the maximal regret. Recall that we can interpret $\max\{\Delta(v_i, \sigma_i, \sigma'_i, \sigma_{-i}), 0\}$ as the regret of player i associated with playing σ_i instead of σ'_i in the last period. For all $v \in V$ and $\sigma \in \Sigma$ define:

$$R(v, \sigma) = \max_{\sigma'_i \in \Sigma(\eta, \sigma), \sigma_{-i} \in \Sigma_{-i}} \max\{\Delta(v, \sigma, \sigma'_i, \sigma_{-i}), 0\}.$$

Thus, $R(v, \sigma)$ is the maximal regret that a player with valuation v can

have if she plays σ in any period. We shall argue next that, in all three models, stochastically stable bids minimize this maximal regret, that is, stochastically stable bids belong to $\arg \min_{\sigma \in \Sigma} R(v, \sigma)$.¹⁵

In the second-price auction, $R(v, \sigma) = |v - \sigma|$ (see Facts 2.1). Thus, bidding truthfully minimizes the maximal regret and we have already seen that it is also the unique stochastically stable bid.

In the first-price auction, $R(v, \sigma) = \max\{\min\{\eta\delta, \sigma - \delta\}, \sigma - v\}$ if $\sigma \geq v$ whereas $R(v, \sigma) = \max\{\min\{\eta\delta, \sigma - \delta\}, v - \sigma - \delta\}$ if $\sigma < v$. (Recall our discussion of the balanced condition, earlier in this section.) For simplicity assume that δ is small enough so that $R(v, \sigma) \approx \max\{\min\{\eta\delta, \sigma\}, |v - \sigma|\}$. Thus for $v < 2\eta\delta$, $R(v, \sigma)$ is minimized at bids that approximately equal $\frac{v}{2}$ whereas for $v \geq 2\eta\delta$, $R(v, \sigma)$ is minimized at bids that approximately lie in the interval $[v - \eta\delta, v + \eta\delta]$. Corollary 3.4 tells us that when $v < 2\eta\delta$, the stochastically stable bids are approximately equal to $\frac{v}{2}$ whereas for $v \geq 2\eta\delta$, stochastically stable bids are approximately equal to $v - \eta\delta$. Thus, stochastically stable bids minimize the maximal regret.

In the k -double auction, $R_s(v_s, \sigma_s) = \max\{(1 - k) \min\{\eta\delta, 1 - \sigma_s\}, v_s - \sigma_s\}$ if $\sigma_s < v_s$ whereas $R_s(v_s, \sigma_s) = \max\{(1 - k) \min\{\eta\delta, 1 - \sigma_s\}, \sigma_s - \delta - v_s\}$ if $\sigma_s \geq v_s$. (Again, the discussion on the balance condition shows this.) Again for simplicity assume that δ is small so that $R_s(v_s, \sigma_s) \approx \max\{(1 - k) \min\{\eta\delta, 1 - \sigma_s\}, |v_s - \sigma_s|\}$. Thus for $1 - v_s < (2 - k)\eta\delta$, $R_s(v_s, \sigma_s)$ is minimized at bids that approximately equal $\frac{1}{2-k}v_s + \frac{1-k}{2-k}$ whereas for $1 - v_s \geq (2 - k)\eta\delta$, $R_s(v_s, \sigma_s)$ is minimized at bids that approximately lie in the interval $[v_s - (1 - k)\eta\delta, v_s + (1 - k)\eta\delta]$. Corollary 4.5 tells us that when $1 - v_s < (2 - k)\eta\delta$, the stochastically stable bids are approximately equal to $\frac{1}{2-k}v_s + \frac{1-k}{2-k}$ whereas for $1 - v_s \geq (2 - k)\eta\delta$, the stochastically stable bids are approximately equal to $v_s + (1 - k)\eta\delta$. A similar argument works for the buyer. Thus, the stochastically stable bids in the k -double auction also minimize the maximal regret.

Hyafil and Boutilier (2004) propose minimizing the maximal regret as a decision criterion when the players do not have any quantifiable prior beliefs over the type space – they refer to this environment as one with strict type uncertainty. They define a strategy of a player to be a minimax best response to a given strategy profile of other players if the action of each type of that

¹⁵Because of the friction η , this is not an “if and only if” statement, that is, there will exist bids that minimize the maximal regret but they are not stochastically stable. However, if $\eta \geq \frac{1}{\delta}$, then one can make it an “if and only if” statement.

player minimizes her maximal regret for any ex-post realization of types of the other players. They go on to define a minimax-regret equilibrium as the profile of strategies such that the strategy of each player is a minimax best response to the strategy profile of the other players. We now argue that the profiles of stable bids define a minimax-regret equilibrium appropriately redefined when the model has obstinate types. In our models, a type of a player is a pair describing her valuation and whether she is obstinate or not. Therefore, let's assume that the environment is that of strict type uncertainty over this extended type space. We have shown above that the stable bids of each valuation type of each non-obstinate type player minimizes her maximal regret with respect to any realization of the bid profile of the other players. Along with the “interiority” assumption on the set of bids of the obstinate type players, this implies that the strategy of every non-obstinate type player in the stable states is a minimax best response to the strategy profile of the other players in the stable states. Thus, the strategy profile of non-obstinate type players in the stable states defines a minimax-regret equilibrium.¹⁶

Our agents are far from being able to play an equilibrium in the corresponding static Bayesian game, let alone its dynamic counterpart, since there are no prior beliefs that are common knowledge among them. In this sense, it is as if they play the game “in the veil of ignorance.” We argue now that, in the three models we have studied, one can find a belief such that the stable bid is an optimal response to that belief, and that belief is that each bid chosen by the opponents is equally likely. We shall refer to this as “uniform belief rationalization.” A bid rationalized in this way satisfies the first-order condition of payoff maximization given uniform beliefs, i.e., marginal gain must equal marginal loss at that bid. As we shall see, the reason why stable bids can be “uniform-belief rationalized” is that the balance condition that is satisfied at such a bid (maximal gain from increasing the bid equal to the maximal gain from decreasing it) is exactly the same as the first-order condition.

The balance condition is $v - \sigma = 0$ in the second-price auction model. This implies that only truthful bid is stable. Clearly, truthful bid is also “rationalizable” for any belief, since it is a dominant strategy to bid truthfully in a second-price sealed-bid auction.

¹⁶Of course, the bid of the obstinate type player is not a minimax best response to the strategy profile of the other players. It is standard to define only the strategy profile of the non-obstinate types as an equilibrium when the model has obstinate types. For instance, see Abreu and Gul (2000).

Thus, the argument of “rationalization” for the second-price auction is fairly robust, independent of the friction η . However, for most of the analysis in the first-price and k -double auctions, we assume in the ensuing discussion that $\eta \geq \frac{1}{\delta}$, so that effectively the friction to change one’s bid in our dynamic models has been removed. Recall that also in these trading procedures, whenever a trader has a dominant strategy, our conclusions for her are independent of η .

The bid $\frac{v}{2}$ is stable for the bidders in the first-price sealed-bid auction when $\eta \geq \frac{1}{\delta}$. In this case, the balance condition is

$$v - \sigma - \delta = \sigma - \delta \quad \implies \quad v - \sigma = \sigma.$$

The bid of $\frac{v}{2}$ can also be “rationalized” in the following sense. Recall that bidding half of one’s valuation is the equilibrium strategy of a two-bidder first-price auction where the valuations are independently drawn from a uniform distribution and the seller’s reservation price is zero. In this case, the belief of bidder i is that the other bidder is bidding half of her valuation. So the first-order condition of that problem is

$$(v - \sigma_i)g(\sigma_i) = G(\sigma_i) \quad \implies \quad (v - \sigma_i)2 = 2\sigma_i \quad \implies \quad v - \sigma_i = \sigma_i.$$

(Here, g and G are the density and distribution functions of the other bidder’s bid.) In our model, a bidder’s incentive to switch her bid is based on an ex-post calculation. So all she cares about is the ex-post realization of σ_{-i}^* . In contrast, a rational bidder in a first-price auction cares about the expected value of σ_{-i}^* , which in particular depends on the number of bidders in the auction. Equivalently, one can argue the “rationalization” of our stable bid if each bidder believes that the distribution of σ_{-i}^* is uniform over an interval $[0, \kappa]$ with $0.5 \leq \kappa \leq 1$ – for example, if she believes there is no shading at all when the underlying distribution is uniform.

We turn now to k -double auctions. If $k = 1$, then the truthful bid is stable for the sellers in the k -double auction. The balance condition at $\sigma_s = v_s$ in this case is

$$\max\{0, v_s - \sigma_s\} = 0 \quad \implies \quad v_s - \sigma_s = 0.$$

Bidding truthfully can also be “rationalized” since doing so is a dominant strategy for a seller in the k -double auction when $k = 1$. All these conclusions are similar for a buyer when $k = 0$.

The bid $\frac{1}{2-k}v_s + \frac{1-k}{2-k}$ is stable for all the sellers in the k -double auction when $\eta \geq \frac{1}{\delta}$. In this case, the balance condition is

$$\sigma_s - \delta - v_s = (1 - k)(1 - \sigma_s).$$

For small δ , we can write it as

$$\sigma_s - v_s \approx (1 - k)(1 - \sigma_s).$$

Bidding $\frac{1}{2-k}v_s + \frac{1-k}{2-k}$ can be “rationalized” since if a seller with valuation v_s believes that the bid of the buyer is distributed uniformly on $[0, 1]$, then that bid is optimal for her. For that belief, the first-order condition of payoff maximization for the seller is

$$(v_s - \sigma_s)\hat{g}(\sigma_s) = (1 - k)(1 - \hat{G}(\sigma_s)) \quad \implies \quad v_s - \sigma_s = (1 - k)(1 - \sigma_s).$$

(Here, \hat{g} and \hat{G} are the density and distribution functions of the buyer’s bid.) Note that if the buyer’s bidding function is linear, then $\hat{G}(\sigma_s) = \hat{g}(\sigma_s)\sigma_s + \text{constant}$. Therefore, the slope of the seller’s bidding function is always $\frac{1}{2-k}$ as long as the buyer’s bidding function is linear. Similarly, the bid $\frac{1}{1+k}v_b$ is stable for all the buyers in the k -double auction when $\eta \geq \frac{1}{\delta}$. It is also “rationalized” since if a buyer with valuation v_b believes that the bid of the seller is distributed uniformly on $[0, 1]$, then the bid that is optimal for her is exactly $\frac{1}{1+k}v_b$.

Finally, it is interesting to compare the linear equilibrium of the k -double auction when the valuation distributions are uniform over the interval $[0, 1]$ with the behavior uncovered by our evolutionary process. The beliefs described to “rationalize” the bids uncovered in our analysis are that no misrepresentation of valuations takes place when traders play the game. Correspondingly, each side plays “tougher” than in equilibrium. The reader will recall that the linear equilibrium strategies for the case of uniform distributions are $\frac{1}{2-k}v_s + \frac{1-k}{2}$ for the seller and $\frac{1}{1+k}v_b + \frac{k(1-k)}{2(1+k)}$ for the buyer. So, in our model, both players are “tougher” than they would be in equilibrium, and the amount of extra toughness over equilibrium bids on each side is $\frac{k(1-k)}{2(2-k)}$ for the seller, and $\frac{k(1-k)}{2(1+k)}$ for the buyer. If $k = 1/2$, this extra toughness exhibited by each side is the same on both sides. The highest toughness by the seller is when $k = 2 - \sqrt{2}$ and by the buyer is when $k = \sqrt{2} - 1$. On the other hand, if either $k = 0$ or $k = 1$, our model supports equilibrium play on both sides – no extra toughness is found.

6 Conclusion

We have applied stochastic stability to the analysis of bidding behavior in private-values auctions and double auctions. The assumed bidding behavior has a strong component of inertia, and it is modified in the direction of regrets. The results obtained are sharp, identifying a unique bid to be used by each type most of the time in the very long run. Such a bid is the one that minimizes the maximal regret. It would be interesting to explore general Bayesian games and see to what extent our conclusions can be generalized.

7 Proofs

Proof of Proposition 2.3:

Case 1: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some $\sigma > v$. Let $\hat{\sigma}$ be the highest $\sigma' > v$ such that $N_B(\omega, v, \sigma') > 0$.

Consider state ω' such that the distribution of bids in ω' is exactly the same as in ω except that $N_B(\omega', v, \hat{\sigma}) = N_B(\omega, v, \hat{\sigma}) - 1$ and $N_B(\omega', v, \hat{\sigma} - \delta) = N_B(\omega, v, \hat{\sigma} - \delta) + 1$. Any transition from ω to ω' involves at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from $\hat{\sigma}$ to something else. Pick any bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid $\hat{\sigma}$ in state ω . Call her bidder i . In state ω , there is a positive probability that bidder i is matched with a seller, a single bidder from \mathcal{O} who bids $\hat{\sigma}$ and other $M' - 2$ bidders from $\mathcal{P} \cap \mathcal{B}_v$ (here we are using the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). Since $\sigma_{-i}^* = \hat{\sigma}$, there is a positive probability that bidder i wins the object and gets a negative payoff $v - \hat{\sigma}$. If she had bid $\hat{\sigma} - \delta$, she would have avoided this loss. Thus, bidder i would have gained $\hat{\sigma} - v$ had she reduced her bid from $\hat{\sigma}$ to $\hat{\sigma} - \delta$ in that auction. Moreover, the state of the market will move from ω to ω' with this single change. Therefore, $r(\omega, \omega') \leq \frac{1}{\hat{\sigma} - v}$. However, $\hat{\sigma} - v$ is the maximum amount that this buyer could have gained by changing her bid from $\hat{\sigma} > v$ (see Facts 2.1). Thus $r(\omega, \omega') = \frac{1}{\hat{\sigma} - v}$ (using Lemma 2.2).

Consider a minimal ω -rooted tree. In it, there must exist an edge from state ω'_j to ω'_{j+1} (on the directed path from ω' to ω) such that at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from $\hat{\sigma} - \delta$ to something else. That is, $N_B(\omega'_{j+1}, v, \hat{\sigma} - \delta) < N_B(\omega'_j, v, \hat{\sigma} - \delta)$. A bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid $\hat{\sigma} - \delta$ could not have gained had she increased her bid in any auction since $\hat{\sigma} - \delta \geq v$ (see Facts 2.1). Also, $\hat{\sigma} - \delta - v$ is the maximum amount that

a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could have gained by decreasing her bid from $\hat{\sigma} - \delta$ (see Facts 2.1). Thus, the maximum gain of a bidder in $\mathcal{P} \cap B_v$ from changing her bid from $\hat{\sigma} - \delta$ is less than $\hat{\sigma} - v$. Therefore, $r(\omega'_j, \omega'_{j+1}) > \frac{1}{\hat{\sigma} - v} = r(\omega, \omega')$ (using Lemma 2.2).

Now, if we add the directed edge (ω, ω') and delete the directed edge $(\omega'_j, \omega'_{j+1})$ from the minimal ω -rooted tree, we get a ω'_j -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction.

The above analysis implies that in any stochastically stable state ω , $N_B(\omega, v, \sigma) = 0$ for all $v \in V$ and all $\sigma > v$. In particular, this implies that in any stochastically stable state ω , $N_B(\omega, 0, 0) = P_0$, that is, all buyers in the population $\mathcal{P} \cap \mathcal{B}_0$ bid exactly 0.

Case 2: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some $\sigma < v$.

Consider state ω'' such that the distribution of bids in ω'' is exactly the same as in ω except that $N_B(\omega'', v, \sigma) = N_B(\omega, v, \sigma) - 1$ and $N_B(\omega'', v, \sigma + \delta) = N_B(\omega, v, \sigma + \delta) + 1$. Any transition from ω to ω'' involves at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from σ to something else. Pick any such bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid σ in state ω . Call her bidder i . In state ω , there is a positive probability that bidder i is matched with a seller, a single bidder from \mathcal{O} who bids σ and other $M' - 2$ bidders from $\mathcal{P} \cap \mathcal{B}_0$ (here again we are using the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). All bidders from $\mathcal{P} \cap \mathcal{B}_0$ bid 0 in state ω since ω is stochastically stable. Therefore, $\sigma_{-i}^* = \sigma$. Thus, there is a positive probability that bidder i loses the auction. If she had bid $\sigma + \delta$, she would have won the auction with a payoff of $v - \sigma_{-i}^* = v - \sigma$. Thus, bidder i would have gained $v - \sigma$ had she increased her bid from σ to $\sigma + \delta$ in that auction. Moreover, the state of the market will move from ω to ω'' with this single change. Thus $r(\omega, \omega'') \leq \frac{1}{v - \sigma}$. Moreover, $v - \sigma$ is the maximum amount that this bidder could have gained by changing her bid from $\sigma < v$ (see Facts 2.1). Thus $r(\omega, \omega'') = \frac{1}{v - \sigma}$ (using Lemma 2.2).

Consider a minimal ω -rooted tree. In it, there must exist an edge from state ω''_j to ω''_{j+1} (on the directed path from ω'' to ω) such that at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from $\sigma + \delta$ to something else. That is, $N_B(\omega''_{j+1}, v, \sigma + \delta) < N_B(\omega''_j, v, \sigma + \delta)$. A bidder in $\mathcal{P} \cap \mathcal{B}_v$ could not have gained had she decreased her bid from $\sigma + \delta$ in any auction since $\sigma + \delta \leq v$ (see Facts 2.1). Also, $v - (\sigma + \delta)$ is the maximum amount that a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could have gained by increasing her bid from $\sigma + \delta$ (see Facts 2.1). Thus, the maximum gain of a bidder in $\mathcal{P} \cap B_v$ from changing her bid from

$\sigma + \delta$ is less than $v - \sigma$. Therefore, $r(\omega''_j, \omega''_{j+1}) > \frac{1}{v-\sigma} = r(\omega, \omega'')$ (using Lemma 2.2).

Now, if we add the directed edge (ω, ω'') and delete the directed edge $(\omega''_j, \omega''_{j+1})$ from the minimal ω -rooted tree, we get a ω''_j -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction.

The above analysis implies that in any stochastically stable state ω , $N_B(\omega, v, \sigma) = 0$ for all $v \in V$ and all $\sigma < v$. ■

Proof of Proposition 3.3:

Case 1: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some $\sigma > v$. Let $\hat{\sigma}$ be the highest $\sigma' > v$ such that $N_B(\omega, v, \sigma') > 0$.

Consider state ω' such that the distribution of bids in ω' is exactly the same as in ω except that $N_B(\omega', v, \hat{\sigma}) = N_B(\omega, v, \hat{\sigma}) - 1$ and $N_B(\omega', v, \hat{\sigma} - \delta) = N_B(\omega, v, \hat{\sigma} - \delta) + 1$. Any transition from ω to ω' involves at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from $\hat{\sigma}$ to something else. Pick any bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid $\hat{\sigma}$ in state ω . Call her bidder i . In state ω , there is a positive probability that this bidder i is matched with a seller and $M' - 1$ bidders from $\mathcal{P} \cap \mathcal{B}_v$ (here we are using again the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). Since $\sigma_{-i}^* \leq \hat{\sigma}$, there is a positive probability that bidder i wins the object and gets a negative payoff $v - \hat{\sigma}$. If she had bid $\hat{\sigma} - \delta$, she would have gained a positive amount irrespective of whether she would have won or lost the auction after such a reduction in her bid (see Facts 3.1). Moreover, the state of the market will move from ω to ω' with this single change. Thus, $r(\omega, \omega') < \frac{1}{\gamma}$.

Consider a minimal ω -rooted tree. Note that $\sum_{\sigma' > \hat{\sigma} - \delta} N_B(\omega', v, \sigma') < \sum_{\sigma' > \hat{\sigma} - \delta} N_B(\omega, v, \sigma')$. Therefore, there must exist states ω'_j and ω'_{j+1} on the directed path from ω' to ω in the ω -tree, such that $\sum_{\sigma' > \hat{\sigma} - \delta} N_B(\omega'_{j+1}, v, \sigma') > \sum_{\sigma' > \hat{\sigma} - \delta} N_B(\omega'_j, v, \sigma')$. That is, at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ increases her bid from some $\tilde{\sigma} \leq \hat{\sigma} - \delta$ to some $\tilde{\sigma} > \hat{\sigma} - \delta$ in the transition from ω'_j to ω'_{j+1} . However, this bidder could not have gained by such an increase in her bid since $\tilde{\sigma} > \hat{\sigma} - \delta \geq v$ (see Facts 3.1). Hence, $r(\omega'_j, \omega'_{j+1}) \geq \frac{1}{\gamma} > r(\omega, \omega')$ (using Lemma 3.2).

Now, if we add the directed edge (ω, ω') and delete the directed edge $(\omega'_j, \omega'_{j+1})$ from the minimal ω -rooted tree, we get a ω'_j -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction.

The above analysis implies that in any stochastically stable state ω ,

$N_B(\omega, v, \sigma) = 0$ for all $v \in V$ and all $\sigma > v$. In particular, this implies that in any stochastically stable state ω , $N_B(\omega, 0, 0) = P_0$, that is, all bidders in the population $\mathcal{P} \cap \mathcal{B}_0$ bid exactly 0.

Case 2: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some σ such that $\max\{v - \eta\delta, \frac{v+\delta}{2}\} < \sigma \leq v$. This inequality will only hold for $v > \delta$, which in turn implies that $\sigma > \delta$.

Consider state ω'' such that the distribution of bids in ω'' is exactly the same as in ω except that $N_B(\omega'', v, \sigma) = N_B(\omega, v, \sigma) - 1$ and $N_B(\omega'', v, \max\{\sigma - \eta\delta, \delta\}) = N_B(\omega, v, \max\{\sigma - \eta\delta, \delta\}) + 1$. Any transition from ω to ω'' involves at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from σ to something else. Pick any such bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid σ in state ω . Call her bidder i . In state ω , there is a positive probability that bidder i is matched with a seller and $M' - 1$ bidders from $\mathcal{P} \cap \mathcal{B}_0$ (here again we use the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). All bidders in $\mathcal{P} \cap \mathcal{B}_0$ use the bid 0 in state ω since ω is stochastically stable. Therefore, $\sigma_{-i}^* = 0$. Thus, bidder i wins the auction and gets $v - \sigma$. If she had bid $\max\{\sigma - \eta\delta, \delta\}$, she would have still won the auction with a payoff of $v - \max\{\sigma - \eta\delta, \delta\}$. Thus, bidder i would have gained $\min\{\eta\delta, \sigma - \delta\}$ had she reduced her bid from σ to $\max\{\sigma - \eta\delta, \delta\}$ in that auction. Moreover, the state of the market will move from ω to ω'' with this single change. Thus $r(\omega, \omega'') \leq \frac{1}{\min\{\eta\delta, \sigma - \delta\}}$.

Consider now an ω -tree of minimal resistance. Since $\sum_{\sigma' < \sigma} N_B(\omega'', v, \sigma') > \sum_{\sigma' < \sigma} N_B(\omega, v, \sigma')$, there must exist states ω''_j and ω''_{j+1} on the directed path from ω'' to ω on the ω -rooted tree, such that $\sum_{\sigma' < \sigma} N_B(\omega''_{j+1}, v, \sigma') < \sum_{\sigma' < \sigma} N_B(\omega''_j, v, \sigma')$. That is, at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from $\check{\sigma} < \sigma$ to some $\tilde{\sigma} \geq \sigma$ in the transition from ω''_j to ω''_{j+1} . Facts 3.1 tell us that this bidder would have gained a positive amount had she increased her bid from $\check{\sigma}$ to $\tilde{\sigma}$ in some auction only if $\tilde{\sigma} < v$. Thus, if $\tilde{\sigma} \geq v$, then $r(\omega''_j, \omega''_{j+1}) \geq \frac{1}{\gamma}$ (using Lemma 3.2). On the other hand, if $\tilde{\sigma} < v$, then the maximum amount she could have gained by increasing her bid from $\check{\sigma}$ to $\tilde{\sigma}$ equals $v - \tilde{\sigma}$ (see Facts 3.1). Thus, in this case, $r(\omega''_j, \omega''_{j+1}) \geq \frac{1}{v - \tilde{\sigma}} \geq \frac{1}{v - \sigma}$ (using Lemma 3.2). Also, $v - \sigma < \min\{\eta\delta, \sigma - \delta\}$ since $\sigma > \max\{v - \eta\delta, \frac{v+\delta}{2}\}$. This implies that $r(\omega''_j, \omega''_{j+1}) > r(\omega, \omega'')$. Therefore, by adding the directed edge (ω, ω'') and deleting the directed edge $(\omega''_j, \omega''_{j+1})$ from the minimal ω -rooted tree, we will get a ω''_j -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction.

Case 3: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some σ such that $\eta\delta \leq \sigma < v - \eta\delta - \delta$.

Consider state $\hat{\omega}$ such that the distribution of bids in $\hat{\omega}$ is exactly the same as in ω except that $N_B(\hat{\omega}, v, \sigma) = N_B(\omega, v, \sigma) - 1$ and $N_B(\hat{\omega}, v, \sigma + \delta) = N_B(\omega, v, \sigma + \delta) + 1$. Any transition from ω to $\hat{\omega}$ involves at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from σ to something else. Pick any such bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid σ in state ω . Call her bidder i . In state ω , there is a positive probability that bidder i is matched with a seller, a single bidder from \mathcal{O} who bids σ and $M' - 2$ bidders from $\mathcal{P} \cap \mathcal{B}_0$ (here again we use the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). All bidders in $\mathcal{P} \cap \mathcal{B}_0$ use the bid 0 in state ω since ω is stochastically stable. Therefore, $\sigma_{-i}^* = \sigma$. Thus, bidder i loses the auction with positive probability and gets 0. If she had bid $\sigma + \delta$, she would have won the auction with a payoff of $v - (\sigma + \delta)$. Thus, bidder i would have gained $v - (\sigma + \delta)$ had she increased her bid from σ to $\sigma + \delta$ in that auction. Moreover, the state of the market will move from ω to $\hat{\omega}$ with this single change. Thus, $r(\omega, \hat{\omega}) \leq \frac{1}{v - (\sigma + \delta)}$.

Consider now an ω -tree of minimal resistance. Note that $N_B(\hat{\omega}, v, \sigma + \delta) > N_B(\omega, v, \sigma + \delta)$. Therefore, there must exist states $\hat{\omega}_j$ and $\hat{\omega}_{j+1}$ on the directed path from $\hat{\omega}$ to ω on the ω -rooted tree such that $N_B(\hat{\omega}_{j+1}, v, \sigma + \delta) < N_B(\hat{\omega}_j, v, \sigma + \delta)$. That is, at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from $\sigma + \delta$ in the transition from $\hat{\omega}_j$ to $\hat{\omega}_{j+1}$. The maximum a bidder in $\mathcal{P} \cap \mathcal{B}_v$ would have gained had she increased her bid from $\sigma + \delta < v - \eta\delta \leq v - \delta$ in any auction equals $v - (\sigma + 2\delta) < v - (\sigma + \delta)$ (see Facts 3.1). On the other hand, the maximum a bidder in $\mathcal{P} \cap \mathcal{B}_v$ would have gained had she reduced her bid from $\sigma + \delta < v$ in any auction equals $\min\{\eta\delta, \sigma\} = \eta\delta < v - (\sigma + \delta)$ (see Facts 3.1). In either case, the gain is less than $v - (\sigma + \delta)$. Therefore, Lemma 3.2 tells us that $r(\hat{\omega}_j, \hat{\omega}_{j+1}) > \frac{1}{v - (\sigma + \delta)} \geq r(\omega, \hat{\omega})$. So, by adding the directed edge $(\omega, \hat{\omega})$ and deleting the directed edge $(\hat{\omega}_j, \hat{\omega}_{j+1})$ from the minimal-resistance ω -rooted tree, we will get a $\hat{\omega}_j$ -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction.

Case 4: Suppose ω is stochastically stable but $N_B(\omega, v, \sigma) > 0$ for some $v \in V$ and some σ such that $\sigma < \min\{\eta\delta, \frac{v - \delta}{2}\}$. Note that there will exist a $\sigma \geq 0$ satisfying this inequality only if $v \geq 2\delta$.

Consider state $\hat{\omega}'$ such that the distribution of bids in $\hat{\omega}'$ is exactly the same as in ω except that $N_B(\hat{\omega}', v, \sigma) = N_B(\omega, v, \sigma) - 1$ and $N_B(\hat{\omega}', v, \sigma + \delta) = N_B(\omega, v, \sigma + \delta) + 1$. Any transition from ω to $\hat{\omega}'$ involves at least one bidder

in $\mathcal{P} \cap \mathcal{B}_v$ changing her bid from σ to some other bid. Pick any such bidder in $\mathcal{P} \cap \mathcal{B}_v$ who is using the bid σ in state ω . Call her bidder i . In state ω , there is a positive probability that bidder i is matched with a seller, a single bidder from \mathcal{O} who bids σ and $M' - 2$ bidders from $\mathcal{P} \cap \mathcal{B}_0$ (here again we use the assumption that $M' \leq \min\{P_0, \dots, P_1\}$). All bidders from $\mathcal{P} \cap \mathcal{B}_0$ use the bid 0 in state ω since ω is stochastically stable. Therefore, $\sigma_{-i}^* = \sigma$. Thus, bidder i loses the auction with positive probability and gets 0. If she had bid $\sigma + \delta$, she would have won the auction with a payoff of $v - (\sigma + \delta)$. Thus, bidder i would have gained $v - (\sigma + \delta)$ had she increased her bid from σ to $\sigma + \delta$ in that auction. Moreover, the state of the market will move from ω to $\hat{\omega}'$ with this single change. Thus $r(\omega, \hat{\omega}) \leq \frac{1}{v - (\sigma + \delta)}$.

Consider now an ω -tree of minimal resistance. Note that $N_B(\hat{\omega}', v, \sigma + \delta) > N_B(\omega, v, \sigma + \delta)$. Therefore, there must exist states $\hat{\omega}'_j$ and $\hat{\omega}'_{j+1}$ on the directed path from $\hat{\omega}'$ to ω on the ω -rooted tree such that $N_B(\hat{\omega}'_{j+1}, v, \sigma + \delta) < N_B(\hat{\omega}'_j, v, \sigma + \delta)$. That is, at least one bidder in $\mathcal{P} \cap \mathcal{B}_v$ changes her bid from $\sigma + \delta$ in the transition from $\hat{\omega}'_j$ to $\hat{\omega}'_{j+1}$. If $\sigma + \delta \geq v - \delta$, a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could not have gained a positive amount by increasing her bid from $\sigma + \delta$ (see Facts 3.1). Whereas, if $\sigma + \delta < v - \delta$, the maximum a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could have gained had she increased her bid from $\sigma + \delta$ in any auction equals $v - (\sigma + 2\delta) < v - (\sigma + \delta)$ (see Facts 3.1). Finally, the maximum a bidder in $\mathcal{P} \cap \mathcal{B}_v$ could have gained had she decreased her bid from $\sigma + \delta < \frac{v + \delta}{2} < v$ (since $v \geq 2\delta$) in any auction equals $\min\{\eta\delta, \sigma\} = \sigma$ (see Facts 3.1). But $\sigma < \min\{\eta\delta, \frac{v - \delta}{2}\}$ implies that $\sigma < v - (\sigma + \delta)$. In any case, the gain is less than $v - (\sigma + \delta)$. So Lemma 3.2 tells us that $r(\hat{\omega}'_j, \hat{\omega}'_{j+1}) > \frac{1}{v - (\sigma + \delta)} \geq r(\omega, \hat{\omega}')$. Therefore, by adding the directed edge $(\omega, \hat{\omega}')$ and deleting the directed edge $(\hat{\omega}'_j, \hat{\omega}'_{j+1})$ from the minimal-resistance ω -rooted tree, we will get a $\hat{\omega}'_j$ -rooted tree with a total resistance less than the stochastic potential of ω , a contradiction. ■

Proof of Proposition 4.4: To avoid repetitions, we shall write down the proof only for the population of sellers. Thus, for the rest of the proof, let v stand for v_s and σ for σ_s .

Case 1: Suppose ω is stochastically stable but $N_S(\omega, v, \sigma) > 0$ for some $v \in V$ and some $\sigma < \min\{v + (1 - k)\eta\delta, \frac{1}{2-k}v + \frac{1-k}{2-k}\}$. Note that $\min\{v + (1 - k)\eta\delta, \frac{1}{2-k}v + \frac{1-k}{2-k}\} \leq 1$ for all $v \in V$, $\eta \geq 1$ and $k \in [0, 1]$.

Sub-case 1: $k = 1$. In this case, $\min\{v + (1 - k)\eta\delta, \frac{1}{2-k}v + \frac{1-k}{2-k}\} = v$. Consider state ω' such that the distribution of bids is exactly the same as in ω except that $N_S(\omega', v, \sigma) = N_S(\omega, v, \sigma) - 1$ and $N_S(\omega', v, \sigma + \delta) = N_S(\omega, v, \sigma + \delta) + 1$. Any transition from ω to ω' must involve at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changing her bid from σ to some other bid. Pick any such seller in $\mathcal{P} \cap \mathcal{S}_v$ who is using the bid σ in state ω . In state ω , there is a positive probability that this seller meets with a buyer in \mathcal{O} who bids σ . The seller trades in this interaction and gets a negative payoff of $\sigma - v$. If she had bid $\sigma + \delta$, she would have avoided this loss. Thus, she would have gained $v - \sigma$ had she increased her bid from σ to $\sigma + \delta$ in that interaction. Moreover, the state of the market will move from ω to ω' with this single change. Therefore, $r(\omega, \omega') \leq \frac{1}{v-\sigma}$. Moreover, $v - \sigma$ is the maximum amount that this seller could have gained by increasing her bid from σ since $k = 1$ (see Facts 4.1). Thus, $r(\omega, \omega') = \frac{1}{v-\sigma}$ (using Lemma 4.2).

Consider the directed path from ω' to ω on an ω -rooted tree of minimal resistance. Note that $N_S(\omega', v, \sigma + \delta) > N_S(\omega, v, \sigma + \delta)$. Thus, there must exist states ω'_j and ω'_{j+1} on the path from ω' to ω on the ω -rooted tree such that $N_S(\omega'_{j+1}, v, \sigma + \delta) < N_S(\omega'_j, v, \sigma + \delta)$. That is, at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changes her bid from $\sigma + \delta$ in the transition from ω'_j to ω'_{j+1} . A seller in $\mathcal{P} \cap \mathcal{S}_v$ would not have gained by reducing her bid from $\sigma + \delta \leq v$ (see Facts 4.1). Moreover, a seller in $\mathcal{P} \cap \mathcal{S}_v$ could not have gained more than $v - (\sigma + \delta)$ by increasing her bid from $\sigma + \delta$ since $k = 1$ (see Facts 4.1). Thus, $r(\omega'_j, \omega'_{j+1}) \geq \frac{1}{v-\sigma-\delta}$ (using Lemma 4.2). Therefore, $r(\omega'_j, \omega'_{j+1}) > r(\omega, \omega')$. Hence, by adding the directed edge (ω, ω') and deleting the directed edge $(\omega'_j, \omega'_{j+1})$ from the minimal ω -rooted tree, we will get a ω'_j -rooted tree with lower total resistance. This is a contradiction.

Sub-case 2: $k < 1$. Let ω'' be the state in which all the players bid the same as in ω except $N_S(\omega'', v, \sigma) = N_S(\omega, v, \sigma) - 1$ and $N_S(\omega'', v, \min\{\sigma + \eta\delta, 1\}) = N_S(\omega, v, \min\{\sigma + \eta\delta, 1\}) + 1$. Any transition from ω to ω'' must involve at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changing her bid from σ to something else. Pick any such seller in $\mathcal{P} \cap \mathcal{S}_v$ who is using the bid σ in state ω . In state ω , there is a positive probability that this seller is matched with a buyer in \mathcal{O} who bids 1. The seller could have gained in this interaction had she increased her bid to $\min\{\sigma + \eta\delta, 1\}$. Her gain in payoff would have been $(1 - k) \min\{\eta\delta, 1 - \sigma\}$. Moreover, the state of the market will move from ω to ω'' with this single change. Thus, $r(\omega, \omega'') \leq \frac{1}{(1-k) \min\{\eta\delta, 1-\sigma\}}$.

Consider the directed path from ω'' to ω on a minimal-resistance ω -rooted tree. Note that $\sum_{\sigma' > \sigma} N_S(\omega'', v, \sigma') > \sum_{\sigma' > \sigma} N_S(\omega, v, \sigma')$. Therefore, there must exist states ω''_j and ω''_{j+1} on the path from ω'' to ω on the ω -rooted tree such that $\sum_{\sigma' > \sigma} N_S(\omega''_{j+1}, v, \sigma') < \sum_{\sigma' > \sigma} N_S(\omega''_j, v, \sigma')$. Thus, at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ reduces her bid from some $\check{\sigma} > \sigma$ to some $\tilde{\sigma} \leq \sigma$ in the transition from ω''_j to ω''_{j+1} . A seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained a positive amount had she reduced her bid from $\check{\sigma}$ to $\tilde{\sigma}$ only if $\check{\sigma} > v + \delta$ and $k(\check{\sigma} - \delta) + (1 - k)\tilde{\sigma} - v > 0$ (see Facts 4.1).

Thus, if either $\check{\sigma} \leq v + \delta$ or $k(\check{\sigma} - \delta) + (1 - k)\tilde{\sigma} - v \leq 0$, then $r(\omega''_j, \omega''_{j+1}) \geq \frac{1}{\gamma}$ (using Lemma 4.3). Then delete the directed edge $(\omega''_j, \omega''_{j+1})$ and add the directed edge (ω, ω'') to the ω -rooted tree. We get a ω''_j -rooted tree with a lower total resistance, a contradiction.

If on the other hand $\check{\sigma} > v + \delta$ and $k(\check{\sigma} - \delta) + (1 - k)\tilde{\sigma} - v > 0$, the maximum possible gain that the seller could have made by reducing her bid from $\check{\sigma}$ to $\tilde{\sigma}$ is $k(\check{\sigma} - \delta) + (1 - k)\tilde{\sigma} - v$ (see Facts 4.1). Thus $r(\omega''_j, \omega''_{j+1}) \geq \frac{1}{k(\check{\sigma} - \delta) + (1 - k)\tilde{\sigma} - v}$ (using Lemma 4.3).

Recall that we know that $\check{\sigma} > v + \delta$ and that $\check{\sigma} \geq \sigma + \delta$. Therefore, one can have two cases:

- (i) Suppose $\check{\sigma} = \sigma + \delta$. Then, $r(\omega''_j, \omega''_{j+1}) \geq \frac{1}{k\sigma + (1 - k)\tilde{\sigma} - v} \geq \frac{1}{\sigma - v}$. By assumption, $\sigma < \min\{v + (1 - k)\eta\delta, \frac{1}{2 - k}v + \frac{1 - k}{2 - k}\}$, which implies that $\sigma - v < (1 - k) \min\{\eta\delta, 1 - \sigma\}$. Thus, $r(\omega''_j, \omega''_{j+1}) > r(\omega, \omega'')$. Then, by deleting the directed edge $(\omega''_j, \omega''_{j+1})$ and adding the directed edge (ω, ω'') to the ω -rooted tree, we get a ω''_j -rooted tree that has a lower total resistance than the minimal ω -rooted tree, a contradiction.
- (ii) Suppose $\check{\sigma} > \max\{\sigma + \delta, v + \delta\}$. Since $\tilde{\sigma} \leq \sigma$, some seller in $\mathcal{P} \cap \mathcal{S}_v$ reduces her bid by more than δ in the transition from ω''_j to ω''_{j+1} . Then consider the state ω''_l such that the distribution of bids in ω''_l is the same as in ω''_j except that $N_S(\omega''_l, v, \max\{\sigma + \delta, v + \delta\}) = N_S(\omega''_j, v, \max\{\sigma + \delta, v + \delta\}) + 1$ and $N_S(\omega''_l, v, \check{\sigma}) = N_S(\omega''_j, v, \check{\sigma}) - 1$. Any transition from ω''_j to ω''_l must involve at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changing her bid from $\check{\sigma}$ to some other bid. Pick any seller in $\mathcal{P} \cap \mathcal{S}_v$ who is using the bid $\check{\sigma}$ in state ω''_j . In state ω''_l , there is a positive probability that this seller meets with a buyer in \mathcal{O} who bids $\check{\sigma} - \delta$. The seller does not trade in this interaction but would have traded had she reduced her bid to $\check{\sigma} - \delta$ and thus gained $\check{\sigma} - \delta - v > 0$. Moreover, the state of the market will move

from ω_j'' to ω_l'' with this single change. Therefore, $r(\omega_j'', \omega_l'') \leq \frac{1}{\tilde{\sigma} - \delta - v}$. Since $\tilde{\sigma} \leq \sigma < \sigma + \delta < \tilde{\sigma}$, we conclude that $r(\omega_j'', \omega_{j+1}'') > r(\omega_j'', \omega_l'')$.

Now, if by deleting the directed edge $(\omega_j'', \omega_{j+1}'')$ and adding the directed edge (ω_j'', ω_l'') to the minimal ω -rooted tree, we get another ω -rooted tree, then we have a contradiction. So suppose by such a deletion and addition we do not get a ω -rooted tree. This is possible only if ω_j'' is in the path from ω_l'' to ω in the minimal ω -rooted tree.

So, consider the path from ω_l'' to ω_j'' on the minimal ω -rooted tree. Note that $N_S(\omega_l'', v, \max\{\sigma + \delta, v + \delta\}) > N_S(\omega_j'', v, \max\{\sigma + \delta, v + \delta\})$. Therefore, there must exist states ω_z'' and ω_{z+1}'' on the path from ω_l'' to ω_j'' , such that $N_S(\omega_{z+1}'', v, \max\{\sigma + \delta, v + \delta\}) < N_S(\omega_z'', v, \max\{\sigma + \delta, v + \delta\})$. That is, at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changes her bid from $\max\{\sigma + \delta, v + \delta\}$ to some other bid in the transition from ω_z'' to ω_{z+1}'' .

- If this seller increases her bid, then the maximum possible gain from any such increase in the bid equals $(1-k) \min\{\eta\delta, 1 - \max\{\sigma + \delta, v + \delta\}\}$ since $\max\{\sigma + \delta, v + \delta\} > v$ (see Facts 4.1). But $(1-k) \min\{\eta\delta, 1 - \max\{\sigma + \delta, v + \delta\}\} \leq (1-k) \min\{\eta\delta, 1 - \sigma\}$.
- If $\sigma \leq v$, then this seller would not have gained a positive amount by a reduction in her bid since $\max\{\sigma + \delta, v + \delta\} = v + \delta$ (see Facts 4.1). If, on the other hand, $\sigma > v$, then the maximum amount a seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained by such a reduction in her bid equals $\sigma - v$ (see Facts 4.1). By assumption, $\sigma < \min\{v + (1-k)\eta\delta, \frac{1}{2-k}v + \frac{1-k}{2-k}\}$, which implies that $\sigma - v < (1-k) \min\{\eta\delta, 1 - \sigma\}$.

The above analysis implies that $r(\omega_z'', \omega_{z+1}'') \geq \frac{1}{(1-k) \min\{\eta\delta, 1 - \sigma\}}$ (using Lemma 4.2). Now, delete the directed edges $(\omega_j'', \omega_{j+1}'')$ and $(\omega_z'', \omega_{z+1}'')$ and add the directed edges (ω, ω'') and (ω_j'', ω_l'') to get an ω_z'' -rooted tree. Since $r(\omega_j'', \omega_{j+1}'') > r(\omega_j'', \omega_l'')$ and $r(\omega_z'', \omega_{z+1}'') \geq r(\omega, \omega'')$, the ω_z'' -rooted tree has a lower total resistance than the minimal ω -rooted tree, also a contradiction.

Case 2: Suppose ω is stochastically stable but $N_S(\omega, v, \sigma) > 0$ for some $v \in V$ and some σ such that $v + (1-k)\eta\delta + \delta < \sigma \leq 1 + \delta - \eta\delta$. Let $\hat{\omega}$ be the state in which all the players bid the same as in ω except $N_S(\hat{\omega}, v, \sigma) = N_S(\omega, v, \sigma) - 1$ and $N_S(\hat{\omega}, v, \sigma - \delta) = N_S(\omega, v, \sigma - \delta) + 1$. Any transition from ω to $\hat{\omega}$ must involve at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changing her bid from σ to something

else. Pick any seller in $\mathcal{P} \cap \mathcal{S}_v$ who is using the bid σ in state ω . In state ω , there is a positive probability that this seller meets with a buyer in \mathcal{O} who bids $\sigma - \delta$. The seller does not trade in this interaction but could have traded by bidding equal to the buyer's bid of $\sigma - \delta$ instead of σ . Thus, her gain in payoff would have been $\sigma - \delta - v > 0$. Moreover, the state of the market will move from ω to $\hat{\omega}$ with this single change. Therefore, $r(\omega, \hat{\omega}) \leq \frac{1}{\sigma - \delta - v}$.

Consider the directed path from $\hat{\omega}$ to ω on an ω -rooted tree of minimal resistance. Note that $N_S(\hat{\omega}, v, \sigma - \delta) > N_S(\omega, v, \sigma - \delta)$. Therefore, there must exist states $\hat{\omega}_j$ and $\hat{\omega}_{j+1}$ on the path from $\hat{\omega}$ to ω on the ω -rooted tree such that $N_S(\hat{\omega}_{j+1}, v, \sigma - \delta) < N_S(\hat{\omega}_j, v, \sigma - \delta)$. Thus, at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changes her bid from $\sigma - \delta$ in the transition from $\hat{\omega}_j$ to $\hat{\omega}_{j+1}$. Note that $\sigma - \delta > v$. If $\sigma - \delta = v + \delta$, then a seller in $\mathcal{P} \cap \mathcal{S}_v$ would not have gained a positive amount had she reduced her bid from $\sigma - \delta$ (see Facts 4.1). If $\sigma - \delta > v + \delta$, then the maximum a seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained had she reduced her bid from $\sigma - \delta$ equals $\sigma - 2\delta - v < \sigma - \delta - v$ (see Facts 4.1). On the other hand, the maximum a seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained by increasing her bid from $\sigma - \delta > v$ equals $(1 - k) \min\{\eta\delta, 1 - \sigma + \delta\} = (1 - k)\eta\delta < \sigma - \delta - v$ (see Facts 4.1). Thus, in any case, the gain is less than $\sigma - \delta - v$. Therefore, Lemma 4.2 tells us that $r(\hat{\omega}_j, \hat{\omega}_{j+1}) > \frac{1}{\sigma - \delta - v} \geq r(\omega, \hat{\omega})$. So, by deleting the directed edge $(\hat{\omega}_j, \hat{\omega}_{j+1})$ and adding the directed edge $(\omega, \hat{\omega})$ to the minimal ω -rooted tree, we will get a $\hat{\omega}_j$ -rooted tree with a lower total resistance, which is a contradiction.

Case 3: Suppose ω is stochastically stable but $N_S(\omega, v, \sigma) > 0$ for some $v \in V$ and some σ such that $\sigma > \max\{1 + \delta - \eta\delta, \frac{1}{2-k}v + \frac{1-k}{2-k} + \delta\}$. Let $\hat{\omega}'$ be the state in which all the players bid the same as in ω except $N_S(\hat{\omega}', v, \sigma) = N_S(\omega, v, \sigma) - 1$ and $N_S(\hat{\omega}', v, \sigma - \delta) = N_S(\omega, v, \sigma - \delta) + 1$. Any transition from ω to $\hat{\omega}'$ must involve at least one seller in $\mathcal{P} \cap \mathcal{S}_v$ changing her bid from σ . Pick a seller in $\mathcal{P} \cap \mathcal{S}_v$ who is using the bid σ in state ω . In state ω , there is a positive probability that this seller meets with a buyer in \mathcal{O} who bids $\sigma - \delta$. The seller does not trade in this interaction but would have traded had she bid equal to the buyer's bid of $\sigma - \delta$ instead of σ . Her gain in payoff would have been $\sigma - \delta - v > 0$. Moreover, the state of the market will move from ω to $\hat{\omega}'$ with this single change. Therefore, $r(\omega, \hat{\omega}') \leq \frac{1}{\sigma - \delta - v}$.

Consider now the directed path from $\hat{\omega}'$ to ω on a minimal-resistance ω -rooted tree. Note that $N_S(\hat{\omega}', v, \sigma - \delta) > N_S(\omega, v, \sigma - \delta)$. Therefore, there must exist states $\hat{\omega}'_j$ and $\hat{\omega}'_{j+1}$ on the path from $\hat{\omega}'$ to ω on the ω -rooted tree such that $N_S(\hat{\omega}'_{j+1}, v, \sigma - \delta) < N_S(\hat{\omega}'_j, v, \sigma - \delta)$. Thus, at least one seller

in $\mathcal{P} \cap \mathcal{S}_v$ changes her bid from $\sigma - \delta$ to something else in the transition from $\hat{\omega}'_j$ to $\hat{\omega}'_{j+1}$. Note that $\sigma - \delta > v$. If $\sigma - \delta = v + \delta$, then a seller in $\mathcal{P} \cap \mathcal{S}_v$ would not have gained a positive amount had she reduced her bid from $\sigma - \delta$ (see Facts 4.1). If $\sigma - \delta > v + \delta$, then the maximum a seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained had she reduced her bid from $\sigma - \delta > v$ equals $\sigma - 2\delta - v < \sigma - \delta - v$ (see Facts 4.1). On the other hand, the maximum a seller in $\mathcal{P} \cap \mathcal{S}_v$ could have gained had she increased her bid from $\sigma - \delta > v$ equals $(1 - k) \min\{\eta\delta, 1 - \sigma + \delta\} = (1 - k)(1 - \sigma + \delta) < \sigma - \delta - v$ (see Facts 4.1). Thus in any case, the gain is less than $\sigma - \delta - v$. Therefore, Lemma 4.2 tells us that $r(\hat{\omega}'_j, \hat{\omega}'_{j+1}) > \frac{1}{\sigma - \delta - v} \geq r(\omega, \hat{\omega})$. So by deleting the directed edge $(\hat{\omega}'_j, \hat{\omega}'_{j+1})$ and adding the directed edge $(\omega, \hat{\omega}')$ to the minimal ω -rooted tree, we will get a $\hat{\omega}'_j$ -rooted tree with lower total resistance, which is our final sought contradiction. ■

References

- Abreu, D., Gul, F., 2000. Bargaining and Reputation. *Econometrica* 68, 85-117.
- Agastya, M., 2004. Stochastic Stability in a Double Auction. *Games and Economic Behavior* 48, 203-222.
- Ben-Ner, A., Putterman, L., 2001. Trusting and Trustworthiness. *Boston Law Review* 81, 523-551.
- Chatterjee, K. and Samuelson, W., 1983. Bargaining Under Incomplete Information. *Operations Research* 31, 835-851.
- Ellingsen, T., 1997. The Evolution of Bargaining Behavior. *Quarterly Journal of Economics* 112, 581-602.
- Foster, D. P., Vohra, R. V., 1998. Asymptotic Calibration. *Biometrika* 85, 379-390.
- Foster, D. P., Young, H. P., 1990. Stochastic Evolutionary Games Dynamics. *Theoretical Population Biology* 38, 219-232.
- Hart, S., 2005. Adaptive Heuristics. *Econometrica* 73, 1401-1430.

- Hart, S., Mas-Colell, A., 2000. A Simple Adaptive Procedure Leading to Correlated Equilibrium. *Econometrica* 68, 1127-1150.
- Hyafil, N., Boutilier, C., 2004. Regret Minimizing Equilibria and Mechanisms for Games with Strict Type Uncertainty. In: Chickering, M., Halpern, J. (Eds.). *Uncertainty in Artificial Intelligence: Proceedings of the 20th Conference*. Arlington: AUAI Press, 268-277.
- Jensen, M., Sloth, B., Wittmann-Jacobsen, H. J., 2005. The Evolution of Conventions Under Incomplete Information. *Economic Theory* 25, 171-185.
- Kandori, M., Mailath, G., Rob, R., 1993. Learning, Mutation, and Long-Run Equilibria in Games. *Econometrica* 61, 29-56.
- Young, H. P., 1993. The Evolution of Conventions. *Econometrica* 61, 57-84.
- Young, H. P., 1998. *Individual Strategy and Social Structure*. Princeton: Princeton University Press.