

Existence of Equilibrium in Large Double Auctions*

Drew Fudenberg, Markus M. Mobius and Adam Szeidl
Harvard University †

April 16, 2003

Abstract

We show the existence of a pure strategy, symmetric, increasing equilibrium in double auction markets with correlated private valuations and many participants. The equilibrium we find is arbitrarily close to fully revealing as the market size grows.

JEL Classification: *C62, C72, D44, D82*

Keywords: *Double auctions, Equilibrium, Existence*

1 Introduction

This paper establishes the existence of pure-strategy equilibria for large double auctions markets with correlated private values. In these equilibria, bids are very close to valuations, and so can be interpreted as approximately truthful reports of the agents' information. Our main contribution is the existence result; Rustichini, Satterthwaite, and Williams (1994) showed that a symmetric equilibrium, if exists, must be close to truth telling in the independent private values case.

Jackson and Swinkels (2001) prove the existence of a (non-trivial) mixed-strategy equilibrium in a variety of auctions by taking the limits of equilibria in auctions with a discretized space of bids. Reny and Perry (2003) show the existence of an equilibrium with non-decreasing bid functions in all sufficiently large discretized auctions; this equilibrium approximates the rational expectations

*Preliminary. We thank Eric Maskin, Andy Postlewaite and Phil Reny for helpful comments and suggestions.

†Department of Economics, Littauer Center, Cambridge, MA 02138. E-mail: dfudenberg@harvard.edu, mobius@fas.harvard.edu, szeidl@fas.harvard.edu.

equilibrium of the continuum limit. Inspired by their conclusion, we obtain monotonicity without the use of grids by perturbing the payoff functions slightly and considering only large auctions. We can then show that best responses in large auctions are close to truth telling and hence monotonic. This monotonicity lets us appeal to a fixed-point theorem on a space of increasing functions. Athey (2001)'s existence theorem for one-sided auction also makes use of monotonicity to appeal to a fixed point theorem; she obtains monotonicity from a single-crossing assumption.

Of course there are trade-offs between our methodology and the alternatives in the literature, and issues that are taken for granted in the finite case require more attention. Most important among these is the existence and uniqueness of best responses, and the continuity of the best response map.¹ On the other hand, the approach based on grids may only lead to mixed equilibria (as in the case of Jackson and Swinkels (2001)) and also needs to take extra care of ties.² Indeed, one main contribution of the paper by Jackson and Swinkels is their clever approach to ties. In our case, tie breaking is resolved with a small perturbation.

Conceptually, there are two main properties that we require in order for our fixed point argument to work, namely that best responses need to exist and be increasing. The first issue is easier to deal with. The only reason that best responses may fail to exist is that ties can have positive probability. To take care of ties, we perturb the model so that the distribution of the realized price conditional on any bid profile is a smooth random variable. This makes the payoffs a continuous functions of the bids, so that best responses exist. As remarked above, we ensure that best responses will be increasing by restricting attention to large auctions, so that equilibrium bids will be close to truthful reporting. This allows us to apply a fixed point theorem to the compact space of increasing functions.

There are also a number of technical difficulties. Prominent among these is that even if best responses stay close to truth-telling, they need not be quite as close as the opponents' strategies, so that the best response map need not send a small neighborhood of the diagonal into itself. To deal with this issue, we introduce a truncated best response map that forces the best response bid function of any player to remain sufficiently close to the diagonal. A related issue is that a best response does not need to be close to truth-telling in case trade for the relevant bidder is a zero probability event. To take care of this, we introduce a second small perturbation, namely that with small probability ε the price is drawn from an exogenous distribution. This guarantees that otherwise indifferent traders bid

¹We need uniqueness because Kakutani's theorem has no parallel for infinite dimensional spaces.

²This may not be exclusively for technical reasons. The nature of tie-breaking can matter for equilibrium existence in certain auctions, as pointed out by Maskin and Riley (2000). Moreover, the McAdams (2003) example of an equilibrium with non-monotonic bids in a one-sided auction also depends on the form of the tie-breaking rule.

their true valuation.³

Given these modifications, we are able to construct a fixed point of the “truncated (ε, φ) -perturbed best response map.” Once we have the fixed point, we take the limit as φ , which measures the extent of the first perturbation, goes to zero. Then we make use of the argument in Rustichini, Satterthwaite, and Williams (1994) to show that the limiting profile is in fact so close to the diagonal that we were not truncating at all. Finally, we send the magnitude of the second perturbation, ε to zero, and this gives a symmetric, pure strategy, increasing equilibrium of the κ -double auction.

The κ -double auction has been introduced to the literature by Chatterjee and Samuelson (1983), and by Wilson (1985) for the multilateral case. The existence of non-trivial equilibria for multiple buyers and sellers has been an open question since then. Recently, there have been two papers that deal with this issue. Reny and Perry (2003) prove existence for large enough double auctions with affiliated values, where bidders are restricted to name prices in a fine, finite grid. Jackson and Swinkels (2001) prove the existence of a mixed equilibrium under very weak assumptions.

The pure, symmetric, increasing equilibria of the auction we consider here were studied by Rustichini, Satterthwaite, and Williams (1994) in the independent private values case. They do not prove existence, but they do show that any symmetric equilibrium has to be arbitrarily close to truth-telling if the number of participants is large enough. We extend their argument to the correlated private-values setting in this paper. Finally, Williams (1991) settles the existence question in a special case, the buyer’s bid double auction. In that market sellers always bid their underlying valuations; this makes the existence argument easier.

2 The Model

Consider the κ -double auction with m buyers and n sellers. The auction mechanism is defined as follows. Each seller has a single unit of the indivisible good, and each buyer wishes to purchase one unit. Sellers and buyers have correlated valuations that are private information. The joint distribution of uncertainty, as well as the structure of the game is common knowledge. Given her realized value, each actor submits a bid to the market. These bids are then ordered from highest to lowest. The market price is determined to be a weighted average of the n -th and $n + 1$ -th bids, with weights κ and $1 - \kappa$ respectively. Buyers whose bids are above and sellers whose bids are below, this market price, buy or sell respectively one unit at the prevailing price. In the case of a tie (i.e., if a bid is equal to the price) some feasible tie-breaking mechanism is applied. The exact nature of

³Jackson and Swinkels (2001) use a similar perturbation to avoid the no trade equilibrium.

tie breaking will not be relevant. Importantly, the mechanism so described is equivalent to constructing the piecewise linear demand and supply curves from buyers' and sellers' bids and then finding one of the many possible market clearing prices. For a more detailed discussion of the κ -double auction please see Rustichini, Satterthwaite, and Williams (1994).

We make a number of assumptions about the joint distribution of valuations. We focus on correlated private values. Let s be the common noise term to all valuations. The distribution of s is $G(s)$ which is assumed to be concentrated on the unit interval (this is without loss of generality) and absolutely continuous. Conditional on s , the valuations of all buyers and sellers are independent. Let the conditional probability distribution of the valuation of a buyer i be $F_B(v_i|s)$, that of a seller j be $F_S(v_j|s)$ and let $f_B(v_i|s)$ and $f_S(v_j|s)$ be the corresponding densities. All valuations are concentrated on the unit interval, with density uniformly bounded away from zero for all s . From these distributions one can calculate the "inverse conditionals." Define $H_i(s|v_i)$ to be the conditional distribution of s given v_i for player i (who can be either a buyer or a seller) and let $h_i(s|v_i)$ be the corresponding density. We assume that $H_i(s|v_i)$ is absolutely continuous with respect to the Lebesgue measure, with full support on the set of values s assumes. Moreover, the density $h_i(s|v_i)$ is assumed to be uniformly Lipschitz in v_i and bounded away from zero. Note that this specification of correlated values includes the well studied independent private values case (see Rustichini, Satterthwaite, and Williams (1994) and Williams (1991)).

In the rest of the paper, we would like to prove that for a large enough number of participants, this auction game has a symmetric Bayesian equilibrium in pure strategies. We need to formalize the notion of large enough auctions. The idea is to increase n and m simultaneously, such that the ratio $\gamma = n/(n + m)$ remains bounded away from both zero and one. For convenience, we will use the notation $N = n + m$. Whenever we use the term "for large enough auctions", what we have in mind is increasing N while keeping γ bounded away from zero and one. Just how large the auction needs to be depends on how tight these bounds are.

We attack the problem by first finding equilibria of slightly different games. Fix $1 > \varepsilon, \varphi \geq 0$, and let us introduce the (ε, φ) -perturbed auction. In that game, payoffs are defined as follows: with probability ε the price \tilde{p} is independent of the bids, and is drawn from a uniform distribution on $[0, 1]$. In this event, the expected payoff of a buyer making bid b_i is $\int_0^{b_i} (v_i - p) dp$ and the expected payoff of a seller is $\int_{b_i}^1 (p - v_i) dp$. With remaining probability $1 - \varepsilon$ the mechanism is as follows. The preliminary price p is determined by the standard κ -double auction but the actual price \tilde{p} will be a smooth random variable which is φ -close to the preliminary

price. Formally, the perturbed price \tilde{p} is defined as

$$\tilde{p}(p, \omega_\varphi) = \begin{cases} 2p(\omega_\varphi - \frac{1}{2}) + p & \text{if } p < \varphi \\ 2\varphi(\omega_\varphi - \frac{1}{2}) + p & \text{if } \varphi \leq p < 1 - \varphi \\ 2(1 - p)(\omega_\varphi - \frac{1}{2}) + p & \text{if } 1 - \varphi \leq p < 1 \end{cases} \quad (1)$$

where ω_φ is a uniform random variable on the unit interval. The price perturbation has the following properties: (1) the random variable \tilde{p} is smooth (absolutely continuous) conditional on the preliminary price p ; (2) the function $\tilde{p}(p, \omega_\varphi)$ is continuous in p ; (3) $\tilde{p}(p, \frac{1}{2}) = p$; (4) the perturbed price $\tilde{p}(p, \omega_\varphi)$ is increasing in both p and ω_φ ; (5) the perturbed price is φ -close to the preliminary price, i.e. $|\tilde{p}(p, \omega_\varphi) - p| < \varphi$ always holds. Any other perturbation which satisfies these five properties is suitable for our purposes. Properties (1) and (2) make the expected payoffs a continuous function of the bids. Property (3) guarantees that the median perturbation is zero. Property (4) ensures that the event $\{\omega_\varphi | \tilde{p}(p, \omega_\varphi) < b\}$ is shrinking as p increases. Property (5) lets the perturbed game converge to the undisturbed game when we later first take φ to zero and then ε to zero. Note that all agents are allowed to make their desired trades at the actual price which need not result in a feasible outcome of the unperturbed game.

The ε perturbation ensures that bidding above one's private value is strictly dominated for a buyer because there is a positive probability that the actual price lies between the private value and the bid in which case the buyer suffers a loss. Likewise, sellers will never submit offers below their private value.

The way the uncertainty is structured in the perturbed auction is as follows. There is a probability space (Ω_s, μ_s) which generates the common signal s . For each agent i there is an independent uniform draw $(\Omega_{0,i}, \mu_{0,i})$ from the unit interval which generates the value $v_i = v_i(\omega_s, \omega_{0,i})$ of agent i . There is another probability space $(\Omega_\varepsilon, \mu_\varepsilon)$ that generates the ε probability event and the uniform price draw. Finally, $(\Omega_\varphi, \mu_\varphi)$ is a uniform draw from the unit interval, that generates the φ -perturbation and the perturbed price $\tilde{p} = \tilde{p}(p, \omega_\varphi)$. The complete probability space is denoted with (Ω, μ) where $\Omega = \prod_i \Omega_{0,i} \times \Omega_s \times \Omega_\varphi \times \Omega_\varepsilon$ and $\mu = \prod_i \mu_{0,i} \times \mu_s \times \mu_\varphi \times \mu_\varepsilon$. Elements of Ω are referred to with ω . We will frequently use the *conditional* measure μ^{v_i} of agent i who knows her private value v_i and which is defined as

$$\mu^{v_i}(A) = \frac{\mu(A \cap \{\omega | v_i = v_i(\omega)\})}{\mu(\{\omega | v_i = v_i(\omega)\})}. \quad (2)$$

We will denote the probability of some event A conditional on v_i with $P^{v_i}(A) = \int_A d\mu^{v_i}(\omega)$ and the conditional expectation over some random variable $X(\omega)$ with $E^{v_i}(X) = \int X(\omega) d\mu^{v_i}(\omega)$.

Formally, the strategy of agent i (buyer or seller) with private value v_i is referred to as $x_i(\cdot)$ where $b_i = x_i(v_i)$ is the bid of agent i . The vector of all agents' strategies

is denoted with $x(\cdot)$ and $x_{-i}(\cdot)$ refers to the strategies of all players except player i . Given the realization of uncertainty ω an agent i who makes bid b_i faces a price $p(b_i, x_{-i}(\cdot), \omega)$. To simplify notation we will sometimes suppress the dependence of the price on other players' bidding strategies and write $p(b_i, \omega)$ or simply $p(b_i)$.

3 Results

Our goal is to establish the following theorem:

Theorem 1 *The κ -double auction with correlated private values has a symmetric equilibrium in increasing pure strategies for all N large enough. In particular, there exists an equilibrium profile where each bid is of order $1/N$ close to being fully revealing.*

We begin by summarizing the main argument.

Fixed point. In order to use Schauder's fixed point theorem for Banach spaces, we need to use a continuous map defined from a compact, convex set of a Banach space to the same set. We choose the space to be $L^1[0, 1]$. Compactness is ensured if all relevant functions are increasing and the φ -perturbation implies that the best response map is continuous. However, the best-response map need not map a given set of increasing strategies into itself, because a player's best response can be further away from truth telling than the opponents' strategies are. In order to deal with this difficulty, we introduce the truncated best response map; Schauder's fixed point theorem gives a fixed point of that map.

Symmetry. We now take φ to zero and consider a sequence of fixed points. This sequence has a limit by compactness. The limit turns out to be strictly increasing everywhere for two reasons. In the regions where trade takes place with positive probability, the limit is strictly increasing because of the symmetry of the profile. In the no-trade regions the limit is strictly increasing because of the ε -perturbation. Because of this strict monotonicity no ties are possible. It follows that the best response property is preserved under taking the limit. Therefore we can find a fixed point of the truncated best response mapping of the ε -perturbed auction.

Relaxing the truncation. Given a symmetric fixed point of the truncated best response mapping, we show that in the fixed point truncation is in fact never used. The argument applied here is borrowed from Rustichini, Satterthwaite, and Williams (1994), who prove that any symmetric equilibrium stays close enough to truth telling (in the independent private values case). Using the logic of that paper, we show that our truncated fixed point never gets far enough from the diagonal for the truncation to become effective.

Taking ε to zero. Given a symmetric equilibrium for all ε , we can select a convergent subsequence. Because of symmetry, the limit will continue to be strictly increasing at all points where trade takes place with positive probability. At values where trade has zero probability, this need not be so. To check whether a limit profile is an equilibrium, we only need to verify the best response property at points where trade has positive probability; and at those points it is guaranteed by the no ties condition.

3.1 Best response of the perturbed game

Lemma 1 *In the (ε, φ) -perturbed κ -double auction, for $\varphi > 0$ a best response to any opponent profile exists.*

Proof: Consider a buyer i with valuation v_i . Her profit function as a function of her bid b_i is continuous because of the perturbation. Hence the profit function has a maximum in b_i for each v_i . A symmetric argument applies for a seller j .⁴ \diamond

This lemma has shown the importance of the φ -perturbation: it guarantees that a best response exists. We now show that a best response to “almost truthful” strategies is also “almost truthful”, where the “modulus of continuity” depends on the number of agents.

Theorem 2 *In the (ε, φ) -perturbed auction, with $\varepsilon, \varphi > 0$, for any $M > 0$ there exists $M' > 0$ such that in a large enough auction, if φ and ε are small enough and the opponents’ strategies are M' close to the diagonal, then the best response of any buyer is at least M close to the diagonal.*

Proof: See appendix A.

The only reason why a player would bid differently from her valuation is to try to influence the price. If all players’ bid functions are close to the diagonal and there are many players, then any bid is unlikely to be pivotal, so the best response must be close to the diagonal as well, though perhaps not as close as the original profile was. We also need that the best response not only stays close but is also increasing. The proof of this requires a simple lemma about the relation between the conditional expectations of a function given a value v_i and the conditional expectation given a different value v'_i .

Lemma 2 *There exists a constant K such that for any positive function $u(s)$*

$$\int_s u(s) |h(s|v'_i) - h(s|v_i)| ds \leq K \cdot \Delta v \cdot \int_s u(s) h(s|v_i) ds. \quad (3)$$

⁴In the following we will only present proofs for buyers. All seller proofs are analogous.

Proof: We have

$$|h(s|v'_i) - h(s|v_i)| \leq K \cdot \Delta v_i \cdot h(s|v_i)$$

for any s and v_i, v'_i because $h(\cdot|\cdot)$ is uniformly Lipschitz in the second argument and bounded away from zero. This completes the proof. \diamond

Theorem 3 *In the (ε, φ) -perturbed auction, with $\varepsilon, \varphi > 0$, if $v_i - x_i(v_i) < 1/4K(K+1)$, then the best response to any opponent strategy profile is weakly increasing.*

Proof: See appendix B.

To get some intuition for this theorem we consider the first-order condition of a buyer i

$$T(b_i, v_i) = (v_i - b_i) r(b_i, v_i) - \kappa R(b_i, v_i) = 0, \quad (4)$$

where $r(b_i, v_i)$ is the density of the n -th highest bid at b_i of all other agents except i , and $R(b_i, v_i)$ is the probability that i has the n -th highest bid (and is therefore pivotal). The first term captures the gain from raising a bid by a small amount (and becoming pivotal in the process) and the second term captures the cost of doing so. We next replace b_i with $x_i(v_i)$ and take the first derivative of the first order condition with respect to v_i to obtain

$$\frac{\partial T}{\partial b_i} x'_i(v_i) + \frac{\partial T}{\partial v_i} = 0. \quad (5)$$

We know that $\frac{\partial T}{\partial b_i} \leq 0$ because this is the second-order condition for a local maximum. Therefore, we only have to show that $\frac{\partial T}{\partial v_i} > 0$ to get increasingness of i 's best response $x_i(v_i)$. We can calculate

$$\frac{\partial T}{\partial v_i} = \underbrace{r(b_i, v_i)}_{\text{Term I}} + \underbrace{(v_i - b_i) \frac{\partial r}{\partial v_i}}_{\text{Term II}} - \underbrace{\kappa \frac{\partial R}{\partial v_i}}_{\text{Term III}}. \quad (6)$$

Also, note that

$$\begin{aligned} r(b_i, v_i) &= \int_0^1 r(b_i, s) h(s|v_i) ds \\ R(b_i, v_i) &= \int_0^1 R(b_i, s) h(s|v_i) ds \end{aligned} \quad (7)$$

where $r(b_i, s)$ is the conditional density for some fixed s and $R(b_i, s)$ is the conditional pivotal probability. Lemma 2 implies that the partial derivatives $\frac{\partial r}{\partial v_i}$ and $\frac{\partial R}{\partial v_i}$ are of the same order of magnitude as $r(b_i, v_i)$ and $R(b_i, v_i)$. By the first order condition (4) $R(b_i, v_i)$ is the same order of magnitude as $(v_i - b_i) r(b_i, v_i)$. Therefore, as long as $v_i - b_i$ is small, both terms II and III are small while term I is large and determines the sign of the partial derivative $\frac{\partial T}{\partial v_i}$ which is indeed positive.⁵

⁵In the special case of independent private values terms II and III are zero - therefore the best response of player i is always increasing in v_i .

Given this result and Theorem 2, we can choose C such for n large enough, if all opponents' strategies are at most C away from the diagonal, then any best response is weakly increasing. Fix such a C .

Theorem 4 *For n large enough, the best response to any profile that is at most C far from the diagonal is unique and increasing.*

Proof: We have just seen increasingness. For uniqueness, suppose that a buyer i has two best response functions, $x_i(v_i)$ and $x'_i(v_i)$. If these differ on a set of positive measure, then there is a point of continuity v_0 of both x_i and x'_i inside the unit interval where they differ (because both are increasing), say $x_i(v_0) > x'_i(v_0)$. But then there is a neighborhood of v_0 where this inequality continues to hold. Define x''_i to be equal to x'_i to the left of v_0 , and equal to x_i to the right of v_0 . Clearly x''_i is a best response, since it is a best response for almost every valuation v_i . However, x''_i is not increasing; to the left of v_0 it approaches $x_i(v_0)$, and to the right of v_0 it approaches $x'_i(v_0)$. Thus x''_i is a non-increasing best response. This is a contradiction. \diamond

3.2 Fixed point

We now introduce the truncated best response mapping of the perturbed game. For the C fixed above, let X be the set of weakly increasing functions defined on the unit interval with values in the C wide corridor around the diagonal. Formally,

$$X = \{x_i : [0, 1] \rightarrow [0, 1] \mid x_i(v_i) \leq x_i(v'_i) \text{ if } v_i \leq v'_i, \text{ and } |x_i(v_i) - v_i| \leq C \forall v_i, v'_i\}.$$

When all opponents play strategies in X the best response of any player is unique and increasing, but it need not be in X . The truncated best response map just truncates this best response by setting it equal to the bound of the corridor at points v_i where it is outside. Formally, if the best response of a buyer is $x_i(v_i)$, the truncated best response at v_i is equal to $\max(x_i(v), v_i - C, 0)$. Importantly, the truncated best response is still increasing, since it is the max (min in the case of a seller) of increasing functions.

Lemma 3 *X is a closed, convex subset of the Banach space $L^1[0, 1]$.*

Proof: The implied topology generated on X is the same as convergence in measure. Indeed, a set of uniformly bounded functions converge to a limit in L^1 if and only if they converge in measure. In addition X is also compact: By Helly's theorem, X is compact in the weak topology, so from any sequence of functions in X we can select a subsequence that is converging to some limit function in all of its points of continuity. But this implies almost everywhere convergence, and that implies convergence in L^1 by Lebesgue's dominated convergence theorem, since all function involved are in X . \diamond

In the rest of the paper we will focus on symmetric profiles, that is, all buyers respectively all sellers using the same bidding function. Therefore we only have to keep track of the pair $x = (x_B(\cdot), x_S(\cdot))$ of the buyer's and the seller's strategy. Then we have $x \in X \times X$.

It is clear that $X \times X$ is a compact, convex subset of the product Banach space $L^1[0, 1] \times L^1[0, 1]$. Now for any positive (ε, φ) we have the truncated best response map $TBR^{\varepsilon, \varphi}(\cdot) : X^2 \rightarrow X^2$. Then $TBR^{\varepsilon, \varphi}(x_B, x_S)$ is the pair of truncated best responses for any buyer respectively seller, when all opponent buyers play $x_B(\cdot)$, and all opponent sellers play $x_S(\cdot)$. We are interested in a fixed point of this map. To get one, first we need continuity.

We will now consider a sequence of perturbed games for $k = 1, 2, \dots$. Game k will have perturbations $(\varepsilon^k, \varphi^k)$. Suppose the strategies played in game k are $x^k = (x_B^k, x_S^k)$, and assume that $y^k = (y_B^k, y_S^k)$ is a best response to x^k in game k . We are interested in whether the best response property is preserved as k goes to infinity. Note that this framework incorporates continuity of the best response map if the sequence $(\varepsilon^k, \varphi^k)$ is constant.

Fix a player i with value v_i . If the player bids b_i in the k -th game the price is denoted with $p^k(b_i)$ which is a random variable. The price in the limit game is denoted by $p(b_i)$.

Lemma 4 *Suppose $x^k \rightarrow x$ in measure, and $y^k \in BR^k(x^k)$ converges in measure to y . Assume for all players i with value v_i and bid $b_i = x_i^k(v_i)$ (1) opponents play strategies $x_{-i}^k(\cdot)$; (2) the price $p^k(b_i)$ converges in probability to $p(b_i)$; (3) the distribution of $p(b_i)$ has no atom at b_i . Then $y \in BR(x)$.*

Proof: See appendix C.

The intuition for this result is the following. Suppose player i is a buyer, and denote the payoff to this buyer with value v_i and bid b_i in the k -th game by $\Pi^k(v_i, b_i)$. If $p^k(b_i)$ converges in probability, it also converges in distribution to $p(b_i)$. Hence in particular the distribution function of $p^k(b_i)$ converges to that of $p(b_i)$ at all points of continuity of the latter, in particular at b_i . In addition, the payoff function to bidding b_i equals

$$\Pi^k(v_i, b_i) = \int_{p < b_i} v_i - p \, dR^k(p(b_i))$$

where $R^k(p(b_i))$ is the price distribution in game k where our buyer bids b_i . By weak convergence, this payoff converges to the limit payoff

$$\Pi(v_i, b_i) = \int_{p < b_i} v_i - p \, dR(p(b_i)).$$

Note that the integrand is not everywhere continuous, so we cannot directly apply the weak convergence result; however, the limit distribution function is continuous at b_i , so the point of discontinuity of the integrand at b_i does not cause a problem.

The lemma will have a number of applications regarding continuity of the best response map and taking the limit as φ and ε are going to zero. In order to state these applications, we need to introduce a concept that we call “positive probability of trade.” We say that a player i in the (ε, φ) -perturbed game with opponent profile x_{-i} has positive probability of trade at value v_i and bid b_i , if in the $1 - \varepsilon$ probability event when the price is not drawn from a uniform distribution, there is positive probability that the player gets a positive payoff (gets to trade). This is equivalent to her expected payoff being strictly larger than $\varepsilon \cdot \int_0^{b_i} (v_i - p) dp$ in case of a buyer, or strictly larger than $\varepsilon \cdot \int_{b_i}^1 (p - v_i) dp$ in the case of a seller.

Corollary 1 *Suppose that either*

- (a) $\varphi^k = \varphi > 0$ and $\varepsilon^k = \varepsilon > 0$ fixed, or
- (b) $\varphi^k \rightarrow 0$, $\varepsilon^k = \varepsilon > 0$ is fixed, and x is strictly increasing, or
- (c) $\varphi^k = 0$, $\varepsilon^k \rightarrow 0$, $y^k = x^k$ and x is strictly increasing for each player i at all values v where there is positive probability of trade.

Then $y \in BR(x)$.

Proof: We will use the previous lemma. The convergence in probability of the conditional prices is obvious in all three cases, given that the bid functions converge in measure, and that the perturbations are continuous. Fixing a buyer, we only need to check whether the limiting price $p(b_i)$ has a distribution that is atomless at b_i . This is obvious in case (a) because the price distribution is completely atomless by the perturbation. In case (b), it follows from the fact that all opponent profiles are strictly increasing. Indeed, any atom that $p(b_i)$ may have at b_i has to come from some opponent bidding b_i with positive probability; this is ruled out.

In case (c), suppose b_i^k converges to b_i , and b_i is not a best response in the limit game, so b_i' does better. We can assume that no opponent bids b_i' with positive probability, otherwise we could have chosen b_i'' that is a little bit larger than b_i' , and still get a higher payoff than that earned by b_i . Therefore, by the argument of the lemma, $\Pi^k(v_i, b_i') \rightarrow \Pi(v_i, b_i')$. Now if no other player bids b with positive probability, then by the argument of the lemma we also have that $\Pi^k(v_i, b_i^k) \rightarrow \Pi(v_i, b_i)$, which yields a contradiction. Thus the only problem we may have is that some opponent bids b_i with positive probability in the limit game. If this opponent is a buyer, then by assumption (c) no trade takes place for a buyer with bid b in the limit game. But then as k goes to infinity, the payoff to bidding b_i^k must be vanishingly small. If the payoff to

bidding b'_i is positive in the limit game, then by $\Pi^k(v_i, b'_i) \rightarrow \Pi(v_i, b'_i)$ bidding b'_i for k large enough is better than bidding b_i^k . This is a contradiction.

If the opponent who bids b_i with positive probability is a seller, then that seller faces no trade in the limit game. Thus all buyers have to bid below b_i with probability one. But note that $b_i = y_B(v_i) = x_B(v_i)$. By assumption, $x_B(v_i)$ is increasing here, because buyers trade with positive probability at this stage. So $x_B(\frac{v_i+1}{2}) > x_B(v_i) = b_i$, but then our seller who bids b should get to trade with positive probability. This is a contradiction.

Note how this final argument hinges on the fact that the x^k profiles were already equilibria of game k . \diamond

Corollary 2 *In the (ε, φ) -perturbed auction, with $\varepsilon, \varphi > 0$, the truncated best-response map is continuous.*

It follows that $TBR^{\varepsilon, \varphi}(\cdot)$ is a continuous self-map of a compact, convex subset of a Banach space.

Theorem 5 *The truncated best response map $TBR^{\varepsilon, \varphi}(\cdot)$ for $\varepsilon, \varphi > 0$ small enough and all large enough auctions has a fixed point.*

Proof: Immediate from Schauder's fixed point theorem. \diamond

3.3 Taking φ to zero

Lemma 5 *Let z^k be a fixed point of $TBR^{\varepsilon_k, \varphi_k}$ and suppose $z^k \rightarrow z$. If $(\varepsilon_k, \varphi_k)$ are such that either*

- (a) $\varepsilon_k = \varepsilon$ and $\varphi_k \rightarrow 0$, or
- (b) $\varepsilon_k \rightarrow 0$ and $\varphi_k = 0$.

then z is strictly increasing at all points where there is positive probability of trade in the limiting game.

Proof: Suppose not, and let $[v_i, v'_i]$ be an interval where the buyer's bid function is constant b_i . Then this has to be a region where the truncated and the non-truncated best responses are the same. Then for k large, the k -th bid function will be δ close to this plateau. Because there is positive probability of trade, for k large enough, with probability bounded away from zero there is a seller's bid below $b_i^k(v_i)$. Then it is an event with probability bounded away from zero that all buyers have values in $[v_i, v'_i]$ and all sellers bid below $b_i^k(v_i)$.

But then by increasing her bid by 2δ , a buyer of value v_i could get an incremental probability of winning that is bounded away from zero. For δ small,

the cost of this bid increment in terms of price impact is arbitrarily small. Thus for k large bidding $b_i^k(v_i)$ cannot be optimal for a buyer of value v_i . This is a contradiction. \diamond

The first part of the Lemma allows us to show that the $(\varepsilon, 0)$ -perturbed double auction has a (truncated) fixed point.

Corollary 3 *In a large enough auction, the truncated best response map $TBR^{\varepsilon, \varphi}(\cdot)$ with $\phi = 0$ and ε small enough has a fixed point.*

Proof: Now fix $\varepsilon > 0$ and pick a sequence $\varphi^k \rightarrow 0$, and let z^k be a fixed point of TBR^k . By compactness, the sequence z^k has a convergent subsequence. By relabelling, we can assume that z^k converges to z . Now at all points where there is positive probability of trade, z is strictly increasing. At a point where there is zero probability of trade, as k goes to infinity there had to be vanishingly small probability of trade. But then the ε probability event would gradually overrule any other consideration. Thus at points with zero probability of trade, bidders have to bid their own value. It follows that z is everywhere strictly increasing.

Then by the above lemma, we have that $z \in TBR(z)$. To see why, note that $w^k = BR^k(z^k)$ also has a convergent subsequence. The limit of that sequence, w , has to be a best response to z because z is everywhere strictly increasing. But then the limit of $TBR^k(z^k) = T(w^k)$ has to be $T(w)$ as truncation is a continuous operation. On the other hand, $TBR^k(z^k) = T(w^k) = z^k$, thus the limit of z^k , which is z , is equal to $T(w)$. It follows that $z = T(w)$, in other words, z is indeed a truncated best response to z . \diamond

3.4 Relaxing the Truncation

Theorem 6 *In the $(\varepsilon, 0)$ -perturbed double auction for all sufficiently large N any fixed point of the truncated best response map is a Bayesian Nash equilibrium.*

We prove this result by showing that for all sufficiently large auctions the truncation $|v_i - x_i(v_i)| \leq C$ does not bind.

Consider a fixed point of the ε -perturbed auction. We denote $b_i = x_B(v_i) = x_S(v_S)$. The first-order condition of a buyer i with valuation v_i becomes

$$(1 - \varepsilon) \int_s \left[(v_i - b_i) \left(nK_{n,m}^s(b_i) \frac{f_S(v_S|s)}{x'_S(v_S)} + (m - 1) L_{n,m}^s(b_i) \frac{f_B(v_i|s)}{x'_B(v_i)} \right) - \kappa M_{n,m}^s(b_i) \right] dH(s|v_i) + \varepsilon (v_i - b_i) = 0 \quad (8)$$

where:

$$\begin{aligned}
K_{n,m}^s(b_i) &\equiv \begin{array}{l} \text{the probability that bid } b_i \text{ lies between the } n\text{th and} \\ \text{(n+1)th highest bid in a sample of } m-1 \text{ buyers not} \\ \text{including buyer } i \text{ and } n-1 \text{ sellers and conditional on} \\ \text{the common signal } s \end{array} \\
L_{n,m}^s(b_i) &\equiv \begin{array}{l} \text{the probability that bid } b_i \text{ lies between the } n\text{th and} \\ \text{(n+1)th highest bid in a sample of } m-2 \text{ buyers not} \\ \text{including buyer } i \text{ and } n \text{ sellers and conditional on the} \\ \text{common signal } s \end{array} \\
M_{n,m}^s(b_i) &\equiv \begin{array}{l} \text{the probability that bid } b_i \text{ lies between the (n-1)th and} \\ \text{nth highest bid in a sample of } m-1 \text{ buyers not including} \\ \text{buyer } i \text{ and } n \text{ sellers and conditional on the common} \\ \text{signal } s \end{array}
\end{aligned}$$

These probabilities are defined for the special case of independent private values as formulae A.6 to A.8 in the appendix of Rustichini, Satterthwaite, and Williams (1994). For correlated private values we simply replace $F_B(v_i)$ with $F_B(v_i|s)$ and $F_S(v_S)$ with $F_S(v_S|s)$. Equation 8 consists of two parts: with probability $1 - \varepsilon$ the price in the auction is determined according to the double auction rules. In this case a slight increase in the bid b_i gives the buyer a chance to win an object and gain $v - b_i$. There are n sellers he can surpass and m buyers and the density of a seller's/ buyer's offer at b_i is $\frac{f_S(v_S|s)}{x'_S(v_S)}$ and $\frac{f_B(v_i|s)}{x'_B(v_i)}$ respectively. The probabilities $K_{n,m}^s$ and $L_{n,m}^s$ are the respective probabilities that the other seller and buyers valuations are ordered accordingly. The cost to buyer i of increasing her bid b_i is increasing the price if she is just pivotal which happens with probability $M_{n,m}^s(b_i)$. The second part of equation 8 is the marginal benefit to buyer i of increasing her bid if the price is a uniform draw from the unit interval.

It follows that

$$\begin{aligned}
v_i - b_i &\leq \frac{\kappa(1 - \varepsilon) \int_s M_{n,m}^s(b_i) dH(s|v_i)}{(m-1)(1 - \varepsilon) \int_s L_{n,m}^s(b_i) f_B(v_i|s) dH(s|v_i) + \varepsilon} x'_B(v_i) \\
v - b_i &\leq \frac{\kappa \int_s M_{n,m}^s(b_i) dH(s|v_i)}{(m-1) \int_s L_{n,m}^s(b_i) f_B(v_i|s) dH(s|v_i)} x'_B(v_i) \tag{9}
\end{aligned}$$

Now define $F_{min}(v) = \min_{s,i} F_i(v|s)$ and $F_{max}(v) = \max_{s,i} F_i(v|s)$. Similarly, define $f_{max} = \max_{i,s,v_i} f_i(v_i|s)$ and $f_{min} = \min_{i,s,v_i} f_i(v_i|s)$. Note, that $0 < f_{min} < f_{max} < \infty$.

Lemma 6 *The ratio $M_{n,m}^v(b_i)/L_{n,m}^v(b_i)$ satisfies*

$$\frac{M_{n,m}^v(b_i)}{L_{n,m}^v(b_i)} \leq 2F_{max}(v_i) + \frac{2n(1 - F_{min}(v_i)) F_{max}(v_i)}{m(1 - F_{max}(v_i))} < D \tag{10}$$

for some constant D and $v_i \in [0, 1]$.

Proof The proof for the independent private values case can be found in the appendix of Rustichini, Satterthwaite, and Williams (1994) (proof of inequality 3.12). The proof for correlated private values is exactly the same - we simply integrate all expressions over s . \diamond

We also know that $f_B(v_i|s) \geq f_{min} > 0$ by assumption. Hence we obtain:

$$v_i - b_i < \frac{\kappa D}{f_{min}(m-1)} x'_B(v_i) \quad (11)$$

This implies that for large enough auctions, the truncated best response remains arbitrarily close to the diagonal at all points v_i where the truncation is not effective. Indeed, at points where $x'_B(v_i)$ is less than 1 this follows from the above inequality. At points where $x'_B(v_i)$ is greater than one, this follows because at those points the trajectory $x_B(v_i)$ is actually getting closer to the diagonal! The same argument extends to seller.

If the truncated best response stays closer than $C/2$ at all point where truncation is ineffective then truncation can never apply. At the no trade region for buyers with low values of v_i , the truncated best response is just bidding one's valuation as we have seen. Once we enter the trade region, the only way the truncated best response can start to be truncated is if it hits the truncation border $v_i - C$. But this never happens, as before hitting the border the above inequality applies, and it guarantees that the trajectory does not even get below the $v_i - C/2$ border. \diamond

3.5 Taking ε to zero

The only remaining step is to take ε to zero. Consider a sequence ε^k going to zero, and let z^k be an equilibrium of the ε^k -perturbed game. As usual, we can select a convergent subsequence with a limit z . By Lemma 5 (b), at all points v that have positive probability of trade, z is increasing. By part (c) of Corollary 1, the best response property is preserved under the limit.

For the sake of completeness, here we show that if at value v_i the limiting bid $x_i(v_i)$ is such that no trade occurs, then no trade would occur for this buyer no matter what her bid would be (given her value is v_i). One might think that a bid higher than $x_i(v_i)$ (still not greater than v_i) may actually secure trade with positive probability. But in that case, she should have bid higher for k large enough, because bidding close to her current bid gives vanishingly small utility, whereas e.g., bidding v_i would have given her utility bounded away from zero for k large. This is a contradiction.

It follows that the limiting profile z is a symmetric, increasing equilibrium of the κ -double auction. We have just proved Theorem 1.

4 Examples and Limitations

Independent private values. This case has been studied in Rustichini, Satterthwaite, and Williams (1994) and Williams (1991) among others. For that auction the best responses are automatically increasing, thus a simpler argument may work. Then the only reason to require that the auction size is large enough is to ensure that the truncated best responses actually stay close. If we do not truncate the best responses, just look at all increasing functions that map into $[0, 1]$, this step is not necessary either. Thus we do not require n large enough for existence.

However, the argument that shows the profile to be close to truth-telling, which is based on Rustichini, Satterthwaite, and Williams (1994), continues to depend on n being large. Can we relax that? Intuitively, the ε -perturbations should help, because in an ε -perturbed equilibrium no trade implies bidding v . Hopefully, making ε vanish does not bring about vanishing trade, just like in Jackson and Swinkels (2001). This brings us to conjecture that the proof can be extended to show that in the independent private values case, an equilibrium exists for all auction sizes.

Relaxing symmetry. The current proof relies on all buyers respectively all sellers being identical, and on the symmetric nature of the equilibrium. The usefulness of symmetry is twofold. First, it buys us the fact that profiles remain strictly increasing if trade occurs with positive probability. Potentially a similar result can be achieved without symmetry too: all we need to rule out is flats in the price distribution. A flat occurs if two bidders have plateaus in their bidding functions at the same level b . In such a situation the same argument we used earlier should give a contradiction, potentially with a tie-breaking rule that throws away everything at ties. Second, we need symmetry to ensure that an actual equilibrium stays close, because otherwise we cannot relax the truncation. To relax this assumption we would either need an argument that asymmetric equilibria also stay close (which would be interesting in itself), or to prove that the best response to a close profile is even closer, in which case we would not even need the truncation. The first of these two is more likely to be true.

5 Conclusion

Our paper is an attempt to bridge the gap between bilateral bargaining theory based on incomplete information and price taking behavior with market clearing prices and efficient allocations. We do this by proving the existence of increasing pure strategy equilibria that are close to fully revealing for large enough markets.

References

- ATHEY, S. (2001): “Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information,” *Econometrica*, 69, 861–890.
- CHATTERJEE, K., AND W. SAMUELSON (1983): “Bargaining under Incomplete Information,” *Operations Research*, 31, 835–851.
- JACKSON, M. O., AND J. M. SWINKELS (2001): “Existence of Equilibrium in Single and Double Private Value Auctions,” Working paper, <http://www.olin.wustl.edu/faculty/swinkels/JacksonSwinkels.pdf>.
- MASKIN, E., AND J. RILEY (2000): “Equilibrium in Sealed High Bid Auctions,” *Review of Economic Studies*, 67, 439–454.
- MCADAMS, D. (2003): “Characterizing Equilibria in Asymmetric First-Price Auctions,” Working paper, MIT.
- RENY, P. J., AND M. PERRY (2003): “Toward a Strategic Foundation For Rational Expectations Equilibrium,” Working paper, University of Chicago.
- RUSTICHINI, A., M. A. SATTERTHWAITE, AND S. R. WILLIAMS (1994): “Convergence of Efficiency in a Simple Market with Incomplete Information,” *Econometrica*, 62, 1041–1063.
- WILLIAMS, S. R. (1991): “Existence and Convergence of Equilibria in the Buyer’s Bid Double Auction,” *Review of Economic Studies*, 58, 351–374.
- WILSON, R. (1985): “Incentive Efficiency of Double Auctions,” *Econometrica*, 53, 1101–1115.

A Proof of Theorem 2

The proof of the theorem requires the following Lemma.

Lemma 7 *In the (ε, φ) -perturbed auction, with $\varepsilon, \varphi > 0$, if opponents’ strategies are M' close to the diagonal, then the following inequality holds:*

$$\Pr(b_i \text{ is the } n + 1\text{-st bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \cdot C \geq \Pr(b_i \text{ is the } n\text{-th bid}).$$

Proof: Fix the common part of all valuations, s . Fix a set of n players i_1, \dots, i_n . Suppose the state of the world (conditional on s) is such that exactly these players have bids above b_i . Denote the probability of this event by

$$P^{i_1, \dots, i_n}.$$

Now pick any of these n players, say i_1 , and consider the event when i_1 bids below b_i , while i_2, \dots, i_n are the players who bid above b_i . Call the corresponding new event the transformed event. This event has a probability of at most

$$P^{i_1, \dots, i_n} \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'}$$

where the f terms estimate the density, and the terms involving b_i are bounds on the set of valuations that make a player bid above, respectively below, b_i .

Now pick instead player i_2 , and let her bid below b_i , while the others continue to bid above. We can estimate the probability of this event the same way as above. Let us add up the bounds we have thus derived, to get

$$n \cdot P^{i_1, \dots, i_n} \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \geq Q(i_1, \dots, i_n).$$

In this inequality, $Q(i_1, \dots, i_n)$ denotes the probability of the event where exactly one of the players in the set i_1, \dots, i_n bids below b_i (all the others in the set bid above, and the rest of the players also bid below).

Now let us sum up this last inequality for all subsets i_1, \dots, i_n . This yields

$$n \cdot \Pr(b_i \text{ is the } n + 1\text{-st bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \geq \sum Q(i_1, \dots, i_n).$$

What is the sum on the right hand side? We argue that it covers all states of the world where there are exactly $n - 1$ bids above b_i , and moreover that it counts each such state m times. It is clear that any such state of the world is covered, because there is at least one i_1, \dots, i_n set that can give rise to it. More importantly, we can count the number of ways that state of the world arises in the summation. Clearly, it arises when one of the bids currently below b_i is shifted there from above b_i . There are exactly m bids below b_i ; thus there are m different P^{i_1, \dots, i_n} cases that give rise to this particular state of the world through the event transformation.

This proves that

$$n \cdot \Pr(b_i \text{ is the } n+1\text{-st bid} \mid s) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \geq m \cdot \Pr(b_i \text{ is the } n+1\text{-st bid} \mid s)$$

and now we can integrate over s , to get

$$\Pr(b_i \text{ is the } n + 1\text{-st bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \cdot \frac{n}{m} \geq \Pr(b_i \text{ is the } n\text{-th bid}).$$

Since n/m is bounded away from zero and infinity the claim follows. \diamond

Let us now turn to the proof of the theorem. Define $M = kM'$. We will show that k can be chosen in such a way that the conclusion of the theorem follows. Suppose a buyer i has optimal bid b_i , and $b_i < v_i - M$ (otherwise we are done). We consider whether the buyer would prefer to bid instead $b_i + 4M'$. If yes, that would be a contradiction, showing that the optimal bid in fact has to be at least M close to the diagonal. Our argument is going to make use of a supplementary technical lemma, stated and proved below.

In order to proceed, first note that for any opponent bidding function $x_j(\cdot)$ satisfying the condition of lemma 7, we have the set inclusion

$$x_j^{-1}((b_i, b_i + 3M')) \supseteq (b_i + M', b_i + 2M').$$

Indeed, any value in the interval on the right hand side would induce a bid (both in the case of a buyer and a seller) that is contained in the interval $(b_i, b_i + 3M')$.

Now consider that the buyer bids $b_i + 4M'$. Her gain will be an increased probability of winning. This gain is realized for example if her bid was the $n + 1$ -th highest previously, the φ -perturbation biased the final price upwards, and by increasing the bid, she overtook some opponent by sufficient distance so that even the φ -perturbation cannot make her lose. The size of the gain in this case is at least the distance between her current bid and her value v_i , minus possibly φ . Therefore the following formula is a lower bound for the expected gain from raising the bid

$$\frac{1}{2} \Pr(b_i \text{ is the } n + 1\text{-st bid and } \exists \text{ opponent value in } (b_i + M', b_i + 2M')) \cdot (k - 5)M'.$$

To see why, note that with probability $1/2$ the φ perturbation is biased upwards. Now if there is an opponent value in $(b_i + M', b_i + 2M')$, that leads to a bid no greater than $b_i + 3M'$. If φ is small enough relative to M' , then the realized price will still be lower than $b_i + 4M'$, thus our buyer wins. She wins at least $(k - 4)M' - \varphi > (k - 5)M'$. This proves the above inequality.

Next note that the probability in this formula will be arbitrarily close to $\Pr(b_i \text{ is the } n + 1\text{-st bid})$ as the auction size increases in the sense that

$$\frac{\Pr(b_i \text{ is the } n + 1\text{-st bid and } \exists \text{ opponent value in } (b_i + M', b_i + 2M'))}{\Pr(b_i \text{ is the } n + 1\text{-st bid})} \rightarrow 1.$$

Therefore the gain is bounded from below by

$$\Pr(b_i \text{ is the } n + 1\text{-st bid}) \cdot \frac{(k - 5)M'}{4}$$

for a large enough auction.

Next consider the loss from increasing the bid. A loss will take place when b_i was exactly the n -th bid; the size of the loss is bounded from above by $4M'$. Thus the total expected loss is not more than

$$\Pr(b_i \text{ is the } n\text{-th bid}) \cdot 4M'.$$

We need to compare our bounds for the gain and the loss. This involves comparing the probability terms in the two formulas. Lemma 7 shows that

$$\Pr(b_i \text{ is the } n+1\text{-st bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \cdot C \geq \Pr(b_i \text{ is the } n\text{-th bid})$$

when C is an appropriate constant that does not vary with the size of the auction. Thus the gain will be greater than the loss if

$$\Pr(b_i \text{ is the } n+1\text{-st bid}) \cdot \frac{(k-5)M'}{4} \geq \Pr(b_i \text{ is the } n+1\text{-st bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} 4M'$$

or

$$\frac{(k-5)}{4} \geq 4 \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'}$$

which is implied by

$$\frac{(k-5)}{4} \geq 4 \frac{f_{\max}}{f_{\min}} \cdot \left(\frac{2}{kM'} - 1 \right)$$

for k large enough. For k large, this condition is satisfied when

$$k^2 \geq 64 \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{M'}.$$

Now we can choose k large so that this final inequality holds. Moreover, we can choose k such that $kM' = M$ is still going to zero as M' is going to zero (because k is of order $M'^{-1/2}$).

For the appropriately chosen M the above argument shows that $b_i + 3M'$ is a better bid than b_i . This completes the proof. \diamond

B Proof of Theorem 3

We show the claim by contradiction. Assume there is an agent i and two private values $v_i < v'_i$ with corresponding best responses b_i and b'_i such that $b_i > b'_i$.

Incentive compatibility implies

$$\begin{aligned} P^{v_i}(p(b_i, \omega) < b_i) E^{v_i} [v_i - p(b_i, \omega) | p(b_i, \omega) < b_i] &\geq \\ P^{v_i}(p(b'_i, \omega) < b'_i) E^{v_i} [v_i - p(b'_i, \omega) | p(b'_i, \omega) < b'_i] &\end{aligned} \quad (12)$$

and

$$\begin{aligned} P^{v'_i}(p(b'_i, \omega) < b'_i) E^{v'_i} [v'_i - p(b'_i, \omega) | p(b'_i, \omega) < b'_i] &\geq \\ P^{v'_i}(p(b_i, \omega) < b_i) E^{v'_i} [v'_i - p(b_i, \omega) | p(b_i, \omega) < b_i]. \end{aligned} \quad (13)$$

We can rewrite these conditions as follows

$$\begin{aligned} 0 &\geq \left[\int_{p(b'_i, \omega) < b'_i} (v_i - p(b'_i, \omega)) \, d\mu^{v_i}(\omega) - \int_{p(b_i, \omega) < b_i} (v_i - p(b_i, \omega)) \, d\mu^{v_i}(\omega) \right] \\ 0 &\geq - \left[\int_{p(b'_i, \omega) < b'_i} (v'_i - p(b'_i, \omega)) \, d\mu^{v'_i}(\omega) - \int_{p(b_i, \omega) < b_i} (v'_i - p(b_i, \omega)) \, d\mu^{v'_i}(\omega) \right]. \end{aligned}$$

We introduce the notation $\Delta v_i = v_i - v'_i$ as well as the operators $\Delta_{v_i} A(\tilde{v}) = A(v_i) - A(v'_i)$ and $\Delta_{b_i} A(\tilde{b}) = A(b_i) - A(b'_i)$. We now add the two inequalities and organize terms to obtain

$$\begin{aligned} -\Delta v_i \left[\int_{p(b'_i, \omega) < b'_i} d\mu^{v'_i}(\omega) - \int_{p(b_i, \omega) < b_i} d\mu^{v'_i}(\omega) \right] &\geq \\ \Delta v_i \left[\int_{p(b'_i, \omega) < b'_i} (v_i - p(b'_i, \omega)) \, d\mu^{\tilde{v}}(\omega) - \int_{p(b_i, \omega) < b_i} (v_i - p(b_i, \omega)) \, d\mu^{\tilde{v}}(\omega) \right]. \end{aligned} \quad (14)$$

Further reorganization yields

$$\begin{aligned} \Delta v_i \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right) &\geq \\ \Delta v_i \left[\underbrace{\int_{p(b'_i, \omega) < b'_i} (p(b_i, \omega) - p(b'_i, \omega)) \, d\mu^{\tilde{v}}(\omega)}_{\text{Term I}} - \underbrace{\int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{\tilde{v}}(\omega)}_{\text{Term II}} \right]. \end{aligned} \quad (15)$$

On the right-hand side we made use of the monotonicity of the disturbance φ which ensures that the ω -set $p(b_i, \omega) < b_i$ is larger than $p(b'_i, \omega) < b'_i$ (note that $b_i > b'_i$ by assumption).

Term I is non-negative for all ω . By the argument in lemma 2, we have

$$\begin{aligned} \left| \Delta v_i \int_{p(b'_i, \omega) < b'_i} (p(b_i, \omega) - p(b'_i, \omega)) \, d\mu^{\tilde{v}}(\omega) \right| &\leq \\ K \cdot |\Delta v_i| \cdot \int_{p(b'_i, \omega) < b'_i} |p(b_i, \omega) - p(b'_i, \omega)| \, d\mu^{v_i}(\omega) \end{aligned} \quad (16)$$

for some constant K which is independent of v_i, v'_i, b_i, b'_i . Now (12) is easily seen to imply that

$$\int_{\substack{p(b'_i, \omega) < b'_i \\ p(b_i, \omega) \geq b_i}} (p(b_i, \omega) - p(b'_i, \omega)) \, d\mu^{v_i}(\omega) \leq \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{v_i}(\omega)$$

Furthermore, lemma 2 also allows us to bound term II

$$\left| \Delta v_i \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{\tilde{v}}(\omega) \right| \leq K \cdot |\Delta v_i| \cdot \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{v_i}(\omega).$$

It follows from all of these above that the right hand side RHS of (15) can be estimated as

$$|RHS| \leq 2K |\Delta v_i| \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{v_i}(\omega) \quad (17)$$

The integral is evaluated over an ω -set where the inequality $v_i > b_i > p(b_i, \omega) > p(b'_i, \omega) > b'_i$ holds. We therefore know that $v_i - p(b_i, \omega) < v_i - b'_i$. This allows us to simplify the above inequality further

$$|RHS| \leq 2K |\Delta v_i| (v_i - b'_i) \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} d\mu^{v_i}(\omega) \quad (18)$$

By Lipschitz continuity we know that

$$|h(s|v_i) - h(s|v'_i)| \leq Kh(s|v'_i). \quad (19)$$

This allows us to simplify the inequality further to get

$$\begin{aligned} |RHS| &\leq 2K(K+1) |\Delta v_i| (v_i - b'_i) \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} d\mu^{v'_i}(\omega) \\ &= 2K(K+1) |\Delta v_i| (v_i - b'_i) \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right). \end{aligned} \quad (20)$$

Now note that $|\Delta v_i| = -\Delta v_i$ such that we can plug it back into (15) and obtain

$$\Delta v_i \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right) \geq 2K(K+1) \Delta v_i (v_i - b'_i) \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right) \quad (21)$$

which yields

$$\Delta v_i \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right) [1 - 2K(K+1)(\Delta v_i + v'_i - b'_i)] \geq 0. \quad (22)$$

We can actually choose $\Delta v_i < \frac{1}{4K(K+1)}$ because if agent i 's best response function is non-monotonic then we can find v and v' arbitrarily close such that $v_i < v'_i$ and $b_i > b'_i$. Therefore the final inequality implies that as long as

$$v'_i - b'_i < \frac{1}{4K(K+1)} \quad (23)$$

holds we have

$$\Delta_{b_i} P^{v'_i} \left(p(\tilde{b}|\omega) < \tilde{b} \right) \leq 0. \quad (24)$$

But this implies $b_i \leq b'_i$ which is a contradiction. \diamond

C Proof of Lemma 4

The argument in the main text shows that the payoff to bidding b_i in the k -th game converges to that of bidding b_i in the limit game. That is, the payoff function is convergent pointwise. However, this is not quite enough to show that the maximum also converges. For this we need a stronger result, which is the following.

Let $b_i^k \rightarrow b_i$. Then

$$\Pi^k(v_i, b_i^k) \rightarrow \Pi(v_i, b_i).$$

In words, the payoff from bidding b_i^k in game k converges to that of bidding b_i in the limiting game. To see why this is true, let us rewrite the payoff as an integral on the probability space:

$$\Pi^k(v_i, b_i^k) = \int_{\{p^k(b_i^k) < b_i\}} v_i - p^k(b_i^k, \omega) d\mu^{v_i}(\omega)$$

where $p^k(b_i^k, \omega)$ is the price (a random variable) in game k . Note that the price explicitly depends on our buyer's bid b_i . Denote the domain of integration in the above formula by $Z^k = \{\omega | p^k(b_i^k, \omega) < b_i^k\}$, and fix δ a small positive number. Then

$$\begin{aligned} \Pi^k(v_i, b_i^k) &= \int_{\{p(b_i) < b_i - \delta\} \cap Z^k} v_i - p^k(b_i^k, \omega) d\mu^{v_i}(\omega) + \\ &+ \int_{\{b_i - \delta \leq p(b_i) \leq b_i + \delta\} \cap Z^k} v_i - p^k(b_i^k, \omega) d\mu^{v_i}(\omega) + \\ &+ \int_{\{p(b_i) > b_i + \delta\} \cap Z^k} v_i - p^k(b_i^k, \omega) d\mu^{v_i}(\omega). \end{aligned}$$

Note that because the distribution of $p(b_i)$ is atomless, the middle term can be made arbitrarily small by an appropriate choice of δ . Fix a δ . We also have that $p^k(b_i^k, \omega)$ converges to $p(b_i, \omega)$ in probability; therefore for k large enough (given δ), outside of a small probability event we will have that $\omega \in \{p(b_i) < b_i - \delta\}$ implies $\omega \in Z_k$. Therefore, controlling for the small approximation error, the first term can be considered to be integrated over $\{p(b_i) < b_i - \delta\}$. Likewise, for k large, the domain of integration of the final term will have arbitrarily small measure. Therefore we can write that

$$\Pi^k(v_i, b_i^k) = \text{small}(\delta) + \text{small}(k, \text{given } \delta) + \int_{\{p(b_i) < b_i - \delta\}} v - p^k(b_i^k, \omega) d\mu^{v_i}(\omega).$$

Furthermore, because the price distribution is atomless, the domain of integration in this formula can be replaced by $\{p(b_i) < b_i\}$, that introduces approximation errors smaller than what we currently have. Finally, because $p^k(b_i^k, \omega)$ converges in probability to $p(b_i, \omega)$, we can write that

$$\begin{aligned} \Pi^k(v_i, b_i^k) &= \text{small}(\delta) + \text{small}(k, \text{given } \delta) + \int_{\{p(b_i) < b_i\}} v - p(b_i, \omega) d\mu^{v_i}(\omega) = \\ &= \text{small}(\delta) + \text{small}(k, \text{given } \delta) + \Pi(v_i, b_i). \end{aligned}$$

This is what we wanted to prove. By choosing δ small enough, and then accordingly k large enough, we can show that $\Pi^k(v_i, b_i^k)$ gets arbitrarily close to $\Pi(v_i, b_i)$.

Now to get the statement of the lemma, assume that b_i^k is the best response in game k , but b_i , the limit, is not a best response in the limiting game. Then there is a b_i' that does better than b_i , so that $\Pi(v_i, b_i) < \Pi(v_i, b_i')$. However, in game k it has to be the case that $\Pi^k(v_i, b_i^k) \geq \Pi^k(v_i, b_i')$. Taking k to infinity, the left hand side converges to $\Pi(v_i, b_i)$, and the right hand side to $\Pi(v_i, b_i')$, thus giving $\Pi(v_i, b_i) \geq \Pi(v_i, b_i')$. This is a contradiction. \diamond