Econometrica Supplementary Material

SUPPLEMENT TO "EQUILIBRIUM SELECTION IN AUCTIONS AND HIGH STAKES GAMES"

(*Econometrica*, Vol. 86, No. 1, January 2018, 219–261)

PAUL MILGROM
Department of Economics, Stanford University

JOSHUA MOLLNER
Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University

1. PACKAGE AUCTION MODEL

THE PACKAGE AUCTION MODEL is a generalization of the menu auction model, in which each bidder cares about and bids for only some part of the allocation. As one example, there may be a set of goods to be allocated among bidders, with each bidder caring about and bidding for only its own allocation or "package."

Does Corollary 5—that test-set equilibrium leads to core payoffs in the menu auction setting—extend to the full set of package auctions? In this section, we demonstrate by example that the answer is no.[1]

As before, there is one auctioneer, who selects a decision that affects himself and N bidders. The possible packages for bidder n are given by the set X_n and the possible choices for the auctioneer by $X \subseteq \bigtimes_{n=1}^{N} X_n$. The gross monetary payoffs that bidder n receives are described by the function $v_n : X_n \to \mathbb{R}$. Similarly, the auctioneer receives gross monetary payoffs described by $v_0 : X \to \mathbb{R}$.[2]

The N bidders simultaneously submit bids. A bid is a menu of payments to the auctioneer, contingent on the package received, which can be expressed as a function $b_n : X_n \to \mathbb{R}_+$. Given bids, the auctioneer chooses a decision $x \in X$ that maximizes his payoff $v_0(x) + \sum_{n=1}^{N} b_n(x_n)$. As before, we assume that if there are several such decisions, then the auctioneer chooses the one that maximizes the total surplus. We also continue to assume that all agents have lexicographic preferences, first preferring outcomes with the highest personal payoff and secondarily preferring ones with higher total surplus. And we also assume as before that against any bid profile in which at least one competing bidder is playing a strictly dominated strategy, each bidder strictly prefers to set its bid vector equal to its value vector over any other bid vector that leads to the same auctioneer decision and the same zero payoff.

Paul Milgrom: milgrom@stanford.edu
Joshua Mollner: joshua.mollner@kellogg.northwestern.edu

For helpful comments, we thank Gabriel Carroll, Ricardo De la O, Piotr Dworczak, Drew Fudenberg, Philippe Jehiel, Peter Klibanoff, Fuhito Kojima, Markus Mobius, Michael Ostrovsky, James Schummer, Erling Skancke, Andrzej Skrzypacz, Joel Sobel, Bruno Strulovici, Péter Vida, Jörgen Weibull, Glen Weyl, Alexander Wolitzky, seminar participants, and anonymous referees. Milgrom thanks the National Science Foundation for support under Grant 1525730. Mollner was a Postdoctoral Researcher at Microsoft Research while part of this research was completed, and he thanks them for their hospitality.

[1]A further generalization is to the class of core-selecting mechanisms (Day and Milgrom (2008)). The example therefore also implies that test-set equilibrium does not necessarily lead to core payoffs in all core-selecting mechanisms.

[2]The package auction model is equivalent to the menu auction in the special case where $X_1 = \cdots = X_N$ and X is the diagonal subset of $\bigtimes_{n=1}^{N} X_n$.

 https://doi.org/10.3982/ECTA12536

While Bernheim and Whinston (1986) demonstrated that truthful equilibrium and coalition-proof equilibrium lead to bidder-optimal core payoffs only in the menu auction setting, these results generalize to all package auctions. In contrast, Corollary 5—that test-set equilibrium leads to core payoffs in the menu auction setting—does not generalize to all package auctions. Example 1, below, describes a package auction that possesses a test-set equilibrium with non-core payoffs.

In the example, there are six bidders, with possible packages $X_n = \{l, w\}$ ("lose" or "win"). The set X includes six combinations of packages, describing which sets of bidders can simultaneously win. First, bidder 1 may win alone. Alternatively, bidder 2 may win together with one of bidders 3, 4, 5, or 6. Finally, bidders 3, 4, 5, and 6 may win together. The last of these possibilities maximizes total surplus. However, there exists a test-set equilibrium implementing the decision in which bidder 1 wins alone.

EXAMPLE 1: Let $N = 6$. For all n, let $X_n = \{l, w\}$. Let

$$X = \left\{ \begin{array}{c} (w, l, l, l, l, l), (l, w, w, l, l, l), (l, w, l, w, l, l), \\ (l, w, l, l, w, l), (l, w, l, l, l, w), (l, l, w, w, w, w) \end{array} \right\}.$$

For all $x \in X$, let $v_0(x) = 0$. Let the payoffs of the bidders be as follows:

$$\begin{array}{ll} v_1(l) = 0, & v_1(w) = 29, \\ v_2(l) = 0, & v_2(w) = 19, \\ v_3(l) = 0, & v_3(w) = 9, \\ v_4(l) = 0, & v_4(w) = 8, \\ v_5(l) = 0, & v_5(w) = 7, \\ v_6(l) = 0, & v_6(w) = 6. \end{array}$$

Then the following bid profile is a test-set equilibrium, which results in the inefficient decision (w, l, l, l, l, l):

$$\begin{array}{ll} b_1(l) = 0, & b_1(w) = 28, \\ b_2(l) = 0, & b_2(w) = 19, \\ b_3(l) = 0, & b_3(w) = 9, \\ b_4(l) = 0, & b_4(w) = 0, \\ b_5(l) = 0, & b_5(w) = 0, \\ b_6(l) = 0, & b_6(w) = 0. \end{array}$$

PROOF OF EXAMPLE 1: It is easy to verify that these bids result in the decision (w, l, l, l, l, l). This decision is inefficient because it yields a total surplus of 29, whereas the decision (l, l, w, w, w, w) yields a total surplus of 30. It is also easy to check that these bids are a Nash equilibrium. Moreover, in the equilibrium no bidder is using a bid that is weakly dominated in the game, by an extension of Lemma 2.

Bidder 1 has a unique best response to the equilibrium bids of the other bidders: its equilibrium bid of $b_1(l) = 0$ and $b_1(w) = 28$. The best responses of bidder 2 are those of the form $b_2(l) = 0$ and $b_2(w) \in [0, 19]$. For all bidders $n \in \{3, 4, 5, 6\}$, their best responses are those of the form $b_n(l) = 0$ and $b_n(w) \in [0, 9]$. Given this, it is easily checked that each bidder's test-set condition is satisfied. *Q.E.D.*

This package auction example helps to sharpen our understanding of how test-set equilibrium promotes core allocations in the original Bernheim–Whinston menu auction model. Test-set equilibrium is effective there because it promotes "high enough" bids for losing decisions. It does so because each bidder n believes that a deviation by a single other bidder playing a different best response might offer an opportunity for a better outcome, provided that n bids high enough. In this package auction example, however, bidders 4, 5, and 6 are not bidding "high enough," yet there is no element in the test set that offers an opportunity for a better outcome. No single deviation can create such an opportunity; only a joint deviation by two or more others could do that.

Coalition-proof equilibrium refines away this package-auction equilibrium because it considers the possibility of a cooperative joint deviation. Truthful strategies work as a refinement in this context, too, because the restriction to truthful bids is a restriction to bids that are high enough for losing decisions. The test-set refinement for these package auction games, however, does not imply high bids for losing decisions.

2. PROPER EQUILIBRIA OF THE AGENT-NORMAL FORM OF THE SECOND-PRICE, COMMON-VALUE AUCTION

Although we have been unable to characterize the pure proper equilibria of the normal form of the second-price, common-value auction described in Section 4.2—which is a discrete version of the motivating example of Abraham et al. (2016)—we do have such a characterization for the agent-normal form. In the agent-normal form of this auction, the pure test-set equilibria and the pure proper equilibria coincide.

PROPOSITION 12: *There exist two pure proper equilibria of the agent-normal form of the discrete second-price*, *common-value auction described in Section* 4.2: $(0, 1, 0)$ *and* $(0, 1, 1)$.

PROOF OF PROPOSITION 12: The proof consists of two parts. First, we construct sequences of trembles that justify each of $(0, 1, 0)$ and $(0, 1, 1)$ as proper equilibria. Second, we argue that no other pure strategy profile is a proper equilibrium. For brevity, we focus, in each part, on the case where the discretized bid set is $\{0, \frac{1}{m}, \dots, \frac{m-1}{m}, 1\}$.

Part One: For all sufficiently small values of ε, the following completely mixed strategy profile is an ε-proper equilibrium: (i) for $k \in \{1, \dots, m\}$, the low-type informed bidder places probability ε^k on k/m and places all remaining probability on 0, (ii) for $k \in \{0, \dots, m-1\}$, the high-type informed bidder places probability ε^{2m-k} on k/m and places all remaining probability on 1, and (iii) for $k \in \{1, \dots, m\}$, the uninformed bidder places probability ε^k on k/m and places all remaining probability on 0. As ε converges to zero, these ε-proper equilibria converge to $(0, 1, 0)$, which is therefore a proper equilibrium.

For all sufficiently small values of ε, the following completely mixed strategy profile is an ε-proper equilibrium: (i) for $k \in \{1, \dots, m\}$, the low-type informed bidder places probability ε^{m+k} on k/m and places all remaining probability on 0, (ii) for $k \in \{0, \dots, m-1\}$, the high-type informed bidder places probability ε^{m-k} on k/m and places all remaining probability on 1, and (iii) for $k \in \{0, \dots, m-1\}$, the uninformed bidder places probability ε^{m-k} on k/m and places all remaining probability on 1. As ε converges to zero, these ε-proper equilibria converge to $(0, 1, 1)$, which is therefore a proper equilibrium.

Part Two: Let $\{\varepsilon_t\}_{t=1}^{\infty}$ be a sequence of positive numbers converging to zero and let $\{(\sigma_0^t, \sigma_1^t, \sigma_U^t)\}_{t=1}^{\infty}$ be a sequence of completely mixed strategy profiles such that for each t, $(\sigma_0^t, \sigma_1^t, \sigma_U^t)$ is an ε_t-proper equilibrium.

Suppose $b, b' \in \{0, \frac{1}{m}, \dots, \frac{m-1}{m}, 1\}$ with $b < b'$. Because σ_U^t is completely mixed, the low-type informed bidder receives a strictly higher payoff from b than from b' against (σ_1^t, σ_U^t).

Thus, ε_t-properness requires $\sigma_0^t(b') \le \varepsilon_t \sigma_0^t(b)$. Likewise, the high-type informed bidder receives a strictly higher payoff from b' than from b against (σ_0^t, σ_U^t), and we conclude $\sigma_1^t(b) \le \varepsilon_t \sigma_1^t(b')$.

The uninformed bidder's payoff from bidding k/m against (σ_0^t, σ_1^t) is

$$\pi_U^t\left(\frac{k}{m}\right) = \frac{1}{2}\sum_{k'<k}\left[\sigma_0^t\left(\frac{k'}{m}\right)\left(-\frac{k'}{m}\right) + \sigma_1^t\left(\frac{k'}{m}\right)\left(1-\frac{k'}{m}\right)\right] + \frac{1}{4}\left[\sigma_0^t\left(\frac{k}{m}\right)\left(-\frac{k}{m}\right) + \sigma_1^t\left(\frac{k}{m}\right)\left(1-\frac{k}{m}\right)\right].$$

We argue that if (σ_0^t, σ_1^t) satisfies the restrictions described above and if the index t is sufficiently large, then π_U^t cannot be maximized at k/m for any $k \in \{2, \dots, m-1\}$. Suppose to the contrary that it were. This implies $\pi_U^t(\frac{k}{m}) \ge \pi_U^t(\frac{k+1}{m})$, or equivalently,

$$\sigma_0^t\left(\frac{k+1}{m}\right)\left(\frac{k+1}{m}\right) + \sigma_0^t\left(\frac{k}{m}\right)\left(\frac{k}{m}\right) \ge \sigma_1^t\left(\frac{k+1}{m}\right)\left(1-\frac{k+1}{m}\right) + \sigma_1^t\left(\frac{k}{m}\right)\left(1-\frac{k}{m}\right).$$

Applying the above restrictions on (σ_0^t, σ_1^t), we obtain

$$\sigma_0^t\left(\frac{k}{m}\right)\left[\frac{k}{m} + \varepsilon_t\left(\frac{k+1}{m}\right)\right] \ge \sigma_1^t\left(\frac{k}{m}\right)\left[\left(1-\frac{k}{m}\right) + \frac{1}{\varepsilon_t}\left(1-\frac{k+1}{m}\right)\right]. \tag{12}$$

Similarly, π_U^t being maximized at k/m with $k \in \{2, \dots, m-1\}$ also implies $\pi_U^t(\frac{k}{m}) \ge \pi_U^t(\frac{k-1}{m})$, or equivalently,

$$\sigma_0^t\left(\frac{k}{m}\right)\left(\frac{k}{m}\right) + \sigma_0^t\left(\frac{k-1}{m}\right)\left(\frac{k-1}{m}\right) \le \sigma_1^t\left(\frac{k}{m}\right)\left(1-\frac{k}{m}\right) + \sigma_1^t\left(\frac{k-1}{m}\right)\left(1-\frac{k-1}{m}\right).$$

Applying again the above restrictions on (σ_0^t, σ_1^t), we obtain

$$\sigma_0^t\left(\frac{k}{m}\right)\left[\frac{k}{m} + \frac{1}{\varepsilon_t}\left(\frac{k-1}{m}\right)\right] \le \sigma_1^t\left(\frac{k}{m}\right)\left[\left(1-\frac{k}{m}\right) + \varepsilon_t\left(1-\frac{k-1}{m}\right)\right]. \tag{13}$$

Together, (12) and (13) further imply

$$\frac{\frac{k}{m} + \varepsilon_t\left(\frac{k+1}{m}\right)}{\left(1-\frac{k}{m}\right) + \frac{1}{\varepsilon_t}\left(1-\frac{k+1}{m}\right)} \ge \frac{\frac{k}{m} + \frac{1}{\varepsilon_t}\left(\frac{k-1}{m}\right)}{\left(1-\frac{k}{m}\right) + \varepsilon_t\left(1-\frac{k-1}{m}\right)},$$

which is a contradiction for sufficiently large indices t. *Q.E.D.*

3. SUFFICIENT CONDITIONS FOR EXISTENCE OF TEST-SET EQUILIBRIUM

Despite the fact that test-set equilibria may fail to exist in finite games (cf. Section 6), there are classes of games in which a test-set equilibrium always exists. Proposition 13 states three conditions, each of which is sufficient to guarantee the existence of a test-set equilibrium. One sufficient condition is for the game to have two players: proper equilibria exist and by Proposition 1 are also test-set equilibria in such games. A second sufficient condition is for the game to have three players, each of whom has two pure strategies. Together, these two results imply that the game in Figure 1 is the smallest possible game in which a test-set equilibrium may fail to exist. A third sufficient condition is for the game to be a potential game (Monderer and Shapley (1996)): any strategy profile that maximizes the potential function is a test-set equilibrium.

In addition, Proposition 14 states that in generic games, *every* Nash equilibrium is a test-set equilibrium.

PROPOSITION 13: *A finite game in normal form has at least one test-set equilibrium if it also satisfies at least one of the following conditions*:

(i) *it is a two-player game*,
(ii) *it is a three-player game in which each player has at most two pure strategies*, *or*
(iii) *it is a potential game*.

PROOF OF PROPOSITION 13:

Part One: This follows immediately from Proposition 1 and the existence of proper equilibria in finite games.

Part Two: We show that for three-player games in which each player has two pure strategies, test-set equilibrium is implied by extended proper equilibrium (Milgrom and Mollner (2017)). The result will then follow from the existence of extended proper equilibrium in finite games, which we establish in that paper. Since extended proper equilibrium requires players to use strategies that are undominated in the game, it suffices to establish that the test-set condition must hold.

Consider a three-player game with strategy sets $S_n = \{a_n, b_n\}$. Suppose by way of contradiction that σ is an extended proper equilibrium of this game that fails the test-set condition. Without loss of generality, suppose it is player 1 for whom the test-set condition fails. Then there exists $\hat{\sigma}_1 \in \Delta(S_1)$ that weakly dominates σ_1 in $T(\sigma)$. Also without loss of generality, suppose that $(\sigma_1, a_2, \sigma_3)$ is an element of the test set against which $\hat{\sigma}_1$ strictly outperforms σ_1. Thus, $a_2 \in BR_2(\sigma_1, \sigma_3)$, and

$$\pi_1(\hat{\sigma}_1, a_2, \sigma_3) > \pi_1(\sigma_1, a_2, \sigma_3). \tag{14}$$

Now if $\sigma_2(a_2) = 1$, then equation (14) contradicts Nash equilibrium. Therefore, $\sigma_2(a_2) < 1$, which implies $b_2 \in BR_2(\sigma_1, \sigma_3)$. Then by the failure of the test-set condition for player 1, we also have

$$\pi_1(\hat{\sigma}_1, b_2, \sigma_3) \geq \pi_1(\sigma_1, b_2, \sigma_3). \tag{15}$$

Now if $\sigma_2(a_2) > 0$, then equations (14) and (15) would together contradict Nash equilibrium. Thus, $\sigma_2 = b_2$. The remainder of the argument splits into two cases. In the first, $BR_3(\sigma_1, \sigma_2) = \{a_3, b_3\}$. In the second, $BR_3(\sigma_1, \sigma_2)$ is a singleton.

The first case is that $BR_3(\sigma_1, \sigma_2) = \{a_3, b_3\}$. Then

$$T(\sigma) = \{(a_1, b_2, \sigma_3), (b_1, b_2, \sigma_3), (\sigma_1, a_2, \sigma_3), (\sigma_1, b_2, \sigma_3), (\sigma_1, b_2, a_3), (\sigma_1, b_2, b_3)\}.$$

Then by the failure of player 1's test-set condition, we also have

$$\pi_1(\hat{\sigma}_1, b_2, a_3) \geq \pi_1(\sigma_1, b_2, a_3), \tag{16}$$

$$\pi_1(\hat{\sigma}_1, b_2, b_3) \geq \pi_1(\sigma_1, b_2, b_3). \tag{17}$$

Suppose that for some $\alpha > 0$, σ^ε is a sequence of (α, ε)-extended proper equilibria converging to σ. Then,

$$\pi_1(\hat{\sigma}_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) - \pi_1(\sigma_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) = \left[\pi_1(\hat{\sigma}_1, a_2, \sigma_3^\varepsilon) - \pi_1(\sigma_1, a_2, \sigma_3^\varepsilon)\right]\sigma_2^\varepsilon(a_2) + \left[\pi_1(\hat{\sigma}_1, b_2, \sigma_3^\varepsilon) - \pi_1(\sigma_1, b_2, \sigma_3^\varepsilon)\right]\sigma_2^\varepsilon(b_2),$$

which is positive for sufficiently small values of ε, and which therefore constitutes a contradiction to σ being an extended proper equilibrium. To see that this is positive:

(i) Equation (14) implies that the first term is positive for completely mixed σ_3^ε sufficiently close to σ_3.

(ii) Equations (16) and (17) imply that the second term is nonnegative.

Without loss of generality, the second case is $BR_3(\sigma_1, \sigma_2) = \{a_3\}$, so that $\sigma_3 = a_3$. Suppose that for some $\alpha > 0$, σ^ε is a sequence of (α, ε)-extended proper equilibria converging to σ. Then,

$$\begin{aligned}
\pi_1(\hat{\sigma}_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) - \pi_1(\sigma_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) &= \left[\pi_1(\hat{\sigma}_1, a_2, a_3) - \pi_1(\sigma_1, a_2, a_3)\right]\sigma_2^\varepsilon(a_2)\sigma_3^\varepsilon(a_3) \\
&\quad + \left[\pi_1(\hat{\sigma}_1, b_2, a_3) - \pi_1(\sigma_1, b_2, a_3)\right]\sigma_2^\varepsilon(b_2)\sigma_3^\varepsilon(a_3) \\
&\quad + \left[\pi_1(\hat{\sigma}_1, a_2, b_3) - \pi_1(\sigma_1, a_2, b_3)\right]\sigma_2^\varepsilon(a_2)\sigma_3^\varepsilon(b_3) \\
&\quad + \left[\pi_1(\hat{\sigma}_1, b_2, b_3) - \pi_1(\sigma_1, b_2, b_3)\right]\sigma_2^\varepsilon(b_2)\sigma_3^\varepsilon(b_3) \\
&= \sigma_2^\varepsilon(a_2)\sigma_3^\varepsilon(a_3)\Bigg\{\left[\pi_1(\hat{\sigma}_1, a_2, a_3) - \pi_1(\sigma_1, a_2, a_3)\right] \\
&\quad + \left[\pi_1(\hat{\sigma}_1, a_2, b_3) - \pi_1(\sigma_1, a_2, b_3)\right]\frac{\sigma_3^\varepsilon(b_3)}{\sigma_3^\varepsilon(a_3)} \\
&\quad + \left[\pi_1(\hat{\sigma}_1, b_2, a_3) - \pi_1(\sigma_1, b_2, a_3)\right]\frac{\sigma_2^\varepsilon(b_2)}{\sigma_2^\varepsilon(a_2)} \\
&\quad + \left[\pi_1(\hat{\sigma}_1, b_2, b_3) - \pi_1(\sigma_1, b_2, b_3)\right]\frac{\sigma_2^\varepsilon(b_2)\sigma_3^\varepsilon(b_3)}{\sigma_2^\varepsilon(a_2)\sigma_3^\varepsilon(a_3)}\Bigg\},
\end{aligned}$$

which is positive for sufficiently small values of ε, and which therefore constitutes a contradiction to σ being an extended proper equilibrium. To see that this is positive:

(i) Equation (14) implies that the first term inside the braces is positive.

(ii) Equation (15) implies that the second term inside the braces is nonnegative.

(iii) Since $\sigma_2 = a_2$, $\frac{\sigma_2^\varepsilon(b_2)}{\sigma_2^\varepsilon(a_2)}$ converges to zero. Because payoffs in the game are bounded, this implies that the third term inside the braces converges to zero as well.

(iv) Finally, since $a_2 \in BR_2(\sigma_1, \sigma_3)$ while $b_3 \notin BR_3(\sigma_1, \sigma_2)$, the definition of (α, ε)-extended proper equilibrium requires that $\frac{\sigma_3^\varepsilon(b_3)}{\sigma_2^\varepsilon(a_2)}$ converges to zero. Moreover, since $\sigma_3^\varepsilon(a_3)$ converges to 1, this implies that the fourth term inside the braces converges to zero as well.

Part Three: We use the fact that any potential game is strategically equivalent to a game in which there exists a function P, referred to as the *potential function* of the game, which is such that for all n and for all $s \in \prod_{n=1}^N S_n$, $\pi_n(s) = P(s)$. It is known that any finite potential game possesses a pure strategy trembling-hand perfect equilibrium $s^* \in \arg\max_{s \in \prod_{n=1}^N S_n} P(s)$ (Carbonell-Nicolau and McLean (2014)). Since s^* is trembling-hand perfect, it is also a Nash equilibrium in undominated strategies. We claim that s^* is a test-set equilibrium; to show this, it only remains to check the test-set condition.

Let $n \in \mathcal{N}$ and $\sigma' \in T(s^*)$. By definition, there exist a player m and a strategy $\hat{s}_m \in BR_m(s^*_{-m})$ such that $\sigma' = \sigma/\hat{s}_m$. Let $\bar{P} = P(s^*)$ be the maximum potential. Since $\hat{s}_m \in BR_m(s^*_{-m})$, we have $P(s^*/\hat{s}_m) = P(s^*) = \bar{P}$. Moreover, for any $\hat{\sigma}_n \in \Delta(S_n)$, we must have $P(\sigma'/\hat{\sigma}_n) \leq \bar{P}$. Therefore, $P(\sigma'/s^*_n) \leq P(\sigma'/\hat{\sigma}_n)$, as desired. *Q.E.D.*

PROPOSITION 14: *For almost all finite games in normal form, every Nash equilibrium is a test-set equilibrium.*[3]

PROOF OF PROPOSITION 14: A Nash equilibrium σ is *quasi-strict* if for each player n, each element of $BR_n(\sigma_{-n})$ is in the support of σ_n. Harsanyi (1973) established that, for almost all finite games in normal form, every Nash equilibrium is quasi-strict.[4] Similarly, Kreps and Wilson (1982) established that, for almost all finite games in normal form, every Nash equilibrium is trembling-hand perfect. We argue that if both of these conditions are satisfied, as is the case for almost all finite games, then every Nash equilibrium is a test-set equilibrium. In such games, every Nash equilibrium is both a trembling-hand perfect equilibrium—and therefore an equilibrium in undominated strategies—and a quasi-strict Nash equilibrium. To complete the proof, we argue that every quasi-strict Nash equilibrium satisfies the test-set condition.

Let σ be quasi-strict equilibrium. Suppose by way of contradiction that the test-set condition does not hold. Then without loss of generality, there exist players 1 and 2 and a strategy $\hat{\sigma}_1 \in \Delta(S_1)$ such that (i) for some $s_2 \in BR_2(\sigma_{-2})$, $\pi_1(\hat{\sigma}_1, s_2, \sigma_{-12}) > \pi_1(\sigma_1, s_2, \sigma_{-12})$, and (ii) for all $s_2 \in BR_2(\sigma_{-2})$, $\pi_1(\hat{\sigma}_1, s_2, \sigma_{-12}) \geq \pi_1(\sigma_1, s_2, \sigma_{-12})$. Because σ is quasi-strict, $\text{supp}(\sigma_2) = BR_2(\sigma_{-2})$. Thus, the above conditions imply that $\pi_1(\hat{\sigma}_1, \sigma_{-1}) > \pi_1(\sigma_1, \sigma_{-1})$, which contradicts that σ is a Nash equilibrium. *Q.E.D.*

REFERENCES

ABRAHAM, I., S. ATHEY, M. BABAIOFF, AND M. GRUBB (2016): "Peaches, Lemons, and Cookies: Designing Auction Markets With Dispersed Information," Working Paper. [3]

BERNHEIM B. D., AND M. D. WHINSTON (1986): "Menu Auctions, Resource Allocation, and Economic Influence," *The Quarterly Journal of Economics*, 101 (1), 1–32. [2]

CARBONELL-NICOLAU, O., AND R. P. MCLEAN (2014): "Refinements of Nash Equilibrium in Potential Games," *Theoretical Economics*, 9 (3), 555–582. [6]

DAY, R., AND P. MILGROM (2008): "Core-Selecting Package Auctions," *International Journal of Game Theory*, 36 (3), 393–407. [1]

HARSANYI, J. C. (1973): "Oddness of the Number of Equilibrium Points: A New Proof," *International Journal of Game Theory*, 2 (1), 235–250. [7]

KREPS, D. M., AND R. WILSON (1982): "Sequential Equilibria," *Econometrica*, 50 (4), 863–894. [7]

MILGROM, P., AND J. MOLLNER (2017): "Extended Proper Equilibrium," Working Paper, http://ssrn.com/abstract=3035565. [5]

MONDERER, D., AND L. S. SHAPLEY (1996): "Potential Games," *Games and Economic Behavior*, 14 (1), 124–143. [5]

Co-editor Joel Sobel handled this manuscript.

Manuscript received 27 May, 2014; final version accepted 6 September, 2017; available online 12 September, 2017.

[3]We define "almost all games" in the same sense as Harsanyi (1973).

[4]Harsanyi (1973) himself referred to such equilibria as "quasi-strong."