Existence of Equilibrium in Large Double Auctions

Citation

Fudenberg, Drew, Markus M. Mobius, and Adam Szeidl. 2007. Existence of equilibrium in large double auctions. *Journal of Economic Theory* 133(1): 550-567.

Published Version

<http://dx.doi.org/10.1016/j.jet.2005.07.014>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3043420>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Existence of Equilibrium in Large Double Auctions*

by

Drew Fudenberg

Department of Economics, Harvard University, Cambridge, MA 02138

Markus Mobius

Department of Economics, Harvard University, Cambridge, MA 02138

and

Adam Szeidl[†]

Department of Economics, University of California at Berkeley, Berkeley, CA 94720

First version: March 13, 2003

This version: October 24, 2005

Abstract

We show the existence of a pure strategy, symmetric, increasing equilibrium in double auction markets with correlated, conditionally independent private values and many participants. The equilibrium we find is arbitrarily close to fully revealing as the market size grows. Our results provide strategic foundations for price-taking behavior in large markets.

JEL Classification: C62, C72, D44, D82

Keywords: price-taking, pure strategy equilibrium, monotone strategies, large economies

*E-mail addresses: dfudenberg@harvard.edu, mobius@fas.harvard.edu, szeidl@econ.berkeley.edu. We thank Eric Maskin, Andy Postlewaite, Phil Reny, Jeroen Swinkels, two anonymous referees and an associate editor for helpful comments and suggestions. This paper was inspired by Reny's presentation of Reny and Perry (2003) in the fall of 2002. Fudenberg acknowledges support from the National Science Foundation grant SES-0426199, Szeidl thanks the Social Science Research Council.

[†]Corresponding author. Tel.: +1 510 642 2603; fax: +1 510 642 6615.

1 Introduction

This paper establishes the existence of pure-strategy equilibria for large double auctions with correlated, conditionally independent private values. In these equilibria, bids are very close to valuations, and so can be interpreted as approximately truthful reports of the agents' information. Thus the equilibrium we find approximates price-taking behavior in large markets.

The main difficulty in our proof is that without independence of values, standard arguments showing the monotonicity of best response bidding functions break down. Even when all opponents make use of increasing bidding functions, Reny and Perry [11] construct an example for the interdependent values case where the best response of an agent is non-monotonic.¹ Intuitively, observing a higher value has two effects on the bidding behavior of a buyer. First, a higher valuation makes submitting higher bids more attractive. Second, a higher valuation reveals information about opponent values and bids, and this information may lead to a lower best response bid.

We proceed by establishing that this “information effect” is small in auctions with many participants. We do this by first showing that, because each agent is rarely pivotal in large auctions, best responses to strategies that are approximately truthful are themselves approximately truthful. We then prove that the size of the information effect is proportional to the misrepresentation of an agent's bid, because the effect operates through changing the likelihood that the agent is pivotal. Since in large auctions bids are approximately truthful, the information effect is small, and monotonicity obtains.

A second difficulty is that when ties occur with positive probability, best responses need not exist, because then a bidder would like to undercut the tie by an arbitrary positive amount. We deal with this problem by introducing small stochastic perturbations to the original game which make ties zero probability events, and thus smooth out the discontinuities in payoff functions. In these perturbed games best response bidding functions exist, and as argued above, with many participants they are also increasing. We then establish existence of equilibrium in the perturbed game using a fixed-point theorem on the Banach space of increasing bidding functions. At this step, we face the additional difficulty that the image of the best response map may not be contained in its domain. We truncate best responses before applying the fixed point theorem to deal with this problem. Then, extending an argument of Rustichini, Satterthwaite and Williams [12], we bound the degree to which agents misrepresent their valuations in the resulting profile. This bound shows that for large auctions the truncation does not bind, hence our profile is an equilibrium of the perturbed game. Because the bound is uniform in the size of the perturbation, sending the perturbations to zero gives an equilibrium of the unperturbed double auction in which all players bid approximately truthfully. The Rustichini, Satterthwaite and Williams argument also provides

¹See footnote 29 of [11].

a rate of convergence to efficiency of the equilibrium.²

As compared to past work, our main contribution is to provide a shorter and hopefully clearer proof of the existence of equilibrium in the case of correlated values. Because the proof relies on the fact that the information effect is small in large economies, it proves monotonicity at the same time that it proves existence, which may better highlight the role of the assumption that the economy is large. The idea of our proof is to restrict attention to strategies in a neighborhood of the (known) equilibrium of the limit game, prove that best responses in the finite economies are well behaved in this neighborhood, and use that fact to show that a fixed point exists. This approach may be helpful in other settings where there is a well-defined notion of a limit economy and where the equilibrium of the limit game is well behaved.

Jackson and Swinkels [6] prove the existence of a (non-trivial) mixed-strategy equilibrium in a variety of auctions with any fixed number of agents by taking the limits of equilibria in auctions with a discretized space of bids. They assume private values, and that the distribution of values is not perfectly correlated, and prove existence by appealing to papers by Jackson, Simons, Swinkels and Zame [7] and Reny [10] on the existence of equilibria in discontinuous games. They prove existence of pure strategy equilibria only for the case where values are independent. Reny and Perry [11], in an affiliated interdependent values setting, use an approximation of the strategy space by finite grids to show that all discretized auctions in sufficiently large economies have an equilibrium with non-decreasing bid functions; this equilibrium approximates the rational expectations equilibrium of the continuum limit. This paper was inspired by an early version of Reny and Perry; we obtain a somewhat stronger form of monotonicity without the use of grids and with a much shorter proof.³ Reny and Perry look at economies where all the values of all agents are drawn from the same distribution. We relax this by allowing for buyers and sellers values to be drawn from a finite number of distributions, which roughly corresponds to the idea of replica economies in the proof of the core convergence theorem (Debreu and Scarf, [5]).

The κ -double auction was introduced to the literature by Chatterjee and Samuelson [3], and by Wilson [15] for the multilateral case. In the independent private values case, Williams [14] shows that a pure-strategy equilibrium exists in the buyer's bid double auction, and Satterthwaite and Williams [13] provide a rate of convergence. In that auction, sellers always bid their valuations, which makes the existence argument easier. Rustichini, Satterthwaite and Williams extended the Satterthwaite and Williams convergence result to general double auctions, showing that if a symmetric equilibrium exists, it must be close to truth telling, and provided a bound on rate of convergence. They did not prove existence; we extend their analysis to the correlated private-values setting in this paper. Cripps and Swinkels [4] showed

²See [12] for a discussion of the importance of establishing a rate.

³Athey's [1] existence theorem for one-sided auctions also makes use of monotonicity to appeal to a fixed point theorem; she obtains monotonicity from a single-crossing assumption.

that in a broad class of private value auctions all non-trivial equilibria are asymptotically efficient. Their assumption of "z-independence" is more general than our assumption of a finite number of buyer and seller classes.⁴ Thus the main contribution of our paper is the existence result; we also extend Rustichini, Satterthwaite and Williams [12] to correlated values.

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 player 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 tie breaking will not be relevant.⁵ For a more detailed discussion of the κ -double auction please see [12].

We make a number of assumptions about the joint distribution of valuations. We focus on correlated private values. We will let s be a random variable that captures the common component of the valuations. The distribution of s is $G(s)$, which is assumed to be concentrated on the unit interval (the substantive part of this assumption is compact support) and absolutely continuous with respect to the Lebesgue measure. Conditional on s , the valuations of all buyers and sellers are independent. To simplify notation, we focus on the case where all buyers' values are drawn from the same conditional probability distribution which we denote $F_B(v|s)$, and that the values of all sellers are drawn from $F_S(v|s)$. We explain at each step why the argument generalizes to the case where there are a fixed, finite number of distributions F_B^k, F_S^l , with a fixed fraction of agents drawing values from each distribution. Let $f_B(v|s)$ and $f_S(v|s)$ be the densities of F_B and F_S . We assume that all valuations are concentrated on the unit interval, with densities 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.

⁴The assumption implies that there is some residual uncertainty about a given agent's value even given the values of an arbitrarily large number of other agents.

⁵This mechanism 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.

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 independent private values case (see [12] and [14]).

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 .

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. Whenever we use the term "for large enough auctions", what we have in mind is increasing $N = n + m$ while keeping $\gamma = n/N$ bounded away from zero and one. Just how large the auction needs to be depends on how tight these bounds are.

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 a sequence of equilibrium profiles x^N for N large enough such that*

$$\max_i \sup_v |x_i^N(v) - v| = O(1/N).$$

Hence in these equilibria, individual bids are of order $1/N$ close to being fully revealing.

3 Proof steps

3.1 Perturbation

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 $\int_{b_i}^1 (p - v_i) dp$. Note that under this perturbation, as well as under the φ -perturbation introduced below, 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. 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

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

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 and is always in the unit interval; (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 strictly increasing in both p and ω_φ except when $p = 0$ or $p = 1$; (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), (2) and (4) make the expected payoffs a continuous function of the bids. Property (3) guarantees that the median perturbation is zero. Property (4) also 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.

In the rest of the paper, we require that buyers submit bids that are not greater than their valuations, and similarly, we require that sellers play strategies that are not smaller than their valuations. Because with $\varepsilon > 0$ the strategies thus ruled out are strictly dominated, this assumption is not substantive, but makes it easier to state some of our results, for example Proposition 1 below, which we would otherwise have to state as applying to best responses to undominated strategies.

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

$$\mu^{v_i}(A) = \frac{\mu(A \cap \{\omega | v_i = v_i(\omega)\})}{\mu(\{\omega | v_i = v_i(\omega)\})}. \quad (1)$$

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 $E^{v_i}(X) = \int X(\omega) d\mu^{v_i}(\omega)$.

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 $p(b_i)$.

3.2 Best response of the perturbed game

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

Proof. All proofs not given in the main text are presented in the appendix. ■

Because the φ -perturbation makes ties a zero probability event, the payoff function of each bidder becomes continuous. Continuous functions defined on compact sets have a maximum, and therefore best responses in the perturbed auction exist.

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.

Proposition 2 *For any $M > 0$ there exist $M', N_0, \varepsilon_0, \varphi_0 > 0$ such that in a perturbed auction with $N > N_0$, $0 < \varphi < \varphi_0$ and $0 < \varepsilon < \varepsilon_0$, for all bidders i , if opponents’ strategies satisfy $\sup_v |x_{-i}(v) - v| < M'$, any best response of i satisfies $\sup_v |x_i(v) - v| < M$.*

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.

The following is 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 1 *There exists a constant K (independent of (ε, φ) and N) such that for any positive function $u(s)$*

$$\int_s u(s) |h(s|v'_i) - h(s|v_i)| ds \leq K \cdot |v'_i - v_i| \cdot \int_s u(s) h(s|v_i) ds. \quad (2)$$

Proof. We have $|h(s|v'_i) - h(s|v_i)| \leq K \cdot |v'_i - v_i| \cdot h(s|v_i)$ for any s and v_i, v'_i because h is uniformly Lipschitz in the second argument and bounded away from zero. ■

Theorem 2 *In the (ε, φ) -perturbed auction with $\varepsilon, \varphi > 0$, for any bidder i , if a best response x_i to some opponent profile x_{-i} satisfies $\sup_{v_i} |v_i - x_i(v_i)| < 1/4K(K+1)$, then $x_i(\cdot)$ is increasing.*

The theorem claims that if a best response is close enough to truth-telling, it has to be increasing. To get some intuition, consider the first-order condition of a buyer i

$$T(b_i, v_i) = (v_i - b_i) r(b_i, v_i) - \kappa R(b_i, v_i) = 0, \quad (3)$$

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 (that is, we use the

Implicit Function Theorem) to obtain

$$\frac{\partial T}{\partial b_i} x'_i(v_i) + \frac{\partial T}{\partial v_i} = 0. \quad (4)$$

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 prove that i 's best response $x_i(v_i)$ is increasing. We can calculate

$$\frac{\partial T}{\partial v_i} = \underbrace{r(b_i, v_i)}_{\text{Term I}} + \underbrace{(v_i - b_i) \frac{\partial r}{\partial v_i}}_{\text{Term II}} - \underbrace{\kappa \frac{\partial R}{\partial v_i}}_{\text{Term III}}. \quad (5)$$

Also, note that

$$r(b_i, v_i) = \int_0^1 r(b_i, s) h(s|v_i) ds \text{ and } R(b_i, v_i) = \int_0^1 R(b_i, s) h(s|v_i) ds \quad (6)$$

where $r(b_i, s)$ is the conditional density for some fixed s and $R(b_i, s)$ is the conditional pivotal probability. Lemma 1 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 (3), $R(b_i, v_i)$ 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.

Terms II and III in this analysis capture the “information effect” of increasing a bidder’s valuation that we discussed in the introduction. Indeed, in the special case of independent private values terms II and III are zero, because changes in player i ’s value do not reveal information about opponents, and therefore have no effect on the conditional probability and density of the pivotal bid. In this case, only term I remains and the best response of player i is always increasing in v_i . With correlated values, changes in v_i reveal information about opponents, and this affects optimal bidding through changes in the distribution of pivotal bids as captured by terms II and III. However, when the best response of i is close to truth-telling, the effect of a change in the probability of being pivotal is small, because in that event, the difference between the valuation and the bid, which is the benefit of winning, is small. Note, this argument did not rely on having only two distributions, one for buyers and one for sellers.

Corollary 1 *There exists $C > 0$ such that for N large enough, φ and ε small enough, for all bidders i , the best response x_i to any profile x_{-i} that satisfies $\sup_v |v_j - x_j(v_j)| < C$ for all $j \neq i$ is unique and increasing.*

Proof. Combining the previous result and Proposition 2 shows that there exists $C > 0$ such that if x_{-i} satisfies the condition of the corollary, any best response x_i is increasing. Fix

such a C , and 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, $v_0 \in (0, 1)$, of both x_i and x'_i , where they differ, say $x_i(v_0) > x'_i(v_0)$. By continuity, 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 which shows that the best response x_i is unique. ■

3.3 Continuity and 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 increasing functions defined on the unit interval with values in the C wide corridor around the diagonal. Formally,

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

Furthermore, define X_B and X_S to be the subsets of X corresponding to profiles that are (weakly) below, respectively above, the diagonal, that is $X_B = X \cap \{x_i : [0, 1] \rightarrow [0, 1] \mid x_i(v_i) \leq v_i\}$ and $X_S = X \cap \{x_i : [0, 1] \rightarrow [0, 1] \mid x_i(v_i) \geq v_i\}$.

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 simply 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 $\max(x_i(v_i), v_i - C, 0)$. Importantly, the truncated best response is still increasing, since it is the max (min in the case of a seller) of increasing functions.

Lemma 2 *X , X_B and X_S are compact, convex subsets of the Banach space $L^1[0, 1]$.*

Proof. By Helly's theorem (see Billingsley [2]), 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 functions involved are in X . Since X_B and X_S are convex and closed subsets of X , the conclusion follows for them too. ■

In the rest of the paper we will focus on symmetric profiles, where all players who have the same distribution of values use the same bidding function. In the case where all buyers and all sellers have identical distributions, we only need to keep track of the pair $x = (x_B(\cdot), x_S(\cdot))$ of a buyer's and a seller's strategy. Then we have $x \in X_B \times X_S$. With multiple groups of players, we would use a direct product of more than two sets.

It is clear that $X_B \times X_S$ 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 $TBR^{\varepsilon, \varphi}(\cdot) : X_B \times X_S \rightarrow X_B \times X_S$. Here $TBR^{\varepsilon, \varphi}(x_B, x_S)$ 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, which requires first establishing that the map is continuous in the L^1 topology. Note that the L^1 topology restricted to a set of uniformly bounded functions like X is the same as convergence in measure. This implies that we can focus on continuity of $TBR^{\varepsilon, \varphi}$ in measure.

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. 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 . The price in game k is a random variable, whose distribution depends on the bid of player i , b_i . To emphasize this dependence, in the following we will denote the price by $p^k(b_i)$. The price in the limit game is denoted by $p(b_i)$.

Proposition 3 *Suppose $x^k \rightarrow x$ in measure, and $y^k \in BR^k(x^k)$ converges in measure to y . Consider the sequence of random variables $p^k(b_i)$ obtained when i bids b_i while opponents play x_{-i}^k . If for all i and b_i , (1) the price $p^k(b_i)$ converges in probability to $p(b_i)$; (2) the distribution of $p(b_i)$ has no atom at b_i ; then $y \in BR(x)$.*

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 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 proposition 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). Equivalently, her expected payoff is strictly larger than $\varepsilon \cdot \int_0^{b_i} (v_i - p) dp$ for a buyer, or strictly larger than $\varepsilon \cdot \int_{b_i}^1 (p - v_i) dp$ for a seller.

Corollary 2 *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 $y \in BR(x)$.

This result shows that when $\varepsilon, \varphi > 0$, $TBR^{\varepsilon, \varphi}(\cdot)$ is a continuous self-map of a compact, convex subset of a Banach space.

Theorem 3 *There exist $\varepsilon_1, \varphi_1, N_1 > 0$ such that in a perturbed auction with $N > N_1$, $0 < \varepsilon < \varepsilon_1$ and $0 < \varphi < \varphi_1$ the truncated best response map $TBR^{\varepsilon, \varphi}(\cdot)$ has a fixed point.*

Proof. Immediate from Schauder’s fixed point theorem. ■

3.4 Relaxing the perturbations

We begin by taking φ to zero. By compactness, we can select a convergent subsequence from the sequence of fixed points associated with some $\varphi^k \rightarrow 0$. The only reason why the best response property may fail in the limit is that ties may become a positive probability event when φ becomes zero. This possibility is ruled out by part (a) of the following lemma.

Lemma 3 *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.

Thus, the limiting profile is everywhere strictly increasing (in the no trade region, this is ensured by the ε perturbation), so even when $\varphi = 0$, ties have zero probability. The best response property of the limiting profile follows:

Proposition 4 *In a large enough auction, the truncated best response map $TBR^{\varepsilon, \varphi}(\cdot)$ with $\varphi = 0$ and ε small enough has a fixed point.*

Next we relax the truncation.

Proposition 5 *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.*

This result follows from the fact that for all sufficiently large auctions the truncation $|v_i - x_i(v_i)| \leq C$ does not bind. The proof builds on an argument of Rustichini, Satterthwaite and Williams, who show that in the independent private values case the first order condition implies that a symmetric equilibrium must satisfy $|v_i - x_i(v_i)| \leq K_3/N$ for some constant K_3 . Their argument extends for the current correlated values environment, even allowing for the ε -perturbation and multiple buyer and seller groups, and implies that for large auctions, any symmetric increasing profile that satisfies the first order condition almost everywhere will satisfy $|v_i - x_i(v_i)| \leq C/2$. One can then apply this result for the fixed point of the truncated best response map to prove that the truncation does not bind. To see why, note that in the no trade region for buyers with low values of v_i , the truncated best response is just bidding one's valuation. 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 first order condition holds almost everywhere, which implies that $|v_i - x_i(v_i)| \leq C/2$. Hence the trajectory does not even get below the $v_i - C/2$ border. The details of the proof are contained in supplementary material at <http://www.nyu.edu/jet/supplementary.html>. The above argument also implies that $|v_i - x_i(v_i)| \leq K_3/N$ holds for the profile that is the fixed point of the truncated best response map. This shows that the distance between valuations and bids is of order $O(1/N)$ in the candidate equilibrium we consider.

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 3 (b), at all points v that have positive probability of trade, z is strictly increasing. By part (c) of Corollary 2, the best response property is preserved under the limit. It follows that the limiting profile z is a symmetric, increasing equilibrium of the κ -double auction. We have just proved Theorem 1.

4 Conclusion

Affiliated values. We conjecture that our proof technique can be extended to the affiliated values case, where the analog of "bidding truthfully" is "bidding one's value conditional on being pivotal," but we have not been able to provide a proof. The main difficulty seems to lie in providing the appropriate extension of the Rustichini, Satterthwaite and Williams characterization: Showing that in equilibrium bids are close to the diagonal is easier than showing that bids are close to players' conditional expected values.

Appendix

Proof of Proposition 1

Fix a buyer i with valuation v_i , and consider the function $W(b, c) = \Pr(p(b) < c)$ where $b = b_i$ is the bid of buyer i . We begin by showing that $W(b, c)$ is continuous in both b and c . Note that

$$W(b, c) = \varepsilon c + (1 - \varepsilon) \int_{p_0 \in [0, 1]} \Pr(\tilde{p}(p_0, \omega_\varphi) < c) \, dQ_0(p_0; b)$$

where p_0 is the preliminary price, $Q_0(\cdot; b)$ is the distribution of p_0 given that buyer i bids b (and all opponents use their strategies) and $p = \tilde{p}(\cdot, \cdot)$ is the φ -perturbed price used to calculate payoffs.

The function $\tilde{p}(\cdot, \cdot)$ is by definition continuous and strictly increasing in both arguments, except when p_0 equals zero or one. It follows that $\Pr(\tilde{p}(p_0, \omega_\varphi) < c)$ is continuous in c for each $p_0 \neq 0, 1$. But the preliminary price is almost surely different from zero and one by our assumption that all opponent buyers and sellers bid weakly below, respectively above, their valuations. Hence the integral is also continuous in c . To check continuity in b , note that the map $b \rightarrow Q_0(\cdot; b)$ from bids to distribution functions is continuous in the weak topology, because the preliminary price is a continuous function of bids. Because $\Pr(\tilde{p}(p_0, \omega_\varphi) < c)$ is a bounded, continuous function (in p_0), it follows that the integral is continuous in b .

Since $W(b, c)$ is continuous in both arguments and increasing in c , it is easy to see that for $b_k \rightarrow b$, we have $W(b_k, \cdot) \rightarrow W(b, \cdot)$ in the uniform topology. The payoff of i from bidding b equals $\int_{p \in [0, 1]} v_i - p \, dW(b, p)$ which is easily shown to be continuous in b , by the uniform convergence of the functions $W(b, \cdot)$. A symmetric argument applies for a seller, or for multiple buyer and seller groups.

Proof of Proposition 2

Define $f_{\min} = \min [\inf_{v, s} f_B(v|s), \inf_{v, s} f_S(v|s)]$ and $f_{\max} = \max [\sup_{v, s} f_B(v|s), \sup_{v, s} f_S(v|s)]$ which constitute a lower and upper bound for the conditional densities of buyers and sellers. By assumption, $0 < f_{\min} \leq f_{\max} < \infty$.

Lemma 4 ⁶ *There exists K_2 independent of N , ε and φ , such that with $\varepsilon, \varphi > 0$, for all i , v_i and b_i , if opponents' strategies are M' close to the diagonal while player i bids b_i , then*

$$\Pr(b_i \text{ is the } n+1\text{-st bid} | v_i) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} \cdot K_2 \geq \Pr(b_i \text{ is the } n\text{-th bid} | v_i).$$

Proof. Fix s , and let q_j denote the probability that player j bids above b_i conditional on s . Then

$$\frac{q_j}{1 - q_j} \geq \frac{f_{\min}}{f_{\max}} \cdot \frac{1 - b_i - 3M'}{b_i + 3M'}$$

⁶We thank an anonymous referee for suggesting the current proof of this lemma, which is much shorter than our earlier version.

for all j . For if j has value greater than $b_i + 2M'$, then she certainly bids above b_i , because bids are M' close to valuations. Thus the conditional probability that j bids above b_i is at least $f_{\min} \cdot (1 - b_i - 3M')$, and likewise, the conditional probability that j bids below b_i is at most $f_{\max} \cdot (b_i + 3M')$.

The probability that $n - 1$ opponents bid above b_i is

$$\left(\prod_{j \neq i} (1 - q_j) \right) \cdot \sum_{C_{n-1}} \left[\prod_{j \in C_{n-1}} \frac{q_j}{1 - q_j} \right]$$

where the summation is over all subsets C_{n-1} of $n - 1$ opponents (excluding player i). Multiplying each term in the above sum by the lower bound for $q_j/(1 - q_j)$ will be an underestimate of the product $\prod_{j \in C_n} q_j/(1 - q_j)$ for some subset of n opponents C_n . Note that there are m subsets of opponents C_n that contain any given set with $n - 1$ elements C_{n-1} , because the total number of opponents is $m + n - 1$. Hence

$$m \cdot \frac{f_{\min}}{f_{\max}} \cdot \frac{1 - b_i - 3M'}{b_i + 3M'} \cdot \left(\prod_{j \neq i} (1 - q_j) \right) \cdot \sum_{C_{n-1}} \left[\prod_{j \in C_{n-1}} \frac{q_j}{1 - q_j} \right] \leq n \left(\prod_{j \neq i} (1 - q_j) \right) \cdot \sum_{C_n} \prod_{j \in C_n} \frac{q_j}{1 - q_j}.$$

The factor n on the right hand side is included because all sets in C_n are covered n times, as each such set has n different subsets with $n - 1$ elements. Taking expectations over s conditional on v_i , the above inequality can be written as

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

Since n/m is bounded away from zero and infinity, the claim follows. This argument can easily be extended for more general replica economies with several different conditional distributions. The only change is that $f_{\max}$ and $f_{\min}$ have to be redefined and the exact values of these terms are irrelevant for the main results. ■

To prove the proposition, define $M = kM'$. We need to find a k such that the result holds. 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.

Note that for any opponent bidding function $x_j(\cdot)$ satisfying the condition of Lemma 4, we have the set inclusion $x_j^{-1}((b_i, b_i + 3M')) \supseteq (b_i + M', b_i + 2M')$. This is because 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 $(b_i, b_i + 3M')$.

Assume 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 $(k-4)M' - \varphi > (k-5)M'$.

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 $\Pr(b_i \text{ is the } n+1\text{-st bid}) \cdot (k-5)M'/4$ 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. Using Lemma 4, the gain will be greater than the loss if

$$\Pr(b_i \text{ is the } n+1\text{-st bid}) \cdot \frac{(k-5)M'}{4} \geq \Pr(b_i \text{ is the } n\text{-th bid}) \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{b_i + 3M'}{1 - b_i - 3M'} K_2 \cdot 4M'$$

where K_2 does not vary with the size of the auction. This inequality is implied by

$$\frac{(k-5)}{4} \geq 4K_2 \frac{f_{\max}}{f_{\min}} \cdot \left(\frac{2}{kM'} - 1 \right)$$

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

$$k^2 \geq 64K_2 \cdot \frac{f_{\max}}{f_{\min}} \cdot \frac{1}{M'}.$$

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 $b_i + 4M'$ is a better bid than b_i .

Proof of Theorem 2

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

$$\begin{aligned} P^{v_i}(p(b_i, \omega) < b_i) E^{v_i} [v_i - p(b_i, \omega) | p(b_i, \omega) < b_i] &\geq \\ P^{v_i}(p(b'_i, \omega) < b'_i) E^{v_i} [v_i - p(b'_i, \omega) | p(b'_i, \omega) < b'_i] \end{aligned} \quad (7)$$

and

$$\begin{aligned} P^{v'_i}(p(b'_i, \omega) < b'_i) E^{v'_i} [v'_i - p(b'_i, \omega) | p(b'_i, \omega) < b'_i] &\geq \\ P^{v'_i}(p(b_i, \omega) < b_i) E^{v'_i} [v'_i - p(b_i, \omega) | p(b_i, \omega) < b_i]. \end{aligned} \quad (8)$$

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

$$\begin{aligned} -\Delta v_i \left[\int_{p(b'_i, \omega) < b'_i} d\mu^{v'_i}(\omega) - \int_{p(b_i, \omega) < b_i} d\mu^{v'_i}(\omega) \right] &\geq \\ \Delta_{v_i} \left[\int_{p(b'_i, \omega) < b'_i} (v_i - p(b'_i, \omega)) \, d\mu^{\tilde{v}}(\omega) - \int_{p(b_i, \omega) < b_i} (v_i - p(b_i, \omega)) \, d\mu^{\tilde{v}}(\omega) \right]. \end{aligned} \quad (9)$$

Further reorganization yields

$$\begin{aligned} \Delta_{v_i} \Delta_{b_i} P^{v'_i} (p(\tilde{b}, \omega) < \tilde{b}) &\geq \\ \Delta_{v_i} \left[\underbrace{\int_{p(b'_i, \omega) < b'_i} (p(b_i, \omega) - p(b'_i, \omega)) \, d\mu^{\tilde{v}}(\omega)}_{\text{Term I}} - \underbrace{\int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{\tilde{v}}(\omega)}_{\text{Term II}} \right]. \end{aligned} \quad (10)$$

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 1 , we have

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

for some constant K which is independent of v_i, v'_i, b_i, b'_i . Now (7) is easily seen to imply that

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

Furthermore, lemma 1 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 (10) can be estimated as

$$|RHS| \leq 2K |\Delta v_i| \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} (v_i - p(b_i, \omega)) \, d\mu^{v_i}(\omega) \quad (11)$$

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

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

By Lipschitz continuity we know that $|h(s|v_i) - h(s|v'_i)| \leq Kh(s|v'_i)$. This allows us to simplify the inequality further to get

$$\begin{aligned} |RHS| &\leq 2K(K+1) |\Delta v_i| (v_i - b'_i) \int_{\substack{p(b_i, \omega) < b_i \\ p(b'_i, \omega) \geq b'_i}} d\mu^{v'_i}(\omega) \\ &= 2K(K+1) |\Delta v_i| (v_i - b'_i) \Delta_{b_i} P^{v'_i} \left(p(\tilde{b}, \omega) < \tilde{b} \right). \end{aligned} \quad (12)$$

Now note that $|\Delta v_i| = -\Delta v_i$ such that we can plug it back into (10) and obtain

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

which yields

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

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

$$v'_i - b'_i < \frac{1}{4K(K+1)} \quad (13)$$

holds we have $\Delta_{b_i} P^{v'_i} (p(\tilde{b}|\omega) < \tilde{b}) \leq 0$. But this implies $b_i \leq b'_i$ which is a contradiction.

Proof of Proposition 3

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 converges pointwise. However, this is not quite enough to show that the maximum also converges.

Let $b_i^k \rightarrow b_i$. We claim that in this case, $\Pi^k(v_i, b_i^k) \rightarrow \Pi(v_i, b_i)$, that is, the payoff from bidding b_i^k in game k converges to that of bidding b_i in the limiting game. The proof is as follows. First, rewrite the payoff as an integral on the probability space:

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

where $p^k(b_i^k, \omega)$ is the price (a random variable) in game k . 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

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

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 $p(b, \omega)$, 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 $\Pi(v_i, b_i)$.

To get the statement of the proposition, 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.

Proof of Corollary 2

We will use Proposition 3. 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 proposition, $\Pi^k(v_i, b'_i) \rightarrow \Pi(v_i, b'_i)$. Now if no other player bids b with positive probability, then by the argument of the proposition 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 .

Proof of Lemma 3

Suppose not, and let $[v_i, v'_i]$ be an interval where the buyer's bid function is constant b_i . This has to be a region where the truncated and the non-truncated best responses are the same. Then for any $\delta > 0$, there is k large enough such that the k -th bid function will be within a distance of δ from plateau in the range $[v_i, v'_i]$. Because there is positive probability of trade, for k large enough, with probability bounded away from zero there is a seller's bid below $b_i^k(v_i)$. Hence the event that all buyers have values in $[v_i, v'_i]$ and all sellers bid below $b_i^k(v_i)$ has probability bounded away from zero.

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. The proof extends to multiple groups because a positive probability of trade implies that there exists an event with probability bounded away from zero where bidding $b_i^k(v_i) + 2\delta$ and thus beating all other buyers of the same group would deliver the good to v_i but bidding only $b_i^k(v_i)$ would not.

Proof of Proposition 4

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 relabeling, 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. Hence for k large, the impact of trading with opponents in the $(1 - \varepsilon)$ probability event has vanishingly small impact on the best response of a bidder. Thus at points with zero probability of trade, as k goes to infinity bidders will only consider the ε probability event and therefore in the limit they have to bid their own value. It follows that z is everywhere strictly increasing.

Then by Lemma 3, 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 .

References

- [1] S. Athey, Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information, *Econometrica*, 69 (2001), 861–890.

- [2] P. Billingsley, Probability and Measure, 2nd. Edition, Wiley, New York, 1986.
- [3] K. Chatterjee, W. Samuelson, Bargaining under Incomplete Information, Operations Research 31 (1983), 835–851.
- [4] M. W. Cripps, J. M. Swinkels, Efficiency in Large Double Auctions, Econometrica, forthcoming, 2004.
- [5] G. Debreu, H. Scarf, A Limit Theorem on the Core of an Economy, Int. Econ. Rev. 4 (1963), 235–246.
- [6] M. O. Jackson, J. M. Swinkels, Existence of Equilibria in Single and Double Private Value Auctions, Econometrica 73 (2005), 93–139.
- [7] M. O. Jackson, L.K. Simon, J.M. Swinkels, W.R. Zame, Communication and Equilibrium in Discontinuous Games of Incomplete Information, Econometrica 70 (2002), 1711 – 1740.
- [8] E. Maskin, J. Riley, Equilibrium in Sealed High Bid Auctions, Rev. Econ. Stud. 67 (2000), 439–454.
- [9] D. McAdams, Characterizing Equilibria in Asymmetric First-Price Auctions, MIT Sloan Working paper 4256-02 (2002).
- [10] P. J. Reny, On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games, Econometrica 67 (1999), 1029–1056.
- [11] P. Reny, M. Perry, Toward a Strategic Foundation For Rational Expectations Equilibrium, Working paper, University of Chicago (2003).
- [12] A. Rustichini, M. A. Satterthwaite, S. R. Williams, Convergence to Efficiency in a Simple Market with Incomplete Information, Econometrica 62 (1994), 1041–1063.
- [13] M. A. Satterthwaite, S. Williams, The Rate of Convergence to Efficiency in the Buyer’s Bid Double Auction as the Market Becomes Large, Rev. Econ. Stud. 56 (1989), 477–498.
- [14] S. R. Williams, Existence and Convergence of Equilibria in the Buyer’s Bid Double Auction, Rev. Econ. Stud. 58 (1991), 351–374
- [15] R. Wilson, Incentive Efficiency of Double Auctions, Econometrica 53 (1985), 1101–1115.