Equilibrium Selection in Auctions and High Stakes Games*

Paul Milgrom Joshua Mollner

September 13, 2017

Abstract

We introduce the *test-set equilibrium* refinement of Nash equilibrium to formalize the idea that players contemplate only deviations from equilibrium play in which a single competitor plays a non-equilibrium best response. We then apply this refinement to three well-known auction games, comparing our findings to similar ones previously obtained by specialized equilibrium selections. We also introduce a theory of high stakes versions of games, in which strategies are first proposed and then subjected to a potentially costly review-and-revise process. We demonstrate a sense in which the test-set equilibria emerge from such processes when the cost of revision is small.

Keywords: equilibrium refinement, test-set equilibrium, quasi-perfect equilibrium, menu auction, generalized second-price auction, second-price common value auction

1 Introduction

Many auction games studied by economists have multiple Nash equilibria, each with the feature that some bidders have many alternative best responses to equilibrium play. In first-price auctions for a single item, a losing bidder's best responses include all sufficiently low bids; in second-price auctions, the winning bidder's best responses include all sufficiently high bids; and in generalized second-price auctions used for internet advertising, every bidder's set of best responses to equilibrium is typically an open interval of bids. For the multi-item auction mechanisms that are commonly used, for example to allocate radio spectrum, electricity or Treasury securities, each bidder can bid for different quantities or different combinations of lots, with some bids winning and others losing. In equilibria of these auctions, a bidder's set of best responses often contains an open set in a multi-dimensional space. All these games have a continuum of pure Nash equilibria.

Previous analyses of these games have often narrowed the set of equilibrium outcomes by applying various specialized refinements, which are defined in terms of the specific game and do not apply to general non-cooperative games—or even to general finite games. This

*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 number 1525730.

suggests several questions: Can one achieve a similar narrowing using generally applicable refinements like proper equilibrium (Myerson, 1978) or strategically stable equilibrium (Kohlberg and Mertens, 1986)? If not, why not? Is there any consistent principle that can reconcile, or possibly unify, the different specialized refinements that have been applied to different auction games? For the first question, our examples below show that both proper equilibrium and strategically stable equilibrium sometimes fail to narrow the equilibrium set enough to generate conclusions like those obtained by previous analyses. For the second, we will argue that the answer lies in the fact that for games with three or more players, these two older tremble-based refinements fail to impose restrictions on the relative probabilities of trembles by different players. Finally, we will argue that the several specialized refinements can often be understood as imposing restrictions of that kind.

This paper, however, focuses mainly on the third question. We introduce a new Nash equilibrium refinement that we call "test-set equilibrium," which implicitly restricts the relative probabilities of trembles by different players. A strategy profile is a test-set equilibrium if it is a Nash equilibrium in undominated strategies with an additional property: each player's strategy is also undominated when tested against a limited set of profiles. That "test set" consists of the equilibrium strategy profile and every profile in which all players but one play their equilibrium strategies, while a single deviator plays a different best response to the equilibrium. Implicit in this definition is that all players believe that a tremble by one player to some best response to equilibrium play is more likely than any tremble to an inferior response, whether by the same player or a different one, and also more likely than any combination of trembles by two or more players.

Test-set equilibrium may have bite when players have many alternative best responses to equilibrium play, as is often the case in auctions. Applying the test-set refinement to three well-known auction games, we show that it delivers results that are closely comparable to those of the previous specialized refinements. Following the previous papers, we focus on refining the set of pure equilibria. Although our equilibrium selections are not exact matches in any of these auction games, the differences are small and highlight possible limitations of the original analyses.

- Test-set equilibrium is *less* restrictive than the *truthful equilibrium* refinement proposed by Bernheim and Whinston (1986) to study their *menu auction*, and unlike their alternative *coalition-proof equilibrium*, it imposes no constraints on the profitability of joint deviations. Nevertheless, test-set equilibrium still implies a central finding of the previous analysis, namely, that the payoff vector associated with any selected equilibrium is in the core of the related cooperative game. The test-set analysis improves on the original analysis by obtaining this core payoffs result without assuming an explicitly cooperative solution concept and without imposing a severe restriction on the form of the equilibrium strategies.

- Test-set equilibrium is *more* restrictive than the *locally envy-free equilibrium* proposed by Edelman, Ostrovsky and Schwarz (2007) and the similar *symmetric equilibrium* proposed by Varian (2007) for their generalized second-price auction games. In particular, test-set equilibrium implies, just as those refinements do, that losing bidders must not bid too low. When a test-set equilibrium exists, the efficiency and revenue predictions of the previous theories are thus affirmed. But the test-set refinement also requires that winning bidders must not bid too high. Depending on exogenous payoff parameters, the combination of these restrictions may be incompatible with the existence of any pure test-set equilibrium. In such cases, this analysis suggests a possible need to qualify the efficiency and revenue predictions of the previous analyses.

- Test-set equilibrium is *less* restrictive than the *tremble robust equilibrium* proposed by Abraham, Athey, Babaioff and Grubb (2016) for their second-price auction game, in which an uncertain common value is known by just one of the two bidders. The tremble robust equilibrium is the undominated Nash equilibrium that is selected by assuming that there is a small probability that a third bidder appears and bids randomly using a full support distribution. In the unique tremble robust equilibrium, the uninformed bidder bids the minimum possible value. That equilibrium is also a test-set equilibrium, but there is a second test-set equilibrium, in which the uninformed bidder bids the maximum possible value. The second equilibrium is selected if one assumes that there is a small probability that the informed bidder has a binding budget constraint.

We also show that test-set equilibrium characterizes the kinds of equilibrium outcomes that survive when each player subjects its decision to a review, of the sort that is common for decisions with large amounts are at stake. For any general finite game with normal form $\Gamma = (\mathcal{N}, \mathcal{S}, \pi)$ and any $c > 0$, we define the "high stakes version" $\bar{\Gamma}(c)$ to be the following extensive game, in which each player moves independently, with knowledge only of its own past moves. Player n's first move is to propose a strategy $s_n \in S_n$, that is, a pure strategy that it might play in the game Γ. Then, it either approves or rejects the proposed strategy. If it approves, then s_n becomes part of the outcome of the high stakes version. If it rejects, then the player's final move is to select a replacement strategy $\hat{s}_n$. Rejecting and replacing a proposed strategy incurs the cost $c > 0$, which is small when the stakes are high. The outcome of a pure strategy profile for $\bar{\Gamma}(c)$ is a profile of strategies s to be played in Γ, which leads to payoffs $\pi(s)$ minus any costs for strategy replacements. For any $c > 0$, a mixed strategy profile σ is a Nash equilibrium of Γ if and only if it is the outcome of a Nash equilibrium of $\bar{\Gamma}(c)$.

Outcomes of the high stakes versions become more interesting when we refine the equilibrium selection using a variant of the quasi-perfect equilibrium of van Damme (1984). Relative to the original concept, our variant—*quasi*-perfect equilibrium*—adds a restriction that players believe it is much less likely in any extensive form that two or more agents tremble than that just one agent trembles. When the stakes are very high, or equivalently when c is very small, the effect of this added restriction is that the trembles most likely to survive the review process in a high stakes version are ones that result in profiles in the test set. In that case, only test-set equilibria of Γ are quasi*-perfect equilibrium outcomes of $\bar{\Gamma}(c)$. Our quasi*-perfect equilibrium also relaxes the restrictions that players must have identical beliefs about the trembles and believe agents' trembles are uncorrelated, which implies that for all $c > 0$, *every* test-set equilibrium survives review. Our main theorem states that when c is sufficiently small, σ is a quasi*-perfect equilibrium outcome of $\bar{\Gamma}(c)$ *if and only if* it is a test-set equilibrium of Γ.

In section 2, we define test-set equilibrium and show that, for two-player games, any proper equilibrium is a test-set equilibrium and any test-set equilibrium is a trembling-hand perfect equilibrium, affirming that the main novelty is for games with at least three players. Section 3 applies the test-set refinement to the three auction games.

In section 4, we apply the properness and stability refinements to the auction applications. For the generalized second-price auction, we show that neither implies the locally envy-free outcome. For the (two-player) second-price auction with common values, proper equilibrium is no weaker than test-set equilibrium, but strategic stability is much weaker: it does not rule out any undominated Nash equilibrium. For the menu auction game, the two tremble-based concepts appear to be intractable.

In section 5, we introduce our theory of high stakes versions and our quasi*-perfect solution concept. We prove that when the review cost c is small, the outcomes of quasi*-perfect equilibria of $\bar{\Gamma}(c)$ are the same as the test-set equilibria of Γ.

Section 6 displays an example of a three-player game with a unique Nash equilibrium that fails the test-set condition. Our analysis of the example establishes that, for some games, Nash equilibrium implicitly requires players to hold the extreme belief that some trembles to best responses are "infinitely" less likely than others, which test-set equilibrium does not allow. Section 7 discusses our results and puts them into context.

2 Test-Set Equilibrium

We define the test set associated with a strategy profile σ, which we denote $T(\sigma)$, to consist of the strategy profiles that are derived from σ by replacing the strategy of any one player with any other best response to σ. Moreover, we say that a strategy profile σ is a test-set equilibrium if (*i*) it is a Nash equilibrium, (*ii*) no player uses a strategy that is weakly dominated in the game, and (*iii*) no player uses a strategy that is weakly dominated in $T(\sigma)$.

2.1 Notation for Games in Normal Form

A game in normal form is denoted $\Gamma = (\mathcal{N}, S, \pi)$, where $\mathcal{N} = \{1, \ldots, N\}$ is a set of players, $S = (S_n)_{n=1}^N$ is a profile of pure strategy sets, and $\pi = (\pi_n)_{n=1}^N$ is a profile of payoff functions. Such a game is *finite* if, for all players n, S_n is a finite set.

A mixed strategy profile is denoted $\sigma = (\sigma_1, \ldots, \sigma_N) \in \prod_{n=1}^N \Delta(S_n)$. We embed S_n in $\Delta(S_n)$ and extend the utility functions π_n to the domain $\prod_{n=1}^N \Delta(S_n)$ in the usual way. We use σ_{-n} for a typical element of $\prod_{m \neq n} \Delta(S_m)$, $BR_n(\sigma_{-n})$ for the set of best responses, and σ/σ_n' for the strategy profile constructed from σ by replacing player n's strategy with σ_n'.

2.2 Definition

Definition 1. Let

$$T(\sigma) = \bigcup_{n=1}^{N} \{\sigma/s_n : s_n \in BR_n(\sigma_{-n})\}.$$

A mixed strategy profile σ satisfies the *test-set condition* if and only if, for all n, there is no $\hat{\sigma}_n \in \Delta(S_n)$ such that both

(i) for all $\sigma' \in T(\sigma)$, $\pi_n(\sigma'/\hat{\sigma}_n) \geq \pi_n(\sigma'/\sigma_n)$, and

(ii) for some $\sigma' \in T(\sigma)$, $\pi_n(\sigma'/\hat{\sigma}_n) > \pi_n(\sigma'/\sigma_n)$.

We refer to $T(\sigma)$ as the *test set* associated with σ. A strategy profile σ satisfies the test-set condition if no player n is using a strategy that is weakly dominated by some $\hat{\sigma}_n$ against all strategy profiles σ'_{-n} for $\sigma' \in T(\sigma)$. When this is the case for player n, we say that its strategy is "undominated in the test set."

Definition 2. A mixed strategy profile σ is a *test-set equilibrium* if and only if it is a Nash equilibrium in undominated strategies that satisfies the test-set condition.

Our first result emphasizes that the test-set condition imposes its most novel restrictions only for games with more than two players.

Proposition 1. *In any finite, two-player game, every proper equilibrium is a test-set equilibrium, and every test-set equilibrium is a trembling-hand perfect equilibrium.*

That test-set equilibrium implies trembling-hand perfect equilibrium in games with two players follows immediately from the fact that in such games, the set of trembling-hand equilibria coincides with the set of equilibria in undominated strategies. Proper equilibrium implies more constraints on the relative probabilities of trembles by any one player than does test-set equilibrium. While test-set equilibrium implies constraints on the relative probabilities of trembles by *different* players that are absent in proper equilibrium, those constraints do not change the selection in two-player games, because neither player's best response calculation depends on them.

Our auction examples, which we analyze over the next two sections, establish that for games with at least three players, test-set equilibrium sometimes restricts outcomes in ways that proper equilibrium and strategic stability do not.

3 Auction Applications

3.1 Menu Auction

Bernheim and Whinston (1986) study the first-price menu auction and propose two refinements of its equilibria: *truthful equilibrium* and *coalition-proof equilibrium*. Payoffs of an equilibrium satisfying either condition lie on the bidder-optimal frontier of the core of the associated cooperative game.

In this section, we recapitulate the menu auction model and study properties of its test-set equilibria. We find that every truthful equilibrium is a test-set equilibrium. Similarly, every coalition-proof equilibrium payoff vector is a test-set equilibrium payoff vector. However, test-set equilibrium implies fewer restrictions, which expands the set of payoffs that can be implemented to a subset of the core containing the bidder-optimal frontier.

3.1.1 Environment

There is one auctioneer, who selects a decision x from a finite set X, which affects himself and $N \geq 2$ bidders. The gross monetary payoff that bidder n receives from any decision is described by the function $v_n : X \to \mathbb{R}$. Similarly, the auctioneer receives a gross monetary payoff described by $v_0 : X \to \mathbb{R}$. We assume that the values are normalized so that for each bidder n, $\min_{x \in X} v_n(x) = 0$, and that no two decisions generate exactly the same total surplus.

Assumption 1. $\sum_{n=0}^{N} v_n(x)$ is injective.

The N bidders simultaneously submit bids, which are offers to make payments to the auctioneer, contingent on the decision chosen. Thus, each bidder chooses a vector $b \in \mathbb{R}_+^{|X|}$, which we may also write as a function $b_n : X \to \mathbb{R}_+$. Given the bids, the auctioneer chooses a decision that maximizes his payoff $v_0(x) + \sum_{n=1}^{N} b_n(x)$. Given the bids and the decision x, an individual bidder's payoff is $v_n(x) - b_n(x)$.

To model this as a game among the bidders, we need a tie-breaking rule for the auctioneer in case two outcomes achieve the same maximal value. Generalizing the usual rule for the Bertrand model, we specify that the auctioneer breaks ties in favor of the decision with the highest total value, leading to some auctioneer decision function $x : (\mathbb{R}_+^{|X|})_{n=1}^{N} \to X$.[1] The menu auction with continuous bid spaces is then the game $\Gamma = \left[\mathcal{N}, (\mathbb{R}_+^{|X|})_{n=1}^{N}, (\pi_n(\cdot))_{n=1}^{N}\right]$, in which $\pi_n(b) = v_n(x(b)) - b_n(x(b))$.

In the Bertrand model, continuous bid spaces can be convenient for characterizing certain pure Nash equilibria, but the continuous model also differs inconveniently from most of its nearby discretized versions in two important ways. First, in the continuous Bertrand model, a bidder can place a bid that renders it indifferent between winning and losing (i.e. by bidding its value), although this is impossible in generic discretized versions. Second, in discrete Bertrand models, there are typically pure equilibria in which a losing bidder makes the highest bid less than its value, which is an undominated strategy. But the limit of such equilibria in the continuous model involves a bid equal to value, which is a dominated strategy. As an extension of the Bertrand model, the menu auction model encounters the same problems.

To conform equilibrium and dominance analyses for the continuous menu auction model with those of nearby discrete models, we fix the two problems by making two corresponding changes. Neither change affects the truthful equilibria, so both are consistent with the Bernheim-Whinston analysis. The first change is to assume that each bidder breaks indifferences among outcomes that lead to the same net payoff but different auctioneer decisions using the same criterion as the auctioneer, namely, in favor of the outcome involving the decision with the highest total value. When combined with Assumption 1, this tie-breaking assumption implies that all of a bidder's best responses to any given pure strategy profile lead to identical outcomes, just as they would in any generic discretized version of the menu auction.[2] The second change breaks another indifference: 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 (and therefore receive a payoff of zero) over any other bid vector that leads to the same auctioneer decision and the same zero payoff. That allows us to conclude below that a bidder whose bids are equal to its values has not chosen a dominated strategy. With these specifications, we obtain the following near-characterization of undominated strategies:

[1] Our approach departs slightly from the approach of Bernheim and Whinston (1986), who instead modify the definition of equilibrium to include the auctioneer's tie-breaking rule. Both approaches accomplish the same end.

[2] This is relevant for our analysis through its effect on the test-set condition.

Lemma 2. *A pure strategy b_n in the menu auction game is undominated if*

(i) $b_n(x) \le v_n(x)$ *for all* $x \in X$*, with strict inequality if* $v_n(x) > 0$*; or*

(ii) $b_n(x) = v_n(x)$ *for all* $x \in X$*.*

A pure strategy b_n in the menu auction game is dominated if $b_n(x) > v_n(x)$ *for some* $x \in X$*.*

3.1.2 Truthful Equilibrium

Bernheim and Whinston (1986) judge many of the Nash equilibria of the menu auction game to be implausible, and suggest that a bidder might limit its search for strategies to a simple, focal set of strategies, in a way that suggests a refinement:

Definition 3. A pure Nash equilibrium of the menu auction $b = (b_1, \ldots, b_N)$ is a *truthful equilibrium* if and only if for all $n \in \mathcal{N}$ and all $x \in X$, letting $x^* = x(b)$,

$$b_n(x) = \max\{0, b_n(x^*) - v_n(x^*) + v_n(x)\}.$$

In words, an equilibrium is truthful if each bidder's bid for each losing decision expresses its full net willingness to pay to switch to that decision instead (subject to nonnegativity constraint on bids). In any truthful equilibrium, all bidders use strategies that are in the class of *profit-target strategies.* In such a strategy, a bidder n sets a profit target $\bar{\pi}_n$ and bids $b_n(x) = \max(0, v_n(x) - \bar{\pi}_n)$. This bid achieves the target payoff of $\bar{\pi}_n$ whenever that is possible, and no other bid does that, so this is a potentially focal class of strategies for a bidder. An additional appeal is that this class of bids always includes a best response to any competing pure strategy profile.

Truthful equilibrium, however, is a specialized refinement. Our goal below is to show that test-set equilibrium, which is a general refinement, can do much of the same work.

3.1.3 Test-Set Equilibrium

This section has two main results. The first affirms that test-set equilibrium is not more restrictive than truthful equilibrium. The main step needed to prove that is the following lemma.

Lemma 3. *A pure Nash equilibrium of the menu auction $b = (b_1, \ldots, b_N)$ satisfies the test-set condition if and only if for all $n \in \mathcal{N}$ and all $x \in X$, letting $x^* = x(b)$,*

$$b_n(x) \ge \max\{0, b_n(x^*) - v_n(x^*) + v_n(x)\}.$$

The right-hand side of the inequality is the truthful bid for x when the winning decision is x^*. So, in words, the lemma says a Nash equilibrium with winning decision x^* satisfies the test-set condition if and only if the bids for losing decisions are at least as high as with truthful bidding. In particular, every truthful equilibrium satisfies the test-set condition. Combining the two lemmas yields the following result.

Theorem 4. *Every truthful equilibrium of the menu auction is a test-set equilibrium.*

Proof of Theorem 4. Suppose that b is a truthful equilibrium, and let $x^* = x(b)$. For all bidders n, $b_n(x^*) \le v_n(x^*)$, or else the bidder could profitably deviate to a constant bid of zero. For any bidder n, if $b_n(x^*) < v_n(x^*)$, then for all x for which $v_n(x) > 0$, we have $b_n(x) = \max\{0, b_n(x^*) - v_n(x^*) + v_n(x)\} < v_n(x)$. And for all x for which $v_n(x) = 0$, we have $b_n(x) = 0$. Thus, by Lemma 2(i) , b_n is undominated. On the other hand, if $b_n(x^*) = v_n(x^*)$, then for all x, $b_n(x) = \max\{0, b_n(x^*) - v_n(x^*) + v_n(x)\} = v_n(x)$, so Lemma 2(ii) implies that b_n is undominated. Furthermore, by Lemma 3, every truthful equilibrium satisfies the test-set condition. □

To see that the inequality in Lemma 3 is necessary, notice that if some bidder's equilibrium bid fails this condition for some decision x, then it is dominated in the test set by an alternative pure bid, which bids slightly higher for that decision. Indeed, this alternative performs no worse than the original bid against any element of the test set, and it performs strictly better in the event that another bidder also deviates by raising its bid on x by a sufficient amount. Such a deviation by another bidder can be a best response and therefore is included in the test set. For sufficiency, notice first that all best responses by any bidder lead the auctioneer to pick the equilibrium decision x^*. Therefore, when play is in the test set, a bidder can bring about a decision $x \neq x^*$ only by raising its bid for x, which from the inequality can never be profitable. Hence b_n is undominated in the test set.

The second main result demonstrates that test-set equilibrium, despite being weaker than truthful equilibrium, is nevertheless strong enough to preserve some of the same restrictions on outcomes and payoffs. To state the result, we introduce some notation. Given a set of bidders $J \subseteq \mathcal{N}$, let $\bar{J} = \mathcal{N} \setminus J$ denote its complement, and let x^J denote a decision that maximizes the payoff of the coalition consisting of J together with the auctioneer:

$$x^J \in \arg\max_{x \in X} \sum_{n \in \{0\} \cup J} v_n(x).$$

In particular, $x^{\mathcal{N}}$ is the decision that maximizes total surplus. In addition, define C to be the set of the payoffs for the bidders that are consistent with an outcome in the core.[3]

$$C = \left\{ \pi \in \mathbb{R}_+^N \;\middle|\; \begin{array}{l} \forall J \subseteq \mathcal{N}: \\ \displaystyle\sum_{n \in J} \pi_n \le \sum_{n=0}^{N} v_n(x^{\mathcal{N}}) - \sum_{n \in \{0\} \cup \bar{J}} v_n(x^{\bar{J}}) \end{array} \right\}.$$

The main results of Bernheim and Whinston (1986) are that both the truthful equilibrium payoffs and the coalition-proof equilibrium payoffs are the bidder-optimal frontier of the core:

$$E = \left\{ \pi \in \mathbb{R}^N \;\middle|\; \pi \in C \text{ and } \nexists \pi' \in C \text{ with } \pi' \geq \pi \right\}.$$

In contrast, the test-set equilibrium payoffs satisfy the related, but weaker, criterion of lying in the core. The conclusion follows from Lemma 3 using an argument similar to that in Bernheim and Whinston (1986). The intuition is as follows. Lemma 3 requires that in

[3] A core payoff includes the auctioneer's payoff. It is a vector (π_0, π) with $\pi \in C$ and $\pi_0 = \sum_{n=0}^{N} v_n(x^{\mathcal{N}}) - \sum_{n=1}^{N} \pi_n$.

any test-set equilibrium, bids for losing decisions must be sufficiently high. In particular, if the equilibrium decision were $x^* \neq x^{\mathcal{N}}$, then the lemma implies that the sum of bids for $x^{\mathcal{N}}$ is so high that the auctioneer would derive a higher payoff from choosing $x^{\mathcal{N}}$ than x^*, which is a contradiction. Likewise, if the inequality corresponding to some coalition J in the definition of C were violated, then the lemma implies that the sum of bids for $x^{\bar{J}}$ is so high that the auctioneer would derive a higher payoff from choosing $x^{\bar{J}}$ than x^*, which is a similar contradiction.

Corollary 5. *In all test-set equilibria of the menu auction, the auctioneer implements the surplus-maximizing decision $x^{\mathcal{N}}$, and the bidders receive payoffs in C.*

Among the Nash equilibria of the menu auction game are ones with inefficient outcomes ($x^* \neq x^{\mathcal{N}}$), and possibly ones that are Pareto ranked for the bidders. Both the inefficient equilibria and the Pareto inferior ones are sometimes called "coordination failures." The approach taken by Bernheim and Whinston (1986) seems to hint that eliminating either type of coordination failure somehow hinges upon either (*i*) restrictions on the strategies that can be played in equilibrium, as in truthful equilibrium, or (*ii*) cooperation in selecting bids, as in coalition-proof equilibrium. The test-set analysis highlights that this is not quite right; the individual choice criterion embodied in the test-set condition is sufficient to select an efficient outcome with core payoffs. Test-set equilibrium delivers that conclusion by implying that, for every bidder n and every decision x, $v_n(x) - b_n(x)$ is weakly greater than n's equilibrium payoff.

Test-set equilibrium does still leave open the possibility for the second kind of coordination failure, in which there is some other equilibrium that all bidders prefer. Truthful equilibrium also rules out those equilibria by requiring in its definition that if $b_n(x) > 0$, then $v_n(x) - b_n(x)$ must equal n's equilibrium payoff. This ensures that bidders do not bid too high for non-equilibrium alternatives, and so they are not forced to bid very high to make the auctioneer select $x^{\mathcal{N}}$. Likewise, coalition-proof equilibrium rules out coordination failures because the concept itself assumes that players try to coordinate. Because test-set equilibrium relies neither on exogenous restrictions on strategies nor on the assumption that bidders try to coordinate, it highlights that the auction game itself promotes efficient outcomes and payoffs in the core, but that coordinating on a core allocation that is best for bidders requires more.[4]

3.2 Generalized Second-Price Auction

Edelman, Ostrovsky and Schwarz (2007) study the generalized second-price (GSP) auction, using a Nash equilibrium refinement that they term *locally envy-free equilibrium.* Varian (2007) studies the same auction and makes the same equilibrium selection, calling these "symmetric" equilibria.

[4] A generalization of the Bernheim-Whinston model can be applied to study combinatorial auctions, in which each bidder cares about, and is restricted to bidding for, only the set of goods allocated to that bidder. Yet, the bidders may be linked by feasibility constraints on the allocation. Truthful bidding is still possible in this generalization, and one can still show that truthful equilibrium payoffs are on the bidder-optimal frontier of the core. For such games, however, a test-set equilibrium payoff may fail to lie in the core. This negative finding suggests that coordination among bidders on a core outcome may be more challenging in such an auction than in one in which bidding is not so restricted. See the online appendix for details.

In this section, we restate the GSP auction model and study the properties of its test-set equilibria. We find that every pure test-set equilibrium is locally envy-free, but that the test-set condition also implies additional restrictions. Whether these additional restrictions preclude the existence of a pure test-set equilibrium depends on the parameters of the game.

3.2.1 Environment

There are I ad positions and N bidders. The click rate of the ith position is $\kappa_i > 0$. The value per click of bidder n is $v_n > 0$. Bidder n's payoff from being in position i is $\kappa_i v_n$ minus its payments to the auctioneer.

The N bidders simultaneously submit bids. Allowable bids are the nonnegative reals, $\mathbb{R}_+$. Let $b^{(i)}$ denote the ith highest bid. It is convenient to define $b^{(N+1)} = 0$ and $\kappa_{I+1} = 0$. Bidders are then sorted in order of their bids, where ties are broken uniformly at random. After ties are broken, let, for $i \leq \min\{I, N\}$, $g(i)$ denote the identity of the ith highest bidder. Let $G(I+1)$ denote the set of all other bidders. The GSP mechanism allocates position i to bidder $g(i)$ at a per-click price of $b^{(i+1)}$, for a total payment of $\kappa_i b^{(i+1)}$. Members of $G(I+1)$ win nothing and pay nothing.

The expected payoff to bidder n under the bid profile $b = (b_1, \ldots, b_N)$ is

$$\pi_n(b) = \mathbb{E}\left[\kappa_{I_n(b)}\left(v_n - b^{(I_n(b)+1)}\right)\right],$$

where the expectation is taken over the random variable $I_n(b)$, the position won by bidder n. A GSP auction is modeled as a game $\Gamma = [\mathcal{N}, (\mathbb{R}_+)_{n=1}^N, (\pi_n(\cdot))_{n=1}^N]$.

For this analysis, we label positions and bidders so that click rates and bidder values are in descending order, from highest to lowest.

Assumption 2. We assume the following:

(i) $\kappa_1 > \cdots > \kappa_I$, and

(ii) $v_1 > \cdots > v_N$.

3.2.2 Locally Envy-Free Equilibrium

Edelman, Ostrovsky and Schwarz (2007) find that there are many implausible Nash equilibria of their GSP game, and they justify a specialized refinement in a way that is time-honored among economists: by making reference to factors outside of the game model.

Definition 4. A pure equilibrium of the GSP auction $b = (b_1, \ldots, b_N)$ is a *locally envy-free equilibrium* if for all $i \in \{2, \ldots, \min\{I+1, N\}\}$,

$$\kappa_i \left[v_{g(i)} - b^{(i+1)}\right] \geq \kappa_{i-1} \left[v_{g(i)} - b^{(i)}\right].$$[5]

[5] This definition is worded with some abuse of notation. First, the definition ignores the possibility of tied bids. However, this is not an issue, as Lemma 11 shows that in any pure equilibrium, there are no ties among the highest $\min\{I+1, N\}$ bids. Second, it would be more correct to say that in the case $i = I+1$, the inequality must hold for all $g(I+1) \in G(I+1)$.

As partial justification for limiting equilibria in this way, Edelman, Ostrovsky and Schwarz (2007) argue that the one-shot, complete information game of their model can be regarded as standing in for the limit point of an underlying, frequently-repeated game of incomplete information. In that game, one deviation that a bidder could undertake would be to raise its bid to b, thereby increasing the price paid by the bidder one position above, in the hopes of forcing that bidder out of its higher position. The higher bidder can sometimes undermine that strategy by slightly undercutting the new bid b. The bid profiles that are immune to this particular type of deviation correspond to the locally envy-free equilibria of the one-shot, complete information game.

Edelman, Ostrovsky and Schwarz (2007) also point out that if the equilibria of the GSP auction are refined in this way, then the auction possesses some attractive economic properties. In all surviving equilibria, (*i*) the equilibrium allocation and payments together constitute a stable assignment;[6] (*ii*) consequently, the equilibrium outcome is Pareto efficient and payments are competitive; and (*iii*) equilibrium revenue is at least as high as that derived from the dominant-strategy equilibrium of the corresponding VCG mechanism.

3.2.3 Test-Set Equilibrium

Our alternative analysis does not rely on an appeal to an unmodeled repeated game. Our main finding is this:

Theorem 6. *Every pure test-set equilibrium of the GSP auction is a locally envy-free equilibrium.*

Our proof of Theorem 6 (in an appendix) shows the contrapositive. The thrust of the proof relies on the following argument. If the winner $g(i)$ of position i envies $g(i-1)$, who is the winner of position $i-1$, then $g(i)$'s equilibrium bid is weakly dominated in the test set by a slightly higher bid. The reason is that the slightly higher bid leaves $g(i)$'s allocation and price unchanged against all elements of the test set except those in which bidder $g(i-1)$ deviates by reducing its bid to fall between $g(i)$'s equilibrium bid and the alternative bid. For those test elements, bidder $g(i)$'s higher bid does strictly better: it causes the allocation to reverse, so that $g(i)$ wins the higher position at approximately its equilibrium bid. So, if a strategy profile is not locally envy-free, then it it is not a test-set equilibrium.

3.2.4 Locally Envy-Free Equilibria that are not Test-Set Equilibria

While every pure test-set equilibrium is a locally envy-free equilibrium, the converse is not true. As a simple illustration of this possibility, consider an example in which there are two bidders and a single ad position. In the single-item case, the GSP auction reduces to a second-price auction, and so it possesses a dominant strategy solution, which, by definition, is also the unique test-set equilibrium. On the other hand, a bid profile is locally envy-free in this example if bidder 1 wins and if bidder 2 bids at least its own value yet no more than bidder 1's value. This includes profiles in which both bidders play dominated strategies.

[6] In this setting, an assignment is stable if for every pair of positions i and j, $\kappa_i v_{g(i)} - p^{(i)} \geq \kappa_j v_{g(i)} - p^{(j)}$, where $p^{(i)}$ is the price paid by the winner of position i. This result is proven in Lemma 1 of Edelman, Ostrovsky and Schwarz (2007).

Another way to compare locally envy-free equilibrium and test-set equilibrium begins by replacing the former with the following equivalent definition: a Nash equilibrium is locally envy-free if each bidder's bid is undominated against an alternative set of profiles that includes the equilibrium profile and any profile in which the bidder whose bid is one position higher deviates to a different best response. As this restatement reveals, locally envy-free equilibrium implicitly assumes that each bidder regards trembles by higher bidders as relatively more likely than trembles by lower bidders. In contrast, the test-set condition treats trembles by higher bidders and by lower bidders as equally plausible: it tests each bid against a set of profiles in which *any* other bidder may deviate to a different best response. This characterization highlights that locally envy-free equilibrium differs in two ways from test-set equilibrium. First, it does not require that the Nash equilibrium bids are undominated in the usual sense, that is, against all possible profiles of bids. Second, the set of deviations against which the equilibrium profile is tested is different. Both of these differences can matter. The example in the previous paragraph shows that a locally envy-free equilibrium can use a dominated strategy, and so can fail to be a test-set equilibrium. The next subsection highlights the second difference. Its result implies that even when a locally envy-free equilibrium uses undominated strategies, it can still fail the test-set condition.

3.2.5 Potential Nonexistence of Pure Test-Set Equilibria

In a GSP game, pure test-set equilibria can fail to exist, depending on the parameters.

Proposition 7. *Let $I = N = 3$. There exists a pure test-set equilibrium of the GSP auction if and only if*

$$\frac{v_3}{v_2} \leq \frac{\kappa_2^2 - \kappa_1\kappa_3}{\kappa_2^2 - \kappa_2\kappa_3}.$$

We sketch the proof, alongside an intuitive explanation. Suppose that $b = (b_1, b_2, b_3)$ is a test-set equilibrium. Because pure test-set equilibria are locally envy-free, the outcome must be an assortative matching, which requires $b_1 > b_2 > b_3$. Test-set equilibrium implies two additional restrictions. It requires that b_2 must exceed some threshold b_2^{low}, for otherwise it is weakly dominated in the test set by $b_2 + \varepsilon$ for some small ε. Intuitively, this higher bid is better when bidder 1 deviates downwards, and otherwise is no worse. But test-set equilibrium also requires that b_2 must lie below some threshold b_2^{high}, for otherwise it is weakly dominated in the test set by $b_2 - \varepsilon$ for some small ε. Intuitively, this lower bid is better when bidder 3 deviates upwards, and otherwise is no worse. A pure test-set equilibrium exists if and only if $b_2^{high} \geq b_2^{low}$. The proposition restates that inequality in terms of the exogenous parameters.

In summary, test-set equilibrium generates a partitioning of the parameter space that was not revealed by previous analysis. In one part of the parameter space, a pure test-set equilibrium exists, which, by Theorem 6, is locally envy-free. In such cases, our analysis supports the reasonableness of the locally envy-free selection, as well as its associated predictions about efficiency and revenue in the GSP auction. But in another part of the parameter space, no pure test-set equilibrium exists. In such cases, it is less clear that bidding behavior should be expected to converge at all, which would suggest a potential need to qualify the predictions of locally envy-free equilibrium.

3.3 Second-Price, Common Value Auction

Abraham, Athey, Babaioff and Grubb (2016) study a second-price auction game with common values and incomplete information, motivated by certain auctions for internet display advertising. They find that there are a continuum of Nash equilibrium outcomes and suggest narrowing the possible outcomes by selecting a single *tremble robust equilibrium*.

They begin their analysis with a two-bidder example. One bidder is informed, receiving a private signal that is either low or high, and the other uninformed. The paper also treats richer settings, but it suffices for our purposes to consider this simple example. We show that test-set equilibrium eliminates all but two pure equilibria, one of which is the unique tremble robust equilibrium.

3.3.1 Environment

Two bidders participate in a second-price auction. The object being auctioned has a common value of v to both bidders, which is either 0 or 1, each with equal probability. One bidder is *informed* and learns the value of v. The other bidder is *uninformed*. Allowable bids are the nonnegative reals, $\mathbb{R}_+$.

This game possesses many equilibria, yet standard refinements do little to focus the set of predictions. Standard refinements do make a focused prediction for the informed bidder: that it will use the dominant strategy of always bidding the value v. For the uninformed bidder, bids outside the unit interval are weakly dominated, and standard refinements rule those out. However, most refinements do little more to discipline the uninformed bidder's strategy.

Although the original definition of trembling-hand perfect equilibrium (Selten, 1975) applies only to finite games, Simon and Stinchcombe (1995) have proposed some ways to extend the application to games with infinite strategy sets. Abraham, Athey, Babaioff and Grubb (2016) find that those extensions are not very restrictive in this application: in equilibrium, the uninformed bidder can make any bid in the interval $[0, 1]$. Intuitively, in this two-player game, perfection adds only the requirement that each bidder plays an undominated strategy. In a later section, we will show that, for finite approximations of this game, strategic stability is similarly unrestrictive.

3.3.2 Tremble Robust Equilibrium

Finding some of these these equilibria to be implausible, Abraham, Athey, Babaioff and Grubb (2016) propose a refinement, "tremble robust equilibrium," justifying it by reference to the economic context of the game. Informally, an equilibrium is tremble robust if it is the limit of equilibria of a sequence of perturbations of the game in which, with vanishingly small probability, an additional bidder submits a randomly chosen bid.

Definition 5. A Nash equilibrium σ is *tremble robust* if there exists a distribution F with continuous, strictly positive density on $[0, 1]$, a sequence of positive numbers $\{\varepsilon_j\}_{j=1}^{\infty}$ converging to zero, and a sequence of strategy profiles $\{\sigma_j\}_{j=1}^{\infty}$ converging in distribution to σ such that for all j:

(i) σ_j is a Nash equilibrium of the perturbation of the game in which with probability ε_j an additional bidder arrives and submits a bid sampled from F, and

(ii) σ_j does not prescribe dominated bids.

Just as for the previous examples, the proposed equilibrium refinement for this game is specialized.

In this game, a pure strategy profile is a triple (b_0, b_1, b_U) giving the bids by the low- and high-types of the informed bidder and by the uninformed bidder. Abraham, Athey, Babaioff and Grubb (2016) show that this game has a unique tremble robust equilibrium: $(0, 1, 0)$. In contrast to trembling-hand perfect equilibrium, tremble robust equilibrium requires the uninformed bidder to place a bid of 0. This leads to another observation made by the authors: expected revenue in the unique tremble robust equilibrium of the second-price auction (which is zero) is strictly lower than that of any equilibrium of the first-price auction.

3.3.3 Test-Set Equilibrium

For this game, we have seen previously that there is a continuum of pure trembling-hand perfect equilibria. Test-set equilibrium is more restrictive.

Proposition 8. *There are two pure test-set equilibria of this game:* $(0, 1, 0)$ *and* $(0, 1, 1)$.

Given any other undominated Nash equilibrium $(0, 1, b_U)$ with $0 < b_U < 1$, the best responses for the informed bidder have the low type make any bid in $[0, b_U)$ and the high type any bid in $(b_U, 1]$. When the uninformed bidder faces this test set, it finds that any alternative bid $\hat{b}_U \neq b_U$ in $[0, 1]$ weakly dominates b_U, thus eliminating that equilibrium. So, the two profiles identified by the proposition are the only remaining candidates, and it is easy to verify that they satisfy the test-set condition.[7]

Thus, test-set equilibrium makes a more focused set of predictions than perfect equilibrium, isolating just two candidate pure equilibria. The equilibrium with $b_U = 0$ is the one that Abraham, Athey, Babaioff and Grubb (2016) had selected by introducing a perturbed game, with a third bidder who appears with low probability and places a bid at random. The third bidder is irrelevant when the informed bidder's type is high, so from the uninformed bidder's perspective, the perturbation is equivalent to assuming that only the low-type informed bidder trembles. The second test-set equilibrium could similarly be selected by a perturbation in which there is a small probability that the informed bidder's budget constraint forces it to bid less than 1, which is equivalent to assuming that only the high-type trembles. In contrast to the tremble robust equilibrium, the auctioneer's revenue in this second test-set equilibrium is strictly higher than in the corresponding first-price auction.

4 Alternative Refinements for the Applications

In this section, we sketch applications of proper equilibrium and strategic stability to a generalized second-price auction game and to the second-price, common value auction considered above. We limit our discussion to those two applications, because the multi-dimensional

[7]The conclusion of Proposition 8 also holds when the game is analyzed in agent-normal form, and the analysis is nearly the same as well.

bid spaces of the menu auction mechanism make it difficult to apply tremble-based concepts like properness and stability. Since these tremble-based refinements are usually defined only for finite games, the following analysis restricts bids to the finite set $\{0, \frac{1}{m}, ..., \frac{mM}{m}\}$ for some positive integers m and M.

4.1 Generalized Second-Price Auction

Consider an example of the generalized second-price auction game with three bidders and three ad positions. The three bidders' values per click are 3, 2 and 1, respectively. The top position attracts 4 clicks; the second position attracts 2 clicks, and the bottom position attracts one click. For the discretized bid set, let m be any integer multiple of 2, and take $M \geq 2$.

We begin by searching for what will turn out to be the unique Nash equilibrium of the GSP auction in which (*i*) the outcome is efficient, (*ii*) bidder 1 is much less likely to tremble than bidder 2 (that is, bidder 1's total probability of trembling is much smaller than bidder 2's probability of trembling to any particular bid), and (*iii*) bidder 2 is much less likely to tremble than bidder 3. To derive the equilibrium bids for bidders 2 and 3, let us temporarily suppose that bidder 1 bids high and that its trembles have probability 0. That results in a second-price auction game between bidder 2 and 3, in which the higher bidder wins the second position and the lower bidder wins the bottom position and pays 0. This game has dominant strategies for its two players $b_3 = \frac{1}{2}$ and $b_2 = 1$. Next, to infer the equilibrium strategy for bidder 1, fix the bid $b_2 = 1$ of bidder 2 and suppose that bidder 2's trembles have probability zero. In the induced game between bidder 1 and 3, bidder 1's dominant strategy is to bid $b_1 = 2$. It is routine to verify that if the zero probabilities used for this intuitive derivation are perturbed to create full support distributions of trembles, then what had been dominant strategy bids become a strict equilibrium (in which each bidder is playing its unique best response). Since proper equilibrium imposes no restrictions on the relative probabilities of trembles by different players, $b^* = (2, 1, \frac{1}{2})$ is a proper equilibrium. Moreover, the uniqueness of equilibrium with these trembles implies that any stable set that includes an efficient equilibrium must include b^*.

Now, we show that b^*, which is both proper and part of any stable set that includes an efficient equilibrium, is not locally envy-free. Indeed, in any locally envy-free equilibrium b, bidder 2 prefers its equilibrium allocation, with its payoff of $2(2 - b_3)$, to the allocation of bidder 1, with its payoff of $4(2 - b_2)$. That requires $b_2 \geq 1 + \frac{b_3}{2}$, which does not hold in the case of b^*. On the other hand, it can be shown that b^* is not a test-set equilibrium even in this discretized setting, provided that the discretization is sufficiently fine ($m \geq 6$).

4.2 Second-Price, Common-Value Auction

As before, there are two bidders, one informed and one uninformed, and the value of the object to either bidder is either 0 or 1. For the discretized bid set, let $m \geq 3$, which ensures that the discretization is sufficiently fine.

As in the continuous model, the undominated Nash equilibria in pure strategies are those in which (*i*) the informed bidder plays its dominant strategy, bidding 0 when the common value is 0 and 1 when it is 1, and (*ii*) the uninformed bidder bids any $b_U = \frac{k}{m} \in [0, 1]$. Both properness and stability imply these restrictions as well.

Stability, however, implies no further restrictions: the unique stable set consists of all the undominated Nash equilibria. To verify that observation, consider a perturbation of the game according to which the uninformed bidder trembles with full support and, for $0 < k < m$, the informed bidder trembles with positive positive probability only to the following strategy: bid $\frac{k+1}{m}$ when the common value is 0 and $\frac{k-1}{m}$ when it is 1. When the informed bidder does not tremble, it must play its dominant strategy, since that is its unique best response in this perturbation. Given that, the uninformed bidder's unique best response in this perturbation is to bid $b_U = \frac{k}{m}$. Therefore, the Nash equilibrium profile $(0, 1, \frac{k}{m})$ is part of any stable set when $0 < k < m$. A similar argument applies when $k = 0$ or m. Moreover, it is again routine to verify that the zero probabilities used in this intuitive derivation can be replaced by very low positive probabilities and full support distributions for the trembles of both bidders.

By arguments similar to those made in the continuous case, there are again two pure test-set equilibria of this auction: $(0, 1, 0)$ and $(0, 1, 1)$.[8] Since this is a two-player game, by Proposition 1, these are the only pure strategy profiles that are candidates to be proper equilibria. In contrast to test-set equilibrium, proper equilibrium imposes some hard-to-analyze restrictions on the relative probabilities of trembles by the two types of the informed bidder. We have been unable to ascertain whether both candidates are proper equilibria. The proper equilibria of the agent-normal form, which impose no such restrictions, coincide exactly with the test-set equilibria. See the online appendix for details.

5 Quasi*-Perfect Equilibrium and High Stakes Games

5.1 Notation for Games in Extensive Form

The solution concept developed in this section is a variant of the quasi-perfect equilibrium of van Damme (1984). To highlight the similarities, the text in this subsection is copied almost verbatim from that paper.

We introduce the following notation. Let $\mathcal{N} = \{1, \ldots, N\}$ denote the set of players. Let $\bar{\Gamma}$ be a finite extensive form game with perfect recall. Let u be an information set of player n in $\bar{\Gamma}$. A local strategy b_{nu} at u is a probability distribution on C_u, where C_u denotes the set of choices at u. The probability that b_{nu} assigns to $c \in C_u$ is denoted by $b_{nu}(c)$ and B_{nu} denotes the set of all local strategies at u.[9] We view C_u as a subset of B_{nu}.

A *behavior strategy* b_n of player n is a mapping that assigns to every information set of player n a local strategy. U_n denotes the set of all information sets of player n and B_n is the set of all behavior strategies of this player. A behavior strategy b_n is *completely mixed* if $b_{nu}(c) > 0$ for all $u \in U_n$ and $c \in C_u$. If $b_n \in B_n$ and $b'_{nu} \in B_{nu}$, then b_n/b'_{nu} is used to denote the behavior strategy which results from b_n if b_{nu} is changed to b'_{nu}, whereas all other local strategies remain unchanged.

[8] The argument for the discrete case is as follows. Suppose that $b_U \in \{\frac{1}{m}, \ldots, \frac{m-1}{m}\}$. To show that $(0, 1, b_U)$ is not a test-set equilibrium, we observe the following. First, a bid of 1 weakly dominates b_U in the test set, unless $b_U = \frac{m-1}{m}$. Second, a bid of 0 weakly dominates b_U in the test set, unless $b_U = \frac{1}{m}$. Since we assume $m \geq 3$, at least one of these two cases applies.

[9] In sections 5.1 and 5.2 only, to conform to van Damme's original formulation, we use c to denote a choice at an information set. In other sections, c denotes the cost of revision in a high stakes version $\bar{\Gamma}(c)$.

For an information set u, we write $Z(u)$ for the set of all endpoints of the tree coming after u and if $u, v \in U_n$, then we write $u \leq v$ if $Z(u) \supset Z(v)$. As usual, $u < v$ stands for $u \leq v$ and $u \neq v$ (hence, $u < v$ means v comes after u). Note that the relation $\leq$ partially orders U_n, since the game has perfect recall. If $b_n, b'_n \in B_n$ and $u \in U_n$, then we use $b_n/_u b'_n$ to denote the behavior strategy b''_n defined by

$$b''_{nv} = \begin{cases} b'_{nv} & \text{if } v \geq u, \\ b_{nv} & \text{otherwise} \end{cases}$$

Furthermore, letting $B = \prod_{n=1}^{N} B_n$, if $b \in B$ and $b'_n \in B_n$, then $b/_u b'_n$ denotes the behavior strategy profile $b/(b_n/_u b'_n)$.

We denote player n's expected payoff given that the behavior strategy profile b is played by $\bar{\pi}_n(b)$. We also denote player n's expected payoff given that information set u is reached and the behavior strategy profile b is played by $\bar{\pi}_{nu}(b)$.

Definition 6 (van Damme, 1984)**.** Let $\bar{\Gamma}$ be a finite game in extensive form with perfect recall. A behavior strategy profile b is a *quasi-perfect equilibrium* of $\bar{\Gamma}$ if there exists a sequence $\{b^t\}_{t=1}^{\infty}$ of completely mixed behavior strategy profiles, converging to b, such that for all n, u, and t,

$$\bar{\pi}_{nu}(b^t/_u b_n) = \max_{b'_n \in B_n} \bar{\pi}_{nu}(b^t/_u b'_n).$$

Thus, quasi-perfect equilibrium requires that at each information set, players take into account past mistakes, as well as potential future mistakes by opponents. However, in contrast to trembling-hand perfect equilibrium of the extensive form $\bar{\Gamma}$, each player assumes that it will make no mistakes at future information sets, even if it had made a mistake in the past.

5.2 Quasi*-Perfect Equilibrium

Our proposed refinement differs from quasi-perfect equilibrium in three ways. First, it allows players to hold differing beliefs about the trembles.[10] Second, it allows players to believe that other players' trembles may be correlated.[11] Third, it adds the restriction that, when taking limits, the "independent components" of the trembles at each information set all converge to zero at the same rate and that the "correlated components" of the trembles converge to zero at a faster rate than that.[12]

[10] Other refinements that relax interpersonal consistency of beliefs include that of Weibull (1992), as well as c-perfect equilibrium and other related refinements in Fudenberg, Kreps and Levine (1988).

[11] Correlated trembles have also been considered elsewhere in the literature. For example, they are the basis of c-perfect equilibrium and other related refinements in Fudenberg, Kreps and Levine (1988). In contrast to their approach, ours allows only a small amount of correlation, requiring the probability of correlated trembles to converge to zero at a faster rate than the probability of any single tremble.

[12] Bagwell and Ramey (1991) employ a related restriction to select equilibrium in a limit-pricing game, in which a potential entrant, who does not know the cost level of the industry, observes the prices set by two informed duopolists. Their refinement specifies that if the entrant finds himself at an unexpected information set that could be explained either by a deviation by a single duopolist (when costs are at one level) or by deviations by both duopolists (when costs are at another level), then he should presume the former.

To accommodate these adjustments, the first step is to extend the payoff functions in the obvious way to allow correlated distributions. A correlated behavior strategy profile is a general distribution d on $\prod_{n=1}^{N} \prod_{u \in U_n} C_u$. Player n's associated payoff, either unconditionally or conditional on reaching an information set u, is denoted by $\bar{\pi}_n(d)$ or $\bar{\pi}_{nu}(d)$, respectively. Additionally, if b'_n is a behavior strategy for player n and if $u \in U_n$, then let $d/_u b'_n$ denote the product of the following two distributions: first, the distribution on $\prod_{v \in U_n : v \geq u} C_v$ that is induced by b'_n; second, the marginal of d on $\prod_{v \in U_n : v \ngeq u} C_v \times \prod_{m \neq n} \prod_{v \in U_m} C_v$. This represents what n believes if it deviates at some information set u and the information sets thereafter, expecting that its deviation will not affect any other player's strategy. These definitions agree with the previous ones in the case that d is the product distribution induced by a behavior strategy profile b.

Using this notation, we describe our proposed modification of quasi-perfect equilibrium.

Definition 7. Let $\bar{\Gamma}$ be a finite game in extensive form with perfect recall. A behavior strategy profile b is a *quasi*-perfect equilibrium* of $\bar{\Gamma}$ if there exist a profile of completely mixed behavior strategy profiles, $(\tau^n)_{n=1}^{N}$; a profile of sequences of distributions on $\prod_{n=1}^{N} \prod_{u \in U_n} C_u$, $(\{d^{t,n}\}_{t=1}^{\infty})_{n=1}^{N}$; as well as sequences of positive real numbers $\{\varepsilon_t\}_{t=1}^{\infty}$ and $\{\delta_t\}_{t=1}^{\infty}$ such that

(i) $\lim_{t \to \infty} \varepsilon_t = 0$ and $\lim_{t \to \infty} \delta_t = 0$,

(ii) for all m, all t, and all $c \in \prod_{n=1}^{N} \prod_{u \in U_n} C_u$,

$$d^{t,m}(c) \geq (1 - \varepsilon_t \delta_t) \prod_{n=1}^{N} \prod_{u \in U_n} \left[(1 - \varepsilon_t) b_{nu}(c_u) + \varepsilon_t \tau^m_{nu}(c_u)\right],$$

(iii) and for all n, all $u \in U_n$, and all t,

$$\bar{\pi}_{nu}(d^{t,n}/_u b_n) = \max_{b'_n \in B_n} \bar{\pi}_{nu}(d^{t,n}/_u b'_n).$$

For each player m, the sequence $\{d^{t,m}\}_{t=1}^{\infty}$ in the definition represents that player's beliefs about trembles. Condition (ii) restricts these beliefs in the following way. With probability $1 - \varepsilon_t \delta_t$, the choices at each information set are made independently. In that case, each player m believes that for information set $u \in U_n$, play follows the equilibrium local strategy b_{nu} with probability $1 - \varepsilon_t$ and follows the independent tremble τ^m_{nu} with probability ε_t. With the remaining $\varepsilon_t \delta_t$ probability, correlated trembles may occur.

This definition incorporates the three properties we have described. First, it permits players to possess different beliefs about trembles. Second, it permits players to believe that trembles are correlated, but that the correlated trembles are relatively rare (happening with total probability $\varepsilon_t \delta_t$) compared to the chance of any single tremble (which has a probability on the order of ε_t). Third, it requires that the probability of every single independent tremble vanishes at the same rate (on the order of ε_t).

5.3 High Stakes Versions

We now introduce a model in which the players, regarding their strategy choices as involving high stakes, subject those choices to review. This might happen, for example, in an auction, if a salesman or a consultant recommends a bid and presents it for review and approval by a senior manager. Alternatively, the agents of any player in the model might all be the same individual, who makes an initial decision and then "sleeps on" it before confirming or changing it.

Given any finite game in normal form, Γ, a *high stakes version* of that game, $\bar{\Gamma}(c)$, is a related extensive form game in which each of the N players selects a strategy to be played in Γ via a three step process with perfect recall. First, each player chooses ("recommends") a pure strategy $s_n \in S_n$. Second, the player observes ("reviews") s_n and makes a binary choice in $\{Approve, Disapprove\}$. In the case of approval, s_n is passed on to be played in Γ. In the case of disapproval, the player chooses ("finally decides") a pure strategy $s'_n \in S_n$, which is then played in Γ. These processes are independent: each player in this extensive game observes only its own previous choices and not those of any other player.

Given a behavior strategy for player n in $\bar{\Gamma}(c)$, b_n, the aforementioned process determines the mixed strategy that will be played in Γ, which we denote $\alpha_n(b_n) \in \Delta(S_n)$. We also use $\beta_n(b_n)$ to denote the probability with which player n's recommendation is disapproved on path when the player uses behavior strategy b_n. Furthermore, given a behavior strategy profile b, we use $\alpha(b)$ to denote the strategy profile $(\alpha_n(b_n))_{n=1}^N$, which we refer to as the "outcome of b."

To complete the description of $\bar{\Gamma}(c)$, we specify a profile of payoff functions $\bar{\pi} = (\bar{\pi}_n)_{n=1}^N$, as follows. Given a behavior strategy profile b, player n receives the payoff $\bar{\pi}_n(b) = \pi_n(\alpha(b)) - \beta_n(b_n)c$. Thus, c is to be understood as the cost that a player incurs by rejecting and replacing its initial choice (or "recommendation").

5.4 Representation Theorem

The next results connect Γ and $\bar{\Gamma}(c)$. For any review cost c, the Nash equilibrium outcomes of $\bar{\Gamma}(c)$ are the Nash equilibria of the game Γ. We interpret this to mean that if the players make no mistakes, then the review process does not change anything. However, for all sufficiently small values of c, the quasi*-perfect equilibrium outcomes of $\bar{\Gamma}(c)$ are the test-set equilibria of the original game Γ. This means that adding the possibility of mistakes in the review process as specified by quasi*-perfect equilibrium leads to test-set equilibrium.[13]

Proposition 9. *Let Γ be a game in normal form and $c \geq 0$. A strategy profile σ is a Nash equilibrium of Γ if and only if there exists a Nash equilibrium b of $\bar{\Gamma}(c)$ such that $\sigma = \alpha(b)$.*

[13] Thus, this framework serves as a behavioral foundation for test-set equilibrium. In a similar spirit, other authors have provided behavioral foundations for proper equilibrium. Myerson and Weibull (2015) consider a model in which each player, rather than having its full strategy sets at its disposal, chooses its strategy from a possibly restricted "consideration set" given by nature. Under certain conditions, the projections of the Nash equilibria of these consideration set games converge to the proper equilibria of the original game. Kleppe, Borm and Hendrickx (2017) consider a model in which each player chooses not only a single strategy but also a list of backups to be attempted in the event that nature "blocks" its initial choice. Under certain conditions, the projections of the Nash equilibria of these fall back games converge to a subset of the proper equilibria of the original game.

Theorem 10. *Let Γ be a finite game in normal form. A strategy profile σ is a test-set equilibrium of Γ if and only if there exists $\bar{c} > 0$ such that for all $c \in (0, \bar{c})$, there exists a quasi*-perfect equilibrium b of $\bar{\Gamma}(c)$ such that $\sigma = \alpha(b)$.*

In one direction, the proof of Theorem 10 (given in an appendix) works by demonstrating that it takes two trembles in $\bar{\Gamma}(c)$ to bring about an outcome in Γ that is outside the test set.[14] In contrast, a single tremble in the initial choice (by the "recommender agent") can survive if it leads to a different best response, because no player finds it worthwhile to incur the cost of c to correct such a tremble. Multiple trembles are infinitely less likely, so only test-set equilibria can survive the review process. That all test-set equilibria survive review is proven by applying the supporting hyperplane theorem to construct each player's beliefs about the independent and correlated components of trembles.

This discussion highlights the reasons we need to modify the quasi-perfect equilibrium concept for the analysis. First, like perfect and proper equilibrium, quasi-perfect equilibrium does not require or imply that, when trembles become rare, single trembles must be much more likely to combinations of trembles. It is by adding this requirement to the definition of quasi*-perfect equilibrium that we obtain necessity of test-set equilibrium. Second, like perfect and proper equilibrium, quasi-perfect equilibrium entails the further restrictions that players agree about the probabilities of trembles and regard trembles as statistically independent. It is by removing those elements from the definition of quasi*-perfect equilibrium that we obtain sufficiency of test-set equilibrium.

6 Test-Set Equilibrium May Not Exist

Figure 1 presents an example of a game that has no test-set equilibrium.[15] (*Up*, *Left*, *West*) is the unique Nash equilibrium of this game, but it is not a test-set equilibrium because *East* weakly dominates *West* in the test set:

$$T(\mathit{Up}, \mathit{Left}, \mathit{West}) = \left\{ \begin{array}{c} (\mathit{Up}, \mathit{Left}, \mathit{West}), (\mathit{Down}, \mathit{Left}, \mathit{West}), (\mathit{Up}, \mathit{Center}, \mathit{West}), \\ (\mathit{Up}, \mathit{Right}, \mathit{West}), (\mathit{Up}, \mathit{Left}, \mathit{East}) \end{array} \right\}.$$

The unique Nash equilibrium of this game (which is also a perfect, proper, and strategically stable equilibrium) is (*Up*, *Left*, *West*). Geo's choice of *West* can be optimal only if Geo believes that a joint tremble to (*Down*, *Right*) is at least as likely as the single tremble to *Center*. That requires the extreme belief that a single tremble to *Center* is "infinitely less likely" than either a single tremble to *Down* or a single tremble to *Right*, even though all

[14] Suppose that c is sufficiently small, and b is a quasi*-perfect equilibrium of $\bar{\Gamma}(c)$ leading to the outcome $\sigma = \alpha(b)$. The proof establishes that for each player n, the behavior strategy b_n must specify: (*i*) playing σ_n at the recommendation information set, (*ii*) approving best responses to σ_{-n} at review information sets while disapproving inferior responses, and (*iii*) playing best responses to σ_{-n} at final decision information sets. As in Proposition 9, σ_n must itself be a best response to σ_{-n}. Given this, a player must tremble twice in $\bar{\Gamma}(c)$ in order to play an inferior response to σ_{-n} in Γ: either at the recommendation and review stages or at the review and final decision stages.

[15] We thank Michael Ostrovsky, Markus Baldauf, and Bernhard von Stengel for helpful suggestions that led us to find this example.

Figure 1: A three player game†

[*West*]	*Left*	*Center*	*Right*
Up	$0,0,0$	$0,0,0$	$0,0,0$
Down	$0,0,0$	$-1,1,0$	$1,-1,0$

[*East*]	*Left*	*Center*	*Right*
Up	$0,0,0$	$0,-1,1$	$0,1,0$
Down	$-1,0,0$	$-1,-1,0$	$1,-1,-1$

†Row's payoffs are listed first. Column's payoffs are listed second. Geo's payoffs are listed third.

three strategies are best responses to equilibrium play. Test-set equilibrium prohibits Geo from holding such extreme beliefs, and therefore does not exist for this game.[16]

Such an extreme belief is also inconsistent with the logic of quasi*-perfect equilibrium in the high stakes version, which specifies that the probabilities of all single trembles must be of the same order of magnitude. In particular, for all $c > 0$, just a single tremble in $\bar{\Gamma}(c)$ (by Column, at the recommendation stage) is required to bring about a single deviation to *Center* in Γ. However, trembles by two different players are required in $\bar{\Gamma}(c)$ to bring about a joint deviation to (*Down*, *Right*) in Γ. Quasi*-perfect equilibrium prevents Geo from believing that the second possibility is as likely as the first.

7 Discussion

Despite doubts among game theorists that Nash equilibrium conditions are either necessary or sufficient for a solution to general non-cooperative games, equilibrium analysis continues to be a standard part of both theoretical and empirical studies of competitive bidding. There is now a huge empirical literature about auction games, with its hypotheses often motivated by theoretical properties of Nash equilibrium. In theoretical work on auctions, many analyses, including the ones reviewed in this paper, have not relied on Nash equilibrium alone. Instead, they have employed specialized equilibrium refinements that do not apply to general finite games.

In this paper, we propose a new refinement of Nash equilibrium for general finite games that may be particularly valuable for limiting multiplicity of equilibria in auctions and other games with high stakes. In offering this refinement, we do not argue that equilibrium approaches are always convincing: our proposed refinement is not a substitute for careful thought and consideration of evidence about the economic setting. What we have shown is that applications of the test-set refinement can mimic, illuminate, and sometimes improve previous analyses that had relied on idiosyncratic criteria.

For games with three or more players, test-set equilibrium introduces a new restriction on player's beliefs that is not found in older tremble-based refinements. The restriction is

[16] Milgrom and Mollner (2017) introduce the related refinement of *extended proper equilibrium* and show that it exists for all finite games and refines proper equilibrium. Like test-set equilibrium, it implies the restriction that any single tremble to a strategy that is a best response to equilibrium is more likely than one to a strategy that is an inferior response to the equilibrium, whether by the same player (as in proper equilibrium) or by another player. Unlike test-set equilibrium, however, extended proper equilibrium does not preclude the extreme belief that a single-player deviation to a strategy that is a best response to equilibrium play may be infinitely less likely than another such strategy.

that the likeliest strategy profiles are ones in which (i) at most one player deviates from equilibrium, where (ii) that deviation is to an alternative best response to equilibrium play. In a similar way, the related extensive-form solution concept that we introduce, quasi*-perfect equilibrium, requires the first of these conditions, although not the second. However, the second condition is satisfied by quasi*-perfect equilibrium *outcomes* of the high stakes version of any finite game (in which players subject their initial strategy choices to a potentially costly review process) when the stakes are sufficiently high.

We have shown that proper equilibrium and strategic stability are too weak to yield the familiar conclusions of two well-known auction analyses, and they are too difficult to apply to a third. In contrast, test-set equilibrium is both easy to apply and effective in all three applications. Moreover, our derivation of test-set equilibria from quasi*-perfect equilibria of the high stakes versions provides an account of how the sort of review processes in common use could promote this selection when the stakes are sufficiently high.

References

Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael Grubb, "Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information," *Working Paper*, 2016.

Arrow, K.J., E.W. Barankin, and D. Blackwell, "Admissible Points of Convex Sets," in H. W. Kuhn and A. W. Tucker, eds., *Contributions to the Theory of Games*, Vol. 2 1953, pp. 87–91.

Bagwell, Kyle and Garey Ramey, "Oligopoly Limit Pricing," *The Rand Journal of Economics*, 1991, *22* (2), 155–172.

Bernheim, B. Douglas and Michael D. Whinston, "Menu Auctions, Resource Allocation, and Economic Influence," *The Quarterly Journal of Economics*, 1986, *101* (1), 1–32.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz, "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 2007, *97* (1), 242–259.

Fudenberg, Drew, David M. Kreps, and David K. Levine, "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 1988, *44* (2), 354–380.

Kleppe, John, Peter Borm, and Ruud Hendrickx, "Fall Back Proper Equilibrium," *TOP*, 2017, *25* (2), 402–412.

Kohlberg, Elon and Jean-Francois Mertens, "On the Strategic Stability of Equilibria," *Econometrica*, 1986, *54* (5), 1003–1037.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green, *Microeconomic Theory*, Oxford University Press, New York, 1995.

Milgrom, Paul and Joshua Mollner, "Extended Proper Equilibrium," *Working Paper*, 2017. http://ssrn.com/abstract=3035565.

Myerson, Roger, "Refinements of the Nash Equilibrium Concept," *International Journal of Game Theory*, 1978, *7* (2), 73–80.

— **and Jörgen Weibull**, "Tenable Strategy Blocks and Settled Equilibria," *Econometrica*, 2015, *83* (3), 943–976.

Selten, Reinhard, "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 1975, *4* (1), 25–55.

Simon, Leo K. and Maxwell B. Stinchcombe, "Equilibrium Refinement for Infinite Normal-Form Games," *Econometrica*, 1995, *63* (6), 1421–1443.

van Damme, Eric, "A Relation Between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games," *International Journal of Game Theory*, 1984, *13* (1), 1–13.

Varian, Hal R., "Position Auctions," *International Journal of Industrial Organization*, 2007, *25* (6), 1163–1178.

Vrahatis, Michael N., "Generalization of the Bolzano Theorem for Simplices," *Topology and its Applications*, 2016, *202*, 40–46.

Weibull, Jörgen, "On Self-Enforcement in Extensive-Form Games," *Games and Economic Behavior*, 1992, *4* (3), 450–462.

A Proofs

A.1 Proofs for Section 2

Proof of Proposition 1. Let Γ be a finite two-player game.

Part One: Let σ be a test-set equilibrium of Γ. Then, σ is a Nash equilibrium in undominated strategies, so it is a trembling-hand perfect equilibrium (Mas-Colell, Whinston and Green, 1995, page 259).

Part Two: Let $\sigma = (\sigma_1, \sigma_2)$ be a proper equilibrium of Γ. However, suppose by way of contradiction that σ is not a test-set equilibrium. Because σ is proper, it must be a Nash equilibrium in undominated strategies. Therefore, it must be the case that the test-set condition fails. Without loss of generality, suppose that player 1's choice of σ_1 is weakly dominated in the test set. Then there exists $\hat{\sigma}_1 \in \Delta(S_1)$ and $\hat{s}_2 \in BR_2(\sigma_1)$ such that (*i*) $\pi_1(\hat{\sigma}_1, \hat{s}_2) > \pi_1(\sigma_1, \hat{s}_2)$, and (*ii*) for all $s_2 \in BR_2(\sigma_1)$, $\pi_1(\hat{\sigma}_1, s_2) \geq \pi_1(\sigma_1, s_2)$.

Because σ is proper, there exist sequences $\{\varepsilon_t\}_{t=0}^{\infty}$ and $\{\sigma^t\}_{t=0}^{\infty}$ such that (*i*) each $\varepsilon_t > 0$ and $\lim_{t\to\infty} \varepsilon_t = 0$, (*ii*) each σ^t is an ε_t-proper equilibrium, and (*iii*) $\lim_{t\to\infty} \sigma^t = \sigma$. By the third criterion, we have that for all $s_2 \in S_2$, $\lim_{t\to\infty} \pi_2(\sigma_1^t, s_2) \to \pi_2(\sigma_1, s_2)$. Therefore, there must exist some T such that for all $t \geq T$ and all $s_2 \notin BR_2(\sigma_1)$, $\pi_2(\sigma_1^t, s_2) < \pi_2(\sigma_1^t, \hat{s}_2)$. By the second criterion, this requires that for all such t and all such s_2, $\sigma_2^t(s_2) \leq \varepsilon_t \sigma_2^t(\hat{s}_2)$.

Now, define $\underline{\Delta} = \pi_1(\hat{\sigma}_1, \hat{s}_2) - \pi_1(\sigma_1, \hat{s}_2) > 0$. Furthermore, let $\bar{\Delta}$ denote the difference between player 1's maximum and minimum payoffs, which exists because the game is finite. Then for all $t \geq T$,

$$\begin{aligned} \pi_1(\hat{\sigma}_1, \sigma_2^t) - \pi_1(\sigma_1, \sigma_2^t) &\geq \underline{\Delta}\sigma_2^t(\hat{s}_2) - \bar{\Delta} \cdot \sum_{s_2 \notin BR_2(\sigma_1)} \sigma_2^t(s_2) \\ &> \left(\underline{\Delta} - \bar{\Delta}|S_2|\varepsilon_t\right) \sigma_2^t(\hat{s}_2), \end{aligned}$$

which is positive for all sufficiently small values of ε_t. Thus, for sufficiently large values of t, σ_1 is not a best response to σ_2^t. This contradicts $\lim_{t\to\infty} \sigma_1^t = \sigma_1$. $\square$

A.2 Proofs for the Menu Auction

Proof of Lemma 2.

Part One: Suppose that b_n denotes a strategy satisfying the conditions of the first claim. Without loss of generality, let $n = 1$. Let $\bar{x} \in X$ be such that $v_1(\bar{x}) = 0$. Note that $b_1(\bar{x}) = 0$. Let $K > \max_x v_0(x) - \min_x v_0(x)$ and $\Delta > \max_x v_1(x)$. Consider any pure bid $\hat{b}_1 \neq b_1$. We will establish that $\hat{b}_1$ does not dominate b_1. Let $\hat{x} \in X$ be such that $\hat{b}_1(\hat{x}) \neq b_1(\hat{x})$. There are two cases:

(i) First, suppose that $v_1(\hat{x}) = 0$. In this case, $b_1(\hat{x}) = 0$, so $\hat{b}_1(\hat{x}) > 0$. Specify a profile of bids for bidders $m \neq 1$ as follows. For all $m > 2$ and all $x \in X$, $b_m(x) = 0$. For bidder 2, define b_2 as follows: $b_2(\hat{x}) = \Delta + K$, and for all $x \neq \hat{x}$, $b_2(x) = 0$. If bidders $m \neq 1$ bid in this way, then b_1 and $\hat{b}_1$ both lead to an auctioneer decision of $\hat{x}$, but bidder 1 makes a smaller payment and therefore receives a higher payoff under b_1 than under $\hat{b}_1$.

(ii) Second, suppose that $v_1(\hat{x}) > 0$. In this case, $b_1(\hat{x}) < v_1(\hat{x})$. If $\hat{b}_1(\hat{x}) > b_1(\hat{x})$, then an appropriate profile of bids for bidders $m \neq 1$ can be specified as in the first case. On the other hand, if $\hat{b}_1(\hat{x}) < b_1(\hat{x})$, then we specify a profile of bids for bidders $m \neq 1$ as follows. For all $m > 2$ and all $x \in X$, $b_m(x) = 0$. For bidder 2, define b_2^δ as follows: $b_2^\delta(\bar{x}) = K + \Delta + b_1(\hat{x})$, $b_2^\delta(\hat{x}) = K + \Delta + v_0(\bar{x}) - v_0(\hat{x}) + \delta$, and for all $x \notin \{\bar{x}, \hat{x}\}$, $b_2^\delta(x) = 0$. If $\delta < b_1(\hat{x}) - \hat{b}_1(\hat{x})$ and if bidders $m \neq 1$ bid in this way, then (i) b_1 leads to an auctioneer decision of $\hat{x}$ and a payoff for bidder 1 of $v_1(\hat{x}) - b_1(\hat{x}) > 0$, but (ii) $\hat{b}_1$ leads to an auctioneer decision of $\bar{x}$ and a payoff for bidder 1 of $v_1(\bar{x}) - \hat{b}_1(\bar{x}) \leq 0$.

As a consequence, b_1 is not dominated by any pure strategy. Furthermore, these arguments can also be extended to rule out dominance by mixed bids.

Part Two: Suppose that b_n denotes a strategy satisfying the conditions of the second claim. Without loss of generality, let $n = 1$. As above, let $\bar{x} \in X$ be such that $v_1(\bar{x}) = 0$, let $K > \max_x v_0(x) - \min_x v_0(x)$, and let $\Delta > \max_x v_1(x)$. Specify the profile of bids for bidders $m \neq 1$ as follows. For all $m > 2$ and all $x \in X$, $b_m(x) = 0$. For bidder 2, define $b_2(\bar{x}) = K + \Delta$ and $b_2(x) = 0$ for all $x \neq \bar{x}$. Bidder 1 has no way to secure a positive payoff against this bid profile and can secure a zero payoff only if the auctioneer decision is $\bar{x}$. Therefore bidder 1's unique best response to this bid profile is to set $b_1(x) = v_1(x)$ for all $x \in X$.

Part Three: The third claim is easily proven by showing that if there exists an $x \in X$ for which $b_n(x) > v_n(x)$, then b_n is dominated by the bid vector $\min(b_n(x), v_n(x))$. □

Proof of Lemma 3. Let b be a Nash equilibrium of the first-price menu auction. Let $x^* = x(b)$. We first argue that for all bidders n, it must be the case that $\hat{b}_n \in BR_n(b_{-n})$ if and only if both $x(\hat{b}_n, b_{-n}) = x^*$ and $\hat{b}_n(x^*) = b_n(x^*)$.

For sufficiency, first observe that since b is a Nash equilibrium, no bidder $m \neq n$ can be using a strictly dominated bid, and so the fact that bidder n strictly prefers setting its bid vector equal to its value vector against certain such bid profiles is irrelevant. Next, observe that the two conditions imply that $(\hat{b}_n, b_{-n})$ leads to the same outcome as (b_n, b_{-n}). Since b is a Nash equilibrium, $b_n \in BR_n(b_{-n})$, and therefore $\hat{b}_n \in BR_n(b_{-n})$ as well. For necessity, recall that, as observed in the text, all of bidder n's best responses must lead to the same decision. Furthermore, all best responses must lead to the same payments for that decision. Since b is a Nash equilibrium, $b_n \in BR_n(b_{-n})$. If also $\hat{b}_n \in BR_n(b_{-n})$, then the two conditions must hold.

Part One (Necessity): Suppose by way of contradiction that there exists some bidder k and some $\hat{x} \in X$ for which $b_k(\hat{x}) < b_k(x^*) - v_k(x^*) + v_k(\hat{x})$. We can therefore define $\varepsilon > 0$ to be

$$\varepsilon = [v_k(\hat{x}) - b_k(\hat{x})] - [v_k(x^*) - b_k(x^*)]\,. \tag{1}$$

Choose some $m \neq k$ and define the following bids for bidders k and m

$$\hat{b}_k(x) = \begin{cases} b_k(x) & \text{if } x \neq \hat{x} \\ b_k(\hat{x}) + \frac{2\varepsilon}{3} & \text{if } x = \hat{x} \end{cases}$$

$$\hat{b}_m(x) = \begin{cases} b_m(x) & \text{if } x \neq \hat{x} \\ v_0(x^*) - v_0(\hat{x}) + \sum_{n=1}^{N} b_n(x^*) - \sum_{n \neq m} b_n(\hat{x}) - \frac{\varepsilon}{3} & \text{if } x = \hat{x} \end{cases}$$

We conclude this direction of the proof by establishing three claims. First, that $\hat{b}_m \in BR_m(b_{-m})$. Second, that $\pi_k(\hat{b}_k, \hat{b}_m, b_{-km}) > \pi_k(b_k, \hat{b}_m, b_{-km})$. Third, for any $n \neq k$ and any $\hat{b}_n \in BR_n(b_{-n})$, that $\pi_k(\hat{b}_k, \hat{b}_n, b_{-kn}) \geq \pi_k(b_k, \hat{b}_n, b_{-kn})$, and in the case of equality, then $x(\hat{b}_k, \hat{b}_n, b_{-kn}) = x(b_k, \hat{b}_n, b_{-kn})$ as well. These three claims constitute a contradiction to b having satisfied the test-set condition.

First, we show that $\hat{b}_m \in BR_m(b_{-m})$. As observed at the beginning of this proof, it suffices to show (*i*) that $\hat{b}_m(x^*) = b_m(x^*)$, which since $x^* \neq \hat{x}$ follows from the definition of $\hat{b}_m$, and (*ii*) $x(\hat{b}_m, b_{-m}) = x^*$. Because $x(b_m, b_{-m}) = x^*$ and because $\hat{b}_m$ agrees with b_m for all $x \neq \hat{x}$, we must have $x(\hat{b}_m, b_{-m}) \in \{x^*, \hat{x}\}$. It therefore only remains to show $x(\hat{b}_m, b_{-m}) \neq \hat{x}$. To see this, we have

$$
\begin{aligned}
&v_0(\hat{x}) + \hat{b}_m(\hat{x}) + \sum_{n \neq m} b_n(\hat{x}) \\
&\quad = v_0(\hat{x}) + \left[v_0(x^*) - v_0(\hat{x}) + \sum_{n=1}^{N} b_n(x^*) - \sum_{n \neq m} b_n(\hat{x}) - \frac{\varepsilon}{3} \right] + \sum_{n \neq m} b_n(\hat{x}) \\
&\quad = v_0(x^*) + b_m(x^*) + \sum_{n \neq m} b_n(x^*) - \frac{\varepsilon}{3} \\
&\quad = v_0(x^*) + \hat{b}_m(x^*) + \sum_{n \neq m} b_n(x^*) - \frac{\varepsilon}{3}.
\end{aligned}
$$

Since $\varepsilon > 0$, this indeed establishes that $x(\hat{b}_m, b_{-m}) \neq \hat{x}$ and completes the proof of this claim.

Second, we show that $\pi_k(\hat{b}_k, \hat{b}_m, b_{-km}) > \pi_k(b_k, \hat{b}_m, b_{-km})$. We have just seen that $x(b_k, \hat{b}_m, b_{-km}) = x^*$. Next, we establish that $x(\hat{b}_k, \hat{b}_m, b_{-km}) = \hat{x}$. Because $x(b_k, \hat{b}_m, b_{-km}) = x^*$ and because $\hat{b}_k$ agrees with b_k for all $x \neq \hat{x}$, we must have $x(\hat{b}_k, \hat{b}_m, b_{-km}) \in \{x^*, \hat{x}\}$. It therefore remains to show $x(\hat{b}_k, \hat{b}_m, b_{-km}) \neq x^*$. To see this, we have

$$
\begin{aligned}
&v_0(\hat{x}) + \hat{b}_k(\hat{x}) + \hat{b}_m(\hat{x}) + \sum_{n \neq k,m} b_n(\hat{x}) \\
&\quad = v_0(\hat{x}) + \left[b_k(\hat{x}) + \frac{2\varepsilon}{3} \right] + \left[v_0(x^*) - v_0(\hat{x}) + \sum_{n=1}^{N} b_n(x^*) - \sum_{n \neq m} b_n(\hat{x}) - \frac{\varepsilon}{3} \right] + \sum_{n \neq k,m} b_n(\hat{x}) \\
&\quad = v_0(x^*) + b_k(x^*) + b_m(x^*) + \sum_{n \neq k,m} b_n(x^*) + \frac{\varepsilon}{3} \\
&\quad = v_0(x^*) + \hat{b}_k(x^*) + \hat{b}_m(\hat{x}^*) + \sum_{n \neq k,m} b_n(x^*) + \frac{\varepsilon}{3}.
\end{aligned}
$$

Since $\varepsilon > 0$, this indeed establishes that $x(\hat{b}_k, \hat{b}_m, b_{-km}) \neq x^*$. Armed with this, we can compare the payoffs of bidder k under the two bid profiles:

$$
\begin{aligned}
\pi_k(\hat{b}_k, \hat{b}_m, b_{-km}) - \pi_k(b_k, \hat{b}_m, b_{-km}) &= \left[v_k(\hat{x}) - \hat{b}_k(\hat{x}) \right] - [v_k(x^*) - b_k(x^*)] \\
&= [v_k(\hat{x}) - b_k(\hat{x})] - [v_k(x^*) - b_k(x^*)] - \frac{2\varepsilon}{3} \\
&= \frac{\varepsilon}{3},
\end{aligned}
$$

where the final step uses the definition of ε in (1). Because $\varepsilon > 0$, this completes the proof of this claim.

Third, we show that for any $n \neq k$ and any $\hat{b}_n \in BR_n(b_{-n})$, that $\pi_k(\hat{b}_k, \hat{b}_n, b_{-kn}) \geq \pi_k(b_k, \hat{b}_n, b_{-kn})$, and in the case of equality, then $x(\hat{b}_k, \hat{b}_n, b_{-kn}) = x(b_k, \hat{b}_n, b_{-kn})$ as well. As observed at the beginning of this proof, if $\hat{b}_n \in BR_n(b_{-n})$, then $x(b_k, \hat{b}_n, b_{-kn}) = x^*$. Because of this and because $\hat{b}_k$ agrees with b_k for all $x \neq \hat{x}$, we must have $x(\hat{b}_k, \hat{b}_n, b_{-kn}) \in \{x^*, \hat{x}\}$. In the first case, $x(\hat{b}_k, \hat{b}_n, b_{-kn}) = \hat{x}$. In this case, the argument from the previous paragraph can be used to show that $\pi_k(\hat{b}_k, \hat{b}_n, b_{-kn}) > \pi_k(b_k, \hat{b}_n, b_{-kn})$. In the second case, $x(\hat{b}_k, \hat{b}_n, b_{-kn}) = x^*$. In this case,

$$\pi_k(\hat{b}_k, \hat{b}_n, b_{-kn}) = v_k(x^*) - \hat{b}_k(x^*) = v_k(x^*) - b_k(x^*) = \pi_k(b_k, \hat{b}_n, b_{-kn}),$$

and in addition, $x(\hat{b}_k, \hat{b}_n, b_{-kn}) = x^* = x(b_k, \hat{b}_n, b_{-kn})$, as required.

Part Two (Sufficiency): Suppose that b violates the test-set condition. A consequence of that is that there exists a bidder k, a pure bid $\hat{b}_k \in BR_k(b_{-k})$, a bidder $m \neq k$, and a pure bid $\hat{b}_m \in BR_m(b_{-m})$ for which $\pi_k(\hat{b}_k, \hat{b}_m, b_{-km}) \geq \pi_k(b_k, \hat{b}_m, b_{-km})$, with equality only if the total surplus from $x(\hat{b}_k, \hat{b}_m, b_{-km})$ exceeds that from $x(b_k, \hat{b}_m, b_{-km})$.

As observed at the beginning of this proof, a necessary condition for $\hat{b}_m \in BR_m(b_{-m})$ is that $x(b_k, \hat{b}_m, b_{-km}) = x^*$. We also define $\hat{x} = x(\hat{b}_k, \hat{b}_m, b_{-km})$. Rewriting the violation of the test-set condition in these terms:

$$v_k(\hat{x}) - \hat{b}_k(\hat{x}) \geq v_k(x^*) - b_k(x^*), \tag{2}$$

with equality only if the total surplus from $\hat{x}$ exceeds that from x^*. We argue that $\hat{x} \neq x^*$. In the case that (2) holds with equality, then the total surplus from $\hat{x}$ exceeds that from x^*, and so we automatically obtain $\hat{x} \neq x^*$. Next, as observed at the beginning of this proof, a necessary condition for $\hat{b}_k \in BR_k(b_{-k})$ is that $\hat{b}_k(x^*) = b_k(x^*)$. Therefore, (2) cannot hold with strict inequality unless $\hat{x} \neq x^*$.

In either case, we have that $x(\hat{b}_k, \hat{b}_m, b_{-km}) = \hat{x} \neq x^* = x(b_k, \hat{b}_m, b_{-km})$. This is only possible if $\hat{b}_k(\hat{x}) > b_k(\hat{x})$. Therefore, (2) implies

$$v_k(\hat{x}) - b_k(\hat{x}) > v_k(x^*) - b_k(x^*),$$

which implies the desired inequality. □

Proof of Corollary 5. Suppose that b is a test-set equilibrium. We argue first that $x(b) = x^{\mathcal{N}}$, and second that b induces payoffs in C.[17]

Part One (Surplus-Maximizing Decision): Define $x^* = x(b)$ and suppose by way of contradiction that $x^* \neq x^{\mathcal{N}}$. By Lemma 3, we have that for all bidders n,

$$b_n(x^{\mathcal{N}}) \geq b_n(x^*) - v_n(x^*) + v_n(x^{\mathcal{N}}).$$

Summing over n,

$$\sum_{n=1}^{N} b_n(x^{\mathcal{N}}) \geq \sum_{n=1}^{N} b_n(x^*) - \sum_{n=1}^{N} v_n(x^*) + \sum_{n=1}^{N} v_n(x^{\mathcal{N}}).$$

[17] This proof draws from the proof of Theorem 2 of Bernheim and Whinston (1986).

Equivalently,

$$v_0(x^{\mathcal{N}}) + \sum_{n=1}^{N} b_n(x^{\mathcal{N}}) \geq v_0(x^*) + \sum_{n=1}^{N} b_n(x^*) - \sum_{n=0}^{N} v_n(x^*) + \sum_{n=0}^{N} v_n(x^{\mathcal{N}}). \tag{3}$$

By the definition of $x^{\mathcal{N}}$, we also have

$$\sum_{n=0}^{N} v_n(x^{\mathcal{N}}) > \sum_{n=0}^{N} v_n(x^*), \tag{4}$$

where the strictness of the inequality is by Assumption 1. Plugging (4) into (3), we obtain

$$v_0(x^{\mathcal{N}}) + \sum_{n=1}^{N} b_n(x^{\mathcal{N}}) > v_0(x^*) + \sum_{n=1}^{N} b_n(x^*),$$

which contradicts $x(b) = x^*$.

Part Two (Core Payoffs): Next, we argue that the test-set equilibrium b generates payoffs $\pi \in C$. Let $J \subseteq \mathcal{N}$. Because $x(b) = x^{\mathcal{N}}$,

$$v_0(x^{\mathcal{N}}) + \sum_{n=1}^{N} b_n(x^{\mathcal{N}}) \geq v_0(x^{\bar{J}}) + \sum_{n=1}^{N} b_n(x^{\bar{J}}).$$

Consequently,

$$v_0(x^{\mathcal{N}}) + \sum_{n\in J} b_n(x^{\mathcal{N}}) + \sum_{n\in \bar{J}} b_n(x^{\mathcal{N}}) \geq v_0(x^{\bar{J}}) + \sum_{n\in \bar{J}} b_n(x^{\bar{J}}).$$

Since $x(b) = x^{\mathcal{N}}$, we have from Lemma 3 that for all bidders n, $b_n(x^{\bar{J}}) \geq b_n(x^{\mathcal{N}}) - v_n(x^{\mathcal{N}}) + v_n(x^{\bar{J}})$. Plugging this into the above,

$$v_0(x^{\mathcal{N}}) + \sum_{n\in J} b_n(x^{\mathcal{N}}) + \sum_{n\in \bar{J}} b_n(x^{\mathcal{N}}) \geq v_0(x^{\bar{J}}) + \sum_{n\in \bar{J}} b_n(x^{\mathcal{N}}) - \sum_{n\in \bar{J}} v_n(x^{\mathcal{N}}) + \sum_{n\in \bar{J}} v_n(x^{\bar{J}}).$$

Rearranging, we obtain

$$\sum_{n\in J} b_n(x^{\mathcal{N}}) \geq \sum_{n\in\{0\}\cup \bar{J}} v_n(x^{\bar{J}}) - \sum_{n\in\{0\}\cup \bar{J}} v_n(x^{\mathcal{N}}).$$

Thus, the the equilibrium payoffs for coalition J are

$$\begin{aligned}\sum_{n\in J} \pi_n &= \sum_{n\in J} v_n(x^{\mathcal{N}}) - \sum_{n\in J} b_n(x^{\mathcal{N}}) \\ &\leq \sum_{n=0}^{N} v_n(x^{\mathcal{N}}) - \sum_{n\in\{0\}\cup \bar{J}} v_n(x^{\bar{J}}),\end{aligned}$$

as desired. $\square$

A.3 Proofs for the GSP Auction

Lemma 11 states that in any pure Nash equilibrium, none of the highest $\min\{I, N\}$ bidders will be tied with any other bidder. It will be helpful in proving Theorem 6, which is the main result for this application.

Lemma 11. *If $b = (b_1, \ldots, b_N)$ is a Nash equilibrium of the GSP auction, then for all $i \in \{1, \ldots, \min\{I, N-1\}\}$, $b^{(i)} > b^{(i+1)}$.*

Proof of Lemma 11. Suppose $i \in \{1, \ldots, \min\{I, N-1\}\}$ and K are such that $b^{(i)} = \cdots = b^{(i+K)} = b^*$ is a "maximal tie." That is to say, suppose there are exactly $K+1$ bidders who bid b^*. We derive a contradiction from $K \geq 1$. First, note that any bidder with per-click value v who is part of the tie earns the payoff

$$\frac{1}{K+1}\left[\kappa_{i+K}\left(v - b^{(i+K+1)}\right) + \sum_{k=0}^{K-1} \kappa_{i+k}(v - b^*)\right].$$

By raising its bid to just above b^*, the bidder would earn the payoff

$$\kappa_i\left(v - b^*\right).$$

If $b^* = 0$, then this deviation would be a profitable one. We therefore assume henceforth that $b^* > 0$. In that case, the bidder can reduce its bid to just below b^*, which would lead to the payoff

$$\kappa_{i+K}\left(v - b^{(i+K+1)}\right).$$

Let $v' < v''$ be the per-click values of two of the tied bidders (Assumption 2(ii) rules out the possibility of equality). Because the bidder with value v' does not find it profitable to deviate to just below b^*,

$$\begin{aligned}
&\frac{1}{K+1}\left[\kappa_{i+K}\left(v' - b^{(i+K+1)}\right) + \sum_{k=0}^{K-1} \kappa_{i+k}(v' - b^*)\right] \geq \kappa_{i+K}\left(v' - b^{(i+K+1)}\right) \\
&\implies \kappa_{i+K}\left(v' - b^{(i+K+1)}\right) + K\kappa_i(v' - b^*) \geq (K+1)\kappa_{i+K}\left(v' - b^{(i+K+1)}\right) \\
&\implies \kappa_i(v' - b^*) \geq \kappa_{i+K}\left(v' - b^{(i+K+1)}\right). \qquad (5)
\end{aligned}$$

Because the bidder with value v'' does not find it profitable to deviate to just above b^*,

$$\begin{aligned}
&\frac{1}{K+1}\left[\kappa_{i+K}\left(v'' - b^{(i+K+1)}\right) + \sum_{k=0}^{K-1} \kappa_{i+k}(v'' - b^*)\right] \geq \kappa_i\left(v'' - b^*\right) \\
&\implies \kappa_{i+K}\left(v'' - b^{(i+K+1)}\right) + K\kappa_i(v'' - b^*) \geq (K+1)\kappa_i\left(v'' - b^*\right) \\
&\implies \kappa_{i+K}\left(v'' - b^{(i+K+1)}\right) \geq \kappa_i\left(v'' - b^*\right). \qquad (6)
\end{aligned}$$

Adding together equations (5) and (6) and canceling like terms, we obtain $\kappa_i v' + \kappa_{i+K} v'' \geq \kappa_{i+K} v' + \kappa_i v''$. Equivalently,

$$(\kappa_i - \kappa_{i+K})(v' - v'') \geq 0.$$

By Assumption 2(i) and the fact that $K \geq 1$, we have $\kappa_i > \kappa_{i+K}$. Furthermore, we previously supposed $v' < v''$. This is therefore a contradiction. $\square$

Proof of Theorem 6. Suppose that b is a pure Nash equilibrium that is not locally envy-free. By Lemma 11, none of the highest $\min\{I, N\}$ bidders will be tied with any other bidder. Therefore, for all $i \in \{1, \ldots, \min\{I, N\}\}$, $G(i)$ is a singleton, the unique element of which we denote $g(i)$. Note, however, that in the case where $N > I$, $G(I+1)$ is defined as the $N - I$ lowest bids, and might therefore not be a singleton.

Let i^* be the largest index for which the locally envy-free inequality is violated, so that for some element of $G(i^*)$, which we henceforth denote $g(i^*)$,

$$\kappa_{i^*-1}\left(v_{g(i^*)} - b^{(i^*)}\right) > \kappa_{i^*}\left(v_{g(i^*)} - b^{(i^*+1)}\right). \tag{7}$$

We use b^* to denote the equilibrium bid of $g(i^*)$. For the case in which $i^* \leq I$, $b^* = b^{(i^*)}$. For the case in which $i^* = I + 1$, $b^* \in \left[0, b^{(i^*)}\right]$.

With this established, the proof consists of three parts. First, we demonstrate that bidders who bid lower than $g(i^*)$ are sorted by their values for clicks. We then show that b fails the test-set condition by demonstrating the existence of an $\varepsilon > 0$ such that $b^{(i^*)} + \varepsilon$ weakly dominates b^* in the test set. The second part establishes this for the case in which $i^* = I + 1$. The third part establishes this for the case in which $i^* \leq I$.

Part One: Let $k \in \{1, \ldots, I + 1 - i^*\}$. Suppose that $g(i^* + k) \in G(i^* + k)$. Because i^* was defined as the largest index for which the locally envy-free inequality is violated, we have

$$\kappa_{i^*+k}\left(v_{g(i^*+k)} - b^{(i^*+k+1)}\right) \geq \kappa_{i^*+k-1}\left(v_{g(i^*+k)} - b^{(i^*+k)}\right).$$

Furthermore, equilibrium requires that bidder $g(i^* + k - 1)$ cannot profit by deviating to just below $b^{(i^*+k)}$. Thus,

$$\kappa_{i^*+k-1}\left(v_{g(i^*+k-1)} - b^{(i^*+k)}\right) \geq \kappa_{i^*+k}\left(v_{g(i^*+k-1)} - b^{(i^*+k+1)}\right).$$

Manipulating these inequalities yields

$$\left(\kappa_{i^*+k-1} - \kappa_{i^*+k}\right)\left(v_{g(i^*+k-1)} - v_{g(i^*+k)}\right) \geq 0.$$

By Assumption 2(i), $\kappa_{i^*+k-1} > \kappa_{i^*+k}$, and so $v_{g(i^*+k-1)} \geq v_{g(i^*+k)}$. By Assumption 2(ii), the inequality must actually be strict. Because we can make this argument for all $k \in \{1, \ldots, I + 1 - i^*\}$, we conclude that the bidders below $g(i^*)$ are sorted by value:

$$v_{g(i^*)} > v_{g(i^*+1)} > \cdots > \max_{g(I+1) \in G(I+1)} v_{g(I+1)}.$$

This observation will be useful in part three of this proof.

Part Two: Suppose that $i^* = I + 1$. Define

$$\varepsilon = \frac{1}{2}\min\left\{b^{(I)} - b^{(I+1)}, v_{g(i^*)} - b^{(I+1)}\right\}.$$

By Lemma 11, $b^{(I)} > b^{(I+1)}$, which implies that the first component of the minimum is positive. That the second component is positive follows from (7), the violation of the locally envy-free inequality for $g(i^*)$. (To see this, recall that we use the convention $\kappa_{I+1} = 0$.) Thus, $\varepsilon > 0$.

We now compare whether bidder $g(i^*)$ is better off by playing $b^{(I+1)} + \varepsilon$ or b^* in the test set. By the definition of ε, $b^{(I+1)} + \varepsilon < b^{(I)}$. Therefore, both bids perform the same in the event that none of the competitors of $g(i^*)$ deviate. We complete the analysis by considering three classes of deviations.

(i) First, b^* and $b^{(I+1)} + \varepsilon$ perform equally well against all elements of the test set in which the deviating bidder does not deviate within the interval $\left[b^*, b^{(I+1)} + \varepsilon\right]$.

(ii) Second, suppose that the deviator is $g(i)$ for some $i \leq I$ and that the deviation is to within the interval $\left[b^*, b^{(I+1)} + \varepsilon\right]$ (i.e. a higher bidder who deviates downward). To begin, assume that the deviation, which we denote $\hat{b}$, is to the interior of the interval. Then the incremental payoff that $g(i^*)$ receives from playing $b^{(I+1)} + \varepsilon$ instead of b^* is

$$\kappa_I \left(v_{g(i^*)} - \hat{b}\right) \geq \kappa_I \left(v_{g(i^*)} - b^{(I+1)} - \varepsilon\right) > 0.$$

The last step in the above uses the fact that by definition, $\varepsilon < v_{g(i^*)} - b^{(I+1)}$. Finally, if the deviation is to one of the endpoints of the interval, then the tie will be broken randomly, so the expected gains are at least half of those above, and therefore still positive.

(iii) Third, suppose that the deviator is another member of $G(I+1)$ and that the deviation is to within the interval $\left[b^*, b^{(I+1)} + \varepsilon\right]$ (i.e. a lower bidder who deviates upward). In these cases, $b^{(I+1)} + \varepsilon$ and b^* perform equally well: both result in a payoff of zero for $g(i^*)$.

Moreover, notice that since $b^{(I+1)} + \varepsilon$ is a best response to equilibrium for bidder $g(I)$, there is at least one element of type (ii) in the test set. Therefore, $b^{(I+1)} + \varepsilon$ weakly dominates b^* in $T(b)$. Consequently, b is not a test-set equilibrium.

Part Three: Suppose that $i^* \leq I$. Define

$$\varepsilon = \frac{1}{2} \min \left\{ b^{(i^*-1)} - b^*, v_{g(i^*)} - b^* - \frac{\kappa_{i^*}}{\kappa_{i^*-1}} \left(v_{g(i^*)} - b^{(i^*+1)}\right) \right\}.$$

By Lemma 11, $b^{(i^*-1)} > b^*$, which implies that the first component of the minimum is positive. That the second component is positive follows from (7), the violation of the locally envy-free inequality for $g(i^*)$. Thus, $\varepsilon > 0$.

We now compare whether bidder $g(i^*)$ is better off by playing $b^* + \varepsilon$ or b^* in the test set. By the definition of ε, $b^* + \varepsilon < b^{(i^*-1)}$. Therefore, both bids perform the same in the event that none of the competitors of $g(i^*)$ deviate. We complete the analysis by considering three classes of deviations.

(i) First, b^* and $b^* + \varepsilon$ perform equally well against all elements of the test set in which the deviating bidder does not deviate within the interval $[b^*, b^* + \varepsilon]$.

(ii) Second, suppose that the deviator is $g(i)$ for some $i < i^*$ and that the deviation is to within the interval $[b^*, b^* + \varepsilon]$ (i.e. a higher bidder who deviates downward). To begin, assume that the deviation, which we denote $\hat{b}$, is to the interior of the interval. Then the incremental payoff that $g(i^*)$ receives from playing $b^* + \varepsilon$ instead of b^* is

$$\begin{aligned}
&\kappa_{i^*-1}\left(v_{g(i^*)} - \hat{b}\right) - \kappa_{i^*}\left(v_{g(i^*)} - b^{(i^*+1)}\right) \\
&> \kappa_{i^*-1}\left(v_{g(i^*)} - b^* - \varepsilon\right) - \kappa_{i^*}\left(v_{g(i^*)} - b^{(i^*+1)}\right) \\
&> 0.
\end{aligned}$$

The last step in the above uses the fact that by definition, $\varepsilon < v_{g(i^*)} - b^* - \frac{\kappa_{i^*}}{\kappa_{i^*-1}}\left(v_{g(i^*)} - b^{(i^*+1)}\right)$. Finally, if the deviation is to one of the endpoints of the interval, then the tie will be broken randomly, so the expected gains are only half of those above, yet are still positive.

(iii) The third remaining possibility is that the deviator is $g(i) \in G(i)$ for some $i > i^*$ and that the deviation is to within the interval $[b^*, b^* + \varepsilon]$ (i.e. a lower bidder who deviates upward). We argue that there are no elements in the test set of this form by showing that this cannot be a best response for $g(i)$.

Because i^* was defined as the largest index for which the locally envy-free inequality is violated, we have that for all $k \in [1, i - i^*]$,

$$\kappa_{i^*+k}\left(v_{g(i^*+k)} - b^{(i^*+k+1)}\right) \geq \kappa_{i^*+k-1}\left(v_{g(i^*+k)} - b^{(i^*+k)}\right).$$

From part one of this proof, we have that for all $k \in [1, i - i^*]$, $v_{g(i^*+k)} \geq v_{g(i)}$, which implies

$$\kappa_{i^*+k}\left(v_{g(i)} - b^{(i^*+k+1)}\right) \geq \kappa_{i^*+k-1}\left(v_{g(i)} - b^{(i^*+k)}\right).$$

Summing the above equation across all $k \in [1, i - i^*]$, then canceling like terms, we obtain

$$\kappa_i\left(v_{g(i)} - b^{(i+1)}\right) \geq \kappa_{i^*}\left(v_{g(i)} - b^{(i^*+1)}\right),$$

which implies

$$\kappa_i\left(v_{g(i)} - b^{(i+1)}\right) > \kappa_{i^*}\left(v_{g(i)} - b^*\right).$$

This implies that the deviation is not a best response for $g(i)$, as desired.

Moreover, notice that since $b^* + \varepsilon$ is a best response to equilibrium for bidder $g(i^* - 1)$, there is at least one element of type (ii) in the test set. Therefore, $b^* + \varepsilon$ weakly dominates b^* in $T(b)$. Consequently, b is not a test-set equilibrium. □

Proof of Proposition 7. For ease of reference, we restate here the condition in the proposition, and we denote it ($\star$):

$$\frac{v_3}{v_2} \leq \frac{\kappa_2^2 - \kappa_1\kappa_3}{\kappa_2^2 - \kappa_2\kappa_3}. \tag{$\star$}$$

The proof proceeds in two parts. First, we demonstrate by construction that ($\star$) is sufficient for the existence of a test-set equilibrium in the environment with three bidders and three positions. Second, we also demonstrate that ($\star$) is necessary for the existence of a test-set equilibrium in this environment. Our strategy for the latter is to use Theorem 6 to establish a lower bound on bidder 2's bid in any test-set equilibrium. We then also establish a corresponding upper bound and demonstrate that both cannot be simultaneously satisfied if ($\star$) is violated.

Part One (Sufficiency): We argue that if ($\star$) holds, then the following is a test-set equilibrium:

$$
\begin{aligned}
b_1 &= \left(1 - \frac{\kappa_2}{\kappa_1}\right) v_1 + \frac{\kappa_2 - \kappa_3}{\kappa_1} v_3 \\
b_2 &= \left(1 - \frac{\kappa_2}{\kappa_1}\right) v_2 + \frac{\kappa_2 - \kappa_3}{\kappa_1} v_3 \\
b_3 &= \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_3
\end{aligned}
$$

We can see from these expressions that $b_1 > b_2 > b_3$. We can also see that the best responses to equilibrium for bidder 3 are $\hat{b}_3 \in [0, b_2)$, for bidder 2 are $\hat{b}_2 \in (b_3, b_1)$, and for bidder 1 are $\hat{b}_1 \in (b_2, \infty)$.

Using this, we now check that bidder 3 does not have an alternate bid that weakly dominates b_3 either in the test set or in the game. For brevity, we limit attention here to pure strategy dominance. (Similar arguments also establish the absence of mixed strategy dominance.) Suppose, by way of contradiction, that such a bid, b_3', exists. Since b_3' must be a best response to equilibrium, $b_3' \in [0, b_2)$. There are two cases.

(i) First, suppose $b_3' < b_3$. Both b_3' and b_3 perform the same against all elements of the test set, so b_3' cannot weakly dominate b_3 in the test set. Moreover, b_3' cannot weakly dominate b_3 in the game because b_3 outperforms b_3' when bidder 2 deviates to any $\hat{b}_2 \in (b_3', b_3)$. Indeed, bidder 3's payoff from playing b_3' against this deviation is $\kappa_3 v_3$. On the other hand, bidder 3's payoff from playing b_3 against this deviation is

$$\kappa_2(v_3 - \hat{b}_2) > \kappa_2(v_3 - b_3) = \kappa_3 v_3.$$

(ii) Second, suppose that $b_3' > b_3$. Then we also have a contradiction, since b_3 outperforms b_3' when bidder 2 deviates to any $\hat{b}_2 \in (b_3, b_3')$, which can happen in the test set. Indeed, bidder 3's payoff from playing b_3 against this deviation is $\kappa_3 v_3$. On the other hand, bidder 3's payoff from playing b_3' against this deviation is

$$\kappa_2(v_3 - \hat{b}_2) < \kappa_2(v_3 - b_3) = \kappa_2 v_3 - \kappa_2 v_3 + \kappa_3 v_3 = \kappa_3 v_3.$$

Similarly, suppose that bidder 2 has an alternate bid, b_2', that weakly dominates b_2 either in the test set or in the game. Since b_2' must be a best response to equilibrium, $b_2' \in (b_3, b_1)$. There are two cases.

(i) First, suppose $b_2' > b_2$. Then we have a contradiction, since b_2 outperforms b_2' when bidder 1 deviates to any $\hat{b}_1 \in (b_2, b_2')$, which can happen in the test set. Indeed, bidder 2's payoff from playing b_2' against this deviation is

$$\kappa_1(v_2 - \hat{b}_1) < \kappa_1(v_2 - b_2) = \kappa_2 v_2 - \kappa_2 v_3 + \kappa_3 v_3.$$

On the other hand, bidder 2's payoff from playing b_2 against this deviation is

$$\kappa_2(v_2 - b_3) = \kappa_2 v_2 - \kappa_2 v_3 + \kappa_3 v_3.$$

(ii) Second, suppose $b_2' < b_2$. Then we also have a contradiction, since b_2 outperforms b_2' when bidder 3 deviates to any $\hat{b}_3 \in (b_2', b_2)$, which can happen in the test set. Indeed, bidder 2's payoff from playing b_2' against this deviation is $\kappa_3 v_2$. On the other hand, bidder 2's payoff from playing b_2 against this deviation is

$$\begin{aligned} \kappa_2(v_2 - \hat{b}_3) &> \kappa_2(v_2 - b_2) \\ &= \frac{\kappa_2^2}{\kappa_1} v_2 - \frac{\kappa_2^2 - \kappa_2\kappa_3}{\kappa_1} v_3 \\ &\geq \frac{\kappa_2^2}{\kappa_1} v_2 - \frac{\kappa_2^2 - \kappa_2\kappa_3}{\kappa_1} \cdot \frac{\kappa_2^2 - \kappa_1\kappa_3}{\kappa_2^2 - \kappa_2\kappa_3} v_2 \\ &= \kappa_3 v_2, \end{aligned}$$

where ($\star$) is used in the penultimate step of the above.

Lastly, suppose that bidder 1 has an alternate bid, b_1', that weakly dominates b_1 in the test set. Since b_1' must be a best response to equilibrium, $b_1' \in (b_2, \infty)$. There are two cases.

(i) First, suppose $b_1' > b_1$. Both b_1' and b_1 perform the same against all elements of the test set, so b_1' cannot weakly dominate b_1 in the test set. Moreover, b_1' cannot weakly dominate b_1 in the game because b_1 outperforms b_1' when bidder 2 deviates to any $\hat{b}_2 \in (b_1, b_1')$. Indeed, bidder 1's payoff from playing b_1' against this deviation is

$$\kappa_1(v_1 - \hat{b}_2) < \kappa_1(v_1 - b_1) = \kappa_2 v_1 - \kappa_2 v_3 + \kappa_3 v_3.$$

On the other hand, bidder 1's payoff from playing b_1 against this deviation is

$$\kappa_2(v_1 - b_3) = \kappa_2 v_1 - \kappa_2 v_3 + \kappa_3 v_3.$$

(ii) Second, suppose that $b_1' < b_1$. Then we also have a contradiction, since b_1 outperforms b_1' when bidder 2 deviates to any $\hat{b}_2 = (b_1', b_1)$, which can happen in the test set. Indeed, bidder 1's payoff from playing b_1' against this deviation is

$$\kappa_2(v_1 - b_3) = \kappa_2 v_1 - \kappa_2 v_3 + \kappa_3 v_3.$$

On the other hand, bidder 1's payoff from playing b_1 against this deviation is

$$\begin{aligned} \kappa_1(v_1 - \hat{b}_2) &> \kappa_1(v_1 - b_1) \\ &= \kappa_2 v_1 - \kappa_2 v_3 + \kappa_3 v_3. \end{aligned}$$

We conclude that (b_1, b_2, b_3) is, indeed, a test-set equilibrium.

Part Two (Necessity): Suppose that $b = (b_1, b_2, b_3)$ is a pure test-set equilibrium. By Theorem 6, the equilibrium is locally envy-free. Locally envy-free equilibria feature assortative matching, so $b_1 > b_2 > b_3$. Furthermore, since b is locally envy-free, bidder 3 must not envy bidder 2. That is, we must have $\kappa_2(v_3 - b_3) \leq \kappa_3 v_3$. Equivalently,

$$b_3 \geq \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_3. \tag{8}$$

Similarly, bidder 2 must not envy bidder 1. That is, we must have $\kappa_1(v_2 - b_2) \leq \kappa_2(v_2 - b_3)$. Equivalently,

$$b_2 \geq v_2 - \frac{\kappa_2}{\kappa_1}(v_2 - b_3). \tag{9}$$

Substituting (8) into (9), we obtain

$$b_2 \geq v_2 - \frac{\kappa_2}{\kappa_1}(v_2 - v_3) - \frac{\kappa_3}{\kappa_1} v_3. \tag{10}$$

Equation (10) is the desired lower bound on b_2. We next establish the following upper bound on b_2:

$$b_2 \leq \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_2. \tag{11}$$

To see that this must be the case, assume by way of contradiction that $b_2 > \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_2$. Then define

$$\varepsilon = \frac{1}{2} \min \left\{ b_2 - \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_2, b_2 - b_3, \left(1 - \frac{\kappa_2}{\kappa_1}\right)(v_1 - v_2) \right\}.$$

By assumption, the first component of the minimum is positive. We also know that $b_2 > b_3$, so the second component of the minimum is positive. That the third component is positive follows from Assumption 2. Thus, $\varepsilon > 0$.

We then compare the performance of $b_2 - \varepsilon$ to that of b_2 in the test set $T(b)$. By the definition of ε, $b_2 - \varepsilon > b_3$. Therefore, both bids perform the same in the event that none of the competitors of bidder 2 deviate. We complete the analysis by considering three classes of deviations.

(i) First, b_2 and $b_2 - \varepsilon$ perform equally well against all elements of the test set in which the deviating bidder does not deviate within the interval $[b_2 - \varepsilon, b_2]$.

(ii) Second, suppose that the deviator is bidder 1 and that the deviation is to within the interval $[b_2 - \varepsilon, b_2]$. For this to be an element of the test set, it must be a best response to equilibrium for bidder 1. This requires $\kappa_1(v_1 - b_2) = \kappa_2(v_1 - b_3)$. To begin, we assume that the deviation, which we denote $\hat{b}_1$, is to the interior of the interval. Then the incremental payoff that bidder 2 receives from playing $b_2 - \varepsilon$ instead of b_2 is

$$\begin{aligned}\kappa_2(v_2 - b_3) - \kappa_1(v_2 - \hat{b}_1) &> \kappa_2(v_2 - b_3) - \kappa_1(v_2 - b_2 + \varepsilon) \\ &= \kappa_2(v_2 - b_3) - \kappa_1(v_2 - b_2 + \varepsilon) - \kappa_2(v_1 - b_3) + \kappa_1(v_1 - b_2) \\ &= (\kappa_1 - \kappa_2)(v_1 - v_2) - \kappa_1\varepsilon \\ &> 0.\end{aligned}$$

The last step in the above uses the fact that by definition, $\varepsilon < \left(1 - \frac{\kappa_2}{\kappa_1}\right)(v_1 - v_2)$. Finally, if the deviation is to one of the endpoints of the interval, then the tie will be broken randomly, so the expected gains are only half of those above, yet are still positive.

(iii) The third remaining possibility is that the deviator is bidder 3 and that the deviation is to within the interval $[b_2 - \varepsilon, b_2]$. To begin, we assume that the deviation, which we denote $\hat{b}_3$, is to the interior of the interval. Then the incremental payoff that bidder 2 receives from playing $b_2 - \varepsilon$ instead of b_2 is

$$\kappa_3 v_2 - \kappa_2(v_2 - \hat{b}_3) > \kappa_3 v_2 - \kappa_2(v_2 - b_2 + \varepsilon) > 0.$$

The last step in the above uses the fact that by definition, $\varepsilon < b_2 - \left(1 - \frac{\kappa_3}{\kappa_2}\right) v_2$. Finally, if the deviation is to one of the endpoints of the interval, then the tie will be broken randomly, so the expected gains are only half of those above, yet are still positive.

Moreover, notice that since $b_2 - \varepsilon$ is a best response to equilibrium for bidder 3, there is at least one element of type (iii) in the test set. Therefore, $b_2 - \varepsilon$ weakly dominates b_2 in the test set. This is the desired contradiction, which establishes the necessity of (11).

Thus we must have both (11), an upper bound on b_2, and (10), a lower bound on b_2. In order for both to be simultaneously satisfied we must have that, as claimed,

$$\frac{v_3}{v_2} \leq \frac{\kappa_2^2 - \kappa_1\kappa_3}{\kappa_2^2 - \kappa_2\kappa_3}. \qquad \square$$

A.4 Proofs for the Second-Price, Common Value Auction

Proof of Proposition 8. We first argue that the stated strategy profiles are in undominated strategies and satisfy the test-set condition. We then complete the proof by arguing that all other pure strategy profiles in undominated strategies fail the test-set condition.

Part One: In both purported test-set equilibria, the strategy of the informed bidder is to bid the value of the object. This is a dominant strategy for the informed bidder, which is therefore both undominated in the test set and undominated in the game. Thus, it suffices to consider the uninformed bidder in what follows.

Consider the equilibrium $(0, 1, 0)$. Let $\hat{b}_U \neq 0$ be an alternative pure bid. For all $\delta \in (0, \hat{b}_U)$, the bid profile $(\delta, 1, 0)$ is in the test set $T(0, 1, 0)$. Moreover, $\pi_U(\delta, 1, \hat{b}_U) < \pi_U(\delta, 1, 0)$. As a consequence, no pure bid dominates a bid of 0 either in the game or in the test set for the uninformed bidder. Furthermore, the argument can be extended to derive a similar conclusion for mixed bids: if $\sigma_U \neq 0$ is an alternative mixed bid, then for all sufficiently small values of δ, $(\delta, 1, 0) \in T(0, 1, 0)$ and $\pi_U(\delta, 1, \sigma_U) < \pi_U(\delta, 1, 0)$, so that no mixed bid dominates a bid of 0 either in the game or in the test set for the uninformed bidder.

Consider the equilibrium $(0, 1, 1)$. Let $\hat{b}_U \neq 1$ be an alternative pure bid. For all $\delta \in (0, |1-\hat{b}_U|)$, the bid profiles $(0, 1-\delta, 1)$ and $(0, 1+\delta, 1)$ are both in the the test set $T(0, 1, 1)$. Moreover, one of the following two conditions holds: (i) $\pi_U(0, 1-\delta, \hat{b}_U) < \pi_U(0, 1-\delta, 1)$, while $\pi_U(0, 1+\delta, \hat{b}_U) = \pi_U(0, 1+\delta, 1)$, or (ii) $\pi_U(0, 1+\delta, \hat{b}_U) < \pi_U(0, 1+\delta, 1)$, while $\pi_U(0, 1-\delta, \hat{b}_U) = \pi_U(0, 1-\delta, 1)$. As a consequence, no pure dominates a bid of 1 either in the game or in the test set for the uninformed bidder. Furthermore, the argument can be extended to derive a similar conclusion for mixed bids: if $\sigma_U \neq 1$ is an alternative mixed bid, then for all sufficiently small values of δ, $\{(0, 1-\delta, 1), (0, 1+\delta, 1)\} \subset T(0, 1, 1)$ and at least one of the following holds: $\pi_U(0, 1-\delta, \sigma_U) < \pi_U(0, 1-\delta, 1)$ or $\pi_U(0, 1+\delta, \sigma_U) < \pi_U(0, 1+\delta, 1)$, so that no mixed bid dominates a bid of 1 either in the game or in the test set for the uninformed bidder.

Part Two: Test-set equilibrium requires players to play undominated strategies. Thus, in any test-set equilibrium, the informed bidder must play its dominant strategy of bidding the value of the object. In addition, all bids $b_U \notin [0, 1]$ are dominated for the uninformed bidder. It therefore remains to show that for any $b_U \in (0, 1)$, the bid profile $(0, 1, b_U)$ is not a test-set equilibrium.

The informed bidder's best responses to the uninformed bidder's bid of b_U are $\{(\hat{b}_0, \hat{b}_1) \mid \hat{b}_0 \in [0, b_U), \hat{b}_1 \in (b_U, \infty)\}$. The uninformed bidder's best responses to the informed bidder's strategy of bidding the value of the object are $\hat{b}_U \in [0, \infty)$. Therefore, the test set is

$$T(0, 1, b_U) = \{(\hat{b}_0, \hat{b}_1, b_U) \mid \hat{b}_0 \in [0, b_U), \hat{b}_1 \in (b_U, \infty)\} \cup \{(0, 1, \hat{b}_U) \mid \hat{b}_U \in [0, \infty)\}.$$

For the uninformed bidder, the alternative bid $\hat{b}_U = 0$ weakly dominates b_U in the test set.[18] In particular, the bid of zero does strictly better than b_U against a strategy for the informed bidder of $(b_U/2, 1)$, and it never does worse in the test set. □

A.5 Proofs for Section 5

Proof of Proposition 9. Let Γ be a game in normal form, and fix any $c \geq 0$.

Sufficiency: Suppose σ is not a Nash equilibrium of Γ. There then exists a player n and a strategy $s'_n \in S_n$ such that $\pi_n(s'_n, \sigma_{-n}) > \pi_n(\sigma_n, \sigma_{-n})$. Let b be a behavior strategy profile of $\bar{\Gamma}(c)$ such that $\sigma = \alpha(b)$. Then player n's payoff in $\bar{\Gamma}(c)$ from playing b_n when all other players play according to b is at most $\pi_n(\sigma_n, \sigma_{-n})$. On the other hand, define an alternative behavior strategy for player n, which we denote b'_n, as follows: (i) play according to s'_n at the recommendation information set, (ii) approve at each review information set, and (iii) play arbitrarily at each final decision information set. Player n's payoff in $\bar{\Gamma}(c)$ from playing b'_n when all other players play according to b is $\pi_n(s'_n, \sigma_{-n})$. Therefore, b'_n is a profitable deviation for player n, and b is not a Nash equilibrium of $\bar{\Gamma}(c)$.

[18] In fact, any $\hat{b}_U \neq b_U$ in the unit interval weakly dominates b_U in the test set.

Necessity: Suppose σ is a Nash equilibrium of Γ. For each player n, define a behavior strategy for $\bar{\Gamma}(c)$, which we denote b_n, as follows: (i) play according to σ_n at the recommendation information set, (ii) approve at each review information set, and (iii) play arbitrarily at each final decision information set. Let $b = (b_1, \ldots, b_N)$. Player n's payoff in $\bar{\Gamma}(c)$ from playing a behavior strategy b'_n when all other players play according to b is

$$\bar{\pi}_n(b'_n, b_{-n}) = \pi_n(\alpha_n(b'_n), \alpha_{-n}(b_{-n})) - \beta_n(b'_n)c = \pi_n(\alpha_n(b'_n), \sigma_{-n}) - \beta_n(b'_n)c.$$

This payoff is maximized when $\alpha_n(b'_n) \in BR_n(\sigma_{-n})$ and $\beta_n(b'_n) = 0$. Since this is indeed the case when $b'_n = b_n$, player n does not possess a profitable deviation, and b is a Nash equilibrium of $\bar{\Gamma}(c)$. $\square$

Proof of Theorem 10 (Sufficiency). Suppose that Γ is a finite game in normal form. Suppose σ is such that there exists a $\bar{c} > 0$ such that for all $c \in (0, \bar{c})$, there exists a quasi*-perfect equilibrium of $\bar{\Gamma}(c)$, b, such that $\sigma = \alpha(b)$. If, for all players n, $BR_n(\sigma_{-n}) = S_n$, then let $\underline{\Delta}$ be any positive number. Otherwise, define $\underline{\Delta} > 0$ as the minimum difference between the payoffs in Γ from best and inferior responses to σ:

$$\underline{\Delta} = \min_{n \in \mathcal{N} : BR_n(\sigma_{-n}) \neq S_n} \left(\max_{s_n \in S_n} \pi_n(s_n, \sigma_{-n}) - \max_{s_n \in S_n \setminus BR_n(\sigma_{-n})} \pi_n(s_n, \sigma_{-n}) \right).$$

Fix some positive $c < \min\{\bar{c}, \underline{\Delta}\}$. Let b be a quasi*-perfect equilibrium of $\bar{\Gamma}(c)$ such that $\sigma = \alpha(b)$. Let $(\tau^n)_{n=1}^N$, $(\{d^{t,n}\}_{t=1}^\infty)_{n=1}^N$, $\{\varepsilon_t\}_{t=1}^\infty$, and $\{\delta_t\}_{t=1}^\infty$ be as in Definition 7.

We prove that σ is a test-set equilibrium through a series of claims. Claim 1 establishes that σ is a Nash equilibrium in undominated strategies. Claim 6 establishes that σ satisfies the test-set condition. Claims 2–5 are intermediate results used in the proof of Claim 6. Specifically, Claim 2 establishes that if an information set of a particular player is reached sufficiently frequently, then the reaching of that information set conveys no information, asymptotically, about the behavior of the other players. Claim 3 establishes that approval must occur at review information sets following recommendations that are best responses to σ, and conversely, disapproval at review information sets following recommendations that are inferior responses to σ. Claims 4 and 5 establish that local strategies at certain final decision information sets, interpreted as strategies in Γ, must be best responses to σ.

Claim 1: Given any player m, let $v \in U_m$ be the recommendation information set in $\bar{\Gamma}(c)$ of that player. Interpreted as a strategy in Γ, the local strategy b_{mv} is equivalent to σ_m.[19] Moreover, σ_m is undominated in Γ and $\sigma_m \in BR_m(\sigma_{-m})$.

Proof: We abuse notation to let α_{-m} map a (potentially non-product) distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$ into the induced distribution on S_{-n}. Let $\pi_m(\sigma_m, \alpha_{-m}(d^{t,m}))$ denote the payoff received in Γ by player m when the distribution of play is determined by the product of σ_m and $\alpha_{-m}(d^{t,m})$.

We have $\bar{\pi}_{mv}(d^{t,m}/_v b'_m) = \pi_m(\alpha_m(b'_m), \alpha_{-m}(d^{t,m})) - \beta_m(b'_m)c$. The fact that $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ for all t implies (i) $\beta_m(b_m) = 0$, and (ii) $\alpha_m(b_m) \in BR_m(\alpha_{-m}(d^{t,m}))$ for all t. By (i), player m's recommendation is never rejected on path, and so $\sigma_m = \alpha_m(b_m)$ is equivalent to b_{mv} when the latter is interpreted as a strategy in Γ.

[19] The local strategy b_{mv} can be interpreted as a mixed strategy in Γ for player m because the available choices at player m's recommendation information set v are exactly the pure strategies S_m.

Moreover, by (ii), $\sigma_m = \alpha_m(b_m) \in BR_m(\alpha_{-m}(d^{t,m}))$ for all t. For all t, $d^{t,m}$ is completely mixed, and therefore $\alpha_{-m}(d^{t,m})$ is a full support distribution on S_{-m}. This immediately implies that σ_m is undominated in Γ. Furthermore, since $d^{t,m}$ converges to the distribution induced by the behavior strategy profile b, we have that $\alpha_{-m}(d^{t,m})$ converges to the distribution induced by σ_{-m}. Taking limits, we conclude $\sigma_m \in BR_m(\sigma_{-m})$.

Claim 2: Given any player m and any information set $v \in U_m$ in $\bar{\Gamma}(c)$, let $d^{t,m}|v$ denote the conditional distribution of $d^{t,m}$ given that v is reached. If the unconditional probability of v being reached under $d^{t,m}$ is $\Omega(\varepsilon_t)$, then $\alpha_{-m}(d^{t,m}|v)$ converges to the distribution induced by σ_{-m}.[20]

Proof: The probability under $d^{t,m}$ that there is a correlated tremble and v is reached is $O(\varepsilon_t\delta_t)$. By assumption, the probability under $d^{t,m}$ that v is reached is $\Omega(\varepsilon_t)$. Consequently, the probability under $d^{t,m}$ that there is no correlated tremble and v is reached is also $\Omega(\varepsilon_t)$. By Bayes' Rule, $d^{t,m}|v$ therefore converges to the conditional distribution of $d^{t,m}$ given that v is reached and there is no correlated tremble. The marginal of that conditional distribution on $\prod_{n\neq m}\prod_{u\in U_n} C_u$ is the distribution induced by b_{-m}. Thus, $\alpha_{-m}(d^{t,m}|v)$ converges to the distribution induced by $\alpha_{-m}(b_{-m}) = \sigma_{-m}$.

Claim 3: Given any player m and any $s'_m \in S_m$, let $v \in U_m$ denote the review information set in $\bar{\Gamma}(c)$ that follows a recommendation of s'_m. If $s'_m \in BR_m(\sigma_{-m})$, then the local strategy b_{mv} specifies approval with probability one. If $s'_m \notin BR_m(\sigma_{-m})$, then the local strategy b_{mv} specifies disapproval with probability one.

Proof: If b_m specifies recommending s'_m with positive probability, then information set v is reached with probability $\Theta(1)$ under $d^{t,m}$; otherwise v is reached with probability $\Theta(\varepsilon_t)$. Let $d^{t,m}|v$ denote the conditional distribution of $d^{t,m}$ given that v is reached. By Claim 2, $\alpha_{-m}(d^{t,m}|v)$ converges to the distribution induced by σ_{-m}. Suppose b'_m is a behavior strategy that specifies approval with probability p at information set v and specifies deciding according to σ''_m at the final decision information set following v. Then $\bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ converges to $p\pi_m(s'_m, \sigma_{-m}) + (1-p)\left[\pi_m(\sigma''_m, \sigma_{-m}) - c\right]$. We use this general form in the payoff calculations below.

First, suppose $s'_m \in BR_m(\sigma_{-m})$. Let b'_m be a behavior strategy that specifies approval with probability 1 at information set v. Let b''_m be a behavior strategy that specifies approval with probability $p < 1$ at information set v and specifies deciding according to some σ''_m at the final decision information set following v. Then

$$\begin{aligned}&\lim_{t\to\infty} \bar{\pi}_{mv}(d^{t,m}/_v b'_m) - \bar{\pi}_{mv}(d^{t,m}/_v b''_m)\\ &\quad = (1-p)\left[\pi_m(s'_m, \sigma_{-m}) - \pi_m(\sigma''_m, \sigma_{-m}) + c\right] \geq (1-p)c > 0,\end{aligned}$$

which implies that $\bar{\pi}_{mv}(d^{t,m}/_v b'_m) > \bar{\pi}_{mv}(d^{t,m}/_v b''_m)$ for sufficiently large values of t. Since $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ for all t, this implies that b_{mv} must specify approval with probability one if $s'_m \in BR_m(\sigma_{-m})$.

Second, suppose $s'_m \notin BR_m(\sigma_{-m})$. Let b'_m be a behavior strategy that specifies approval with probability 0 at information set v and specifies deciding according to some $\sigma'_m \in$

[20] A sequence $\{x_t\}_{t=1}^{\infty}$ is $\Omega(\varepsilon_t)$ if $(\exists k > 0)(\exists T)(\forall t > T)|x_t| \geq k\varepsilon_t$. Similarly, a sequence $\{x_t\}_{t=1}^{\infty}$ is $O(\varepsilon_t)$ if $(\exists k > 0)(\exists T)(\forall t > T)|x_t| \leq k\varepsilon_t$. In addition a sequence $\{x_t\}_{t=1}^{\infty}$ is $\Theta(\varepsilon_t)$ if it is both $\Omega(\varepsilon_t)$ and $O(\varepsilon_t)$.

$BR_m(\sigma_{-m})$ at the final decision information set following v. Let b''_m be a behavior strategy that specifies approval with probability $p > 0$ at information set v and specifies deciding according to some σ''_m at the final decision information set following v. Then

$$\lim_{t\to\infty} \bar{\pi}_{mv}(d^{t,m}/_v b'_m) - \bar{\pi}_{mv}(d^{t,m}/_v b''_m)$$
$$= \pi_m(\sigma'_m, \sigma_{-m}) - p\pi_m(s'_m, \sigma_{-m}) - (1-p)\pi_m(\sigma''_m, \sigma_{-m}) - pc \geq p\underline{\Delta} - pc > 0,$$

which implies that $\bar{\pi}_{mv}(d^{t,m}/_v b'_m) > \bar{\pi}_{mv}(d^{t,m}/_v b''_m)$ for sufficiently large values of t. Since $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ for all t, this implies that b_{mv} must specify disapproval with probability one if $s'_m \notin BR_m(\sigma_{-m})$.

Claim 4: Given any player m and any $s'_m \notin BR_m(\sigma_{-m})$, let $v \in U_m$ be the final decision information set in $\bar{\Gamma}(c)$ following the disapproval of a recommendation of s'_m. Interpreted as a strategy in Γ, $b_{mv} \in BR_m(\sigma_{-m})$.

Proof: By Claim 1, reaching information set v requires a tremble at the recommendation information set of player m. However, by Claim 3, reaching information set v does not require a tremble at the following review information set. Consequently, information set v is reached with probability $\Theta(\varepsilon_t)$ under $d^{t,m}$. Let $d^{t,m}|v$ denote the conditional distribution of $d^{t,m}$ given that v is reached. By Claim 2, $\alpha_{-m}(d^{t,m}|v)$ converges to the distribution induced by σ_{-m}. Then, interpreting b_{mv} as a strategy in Γ, $\bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ converges to $\pi_m(b_{mv}, \sigma_{-m}) - c$. We must have $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$. Taking limits, this requires $b_{mv} \in BR_m(\sigma_{-m})$.

Claim 5: Given any player m and any $s'_m \in S_m$ that is recommended under b_m with positive probability, let $v \in U_m$ denote the final decision information set in $\bar{\Gamma}(c)$ following the disapproval of a recommendation of s'_m. Interpreted as a strategy in Γ, $b_{mv} \in BR_m(\sigma_{-m})$.

Proof: By assumption, reaching information set v does not require a tremble at the recommendation information set of player m, which by Claim 1 implies that $s'_m \in BR_m(\sigma_{-m})$. However, by Claim 3, reaching information set v therefore requires a tremble at the following review information set. Consequently, information set v is reached with probability $\Theta(\varepsilon_t)$ under $d^{t,m}$. Let $d^{t,m}|v$ denote the conditional distribution of $d^{t,m}$ given that v is reached. By Claim 2, $\alpha_{-m}(d^{t,m}|v)$ converges to the distribution induced by σ_{-m}. As in the proof of Claim 4, this requires $b_{mv} \in BR_m(\sigma_{-m})$.

Claim 6: For any player m, σ_m is undominated in $T(\sigma)$.

Proof: Letting $L = \sum_{n=1}^{N} |U_n|$, the following statements are true of the distribution of play in $\bar{\Gamma}(c)$ under $d^{t,m}$. First, the probability of no tremble occurring is $(1-\varepsilon_t\delta_t)(1-\varepsilon_t)^L$, in which case play is distributed according to b. Second, for any player n, any information set $u \in U_n$, and any $c_u \in C_u$, the probability of a "single independent tremble" to c_u is $(1-\varepsilon_t\delta_t)\varepsilon_t(1-\varepsilon_t)^{L-1}\tau^m_{nu}(c_u)$, in which case play is distributed according to b/c_u. Third, the total probability of "joint independent trembles" is $(1-\varepsilon_t\delta_t)\left[1-(1-\varepsilon_t)^L - L\varepsilon_t(1-\varepsilon_t)^{L-1}\right]$, in which cases τ^m dictates the play at two or more information sets while play at the other information sets is distributed according to b. Fourth, the probability of a "correlated tremble" is $\varepsilon_t\delta_t$, in which case play may be distributed arbitrarily.

We next argue that every "single independent tremble" leads to play in the test set. That is, for any player n, any information set $u \in U_n$, and any $c_u \in C_u$, the strategy profile $\alpha_{-m}(b/c_u)$ is a convex combination of strategy profiles in the set $T_m(\sigma) =$

$\cup_{n \neq m} \{(s_n, \sigma_{-mn}) : s_n \in BR_n(\sigma_{-n})\}$. First, suppose $n = m$. Then the tremble is irrelevant, so that $\alpha_{-m}(b/c_u) = \sigma_{-m}$. Second, suppose u is the recommendation information set of a player $n \neq m$ and c_u specifies recommending some $s_n \in BR_n(\sigma_{-n})$. By Claim 3, b_n specifies approving it with probability one. Hence, $\alpha_{-m}(b/c_u) = (s_n, \sigma_{-nm})$. Third, suppose u is the recommendation information set of a player $n \neq m$ and c_u specifies recommending some $s_n \notin BR_n(\sigma_{-n})$. By Claim 3, b_n specifies disapproving it with probability one, and by Claim 4, b_n specifies deciding according to some $\sigma'_n \in BR_n(\sigma_{-n})$ at the subsequent information set. Hence, $\alpha_{-m}(b/c_u) = (\sigma'_n, \sigma_{-nm})$. Fourth, suppose u is some review information set of a player $n \neq m$ that follows a recommendation that is not made with positive probability under b_n. Then the tremble is irrelevant, so $\alpha_{-m}(b/c_u) = \sigma_{-m}$. Fifth, suppose u is some review information set of a player $n \neq m$ that follows a recommendation that is made with positive probability under b_n, and c_u specifies approving the recommendation. By Claim 1, $\alpha_{-m}(b/c_u) = \sigma_{-m}$. Sixth, suppose u is some review information set of a player $n \neq m$ that follows a recommendation that is made with positive probability under b_n, and c_u specifies disapproving the recommendation. By Claim 5, b_n specifies deciding according to some $\sigma'_n \in BR_n(\sigma_{-n})$ at the subsequent final decision information set. The induced play for player n is then some σ''_n whose support is a subset of $\text{supp}(\sigma_n) \cup \text{supp}(\sigma'_n)$. Thus, $\sigma''_n \in BR_n(\sigma_{-n})$, and $\alpha_{-m}(b/c_u) = (\sigma''_n, \sigma_{-nm})$. Seventh, suppose u is some final decision information set of a player $n \neq m$. This means that there are no trembles elsewhere in the game tree. By Claims 1 and 3, the final decision information set will then remain unreached, so the tremble is irrelevant and $\alpha_{-m}(b/c_u) = \sigma_{-m}$.

Suppose there exists some strategy $\hat{\sigma}_m \in \Delta(S_m)$ that weakly dominates σ_m in the test set. We will derive a contradiction in the following way. When the distribution of play in $\bar{\Gamma}(c)$ is governed by $d^{t,m}$, but no trembles occur, then the induced play in Γ is distributed according to σ_{-m}, against which $\hat{\sigma}_m$ and σ_m perform equally well. When a "single independent tremble" occurs, then, as argued above, the induced play in Γ lies in the test set, where $\hat{\sigma}_m$ performs weakly better than σ_m and sometimes strictly so. When a "joint independent tremble" or a "correlated tremble" occurs, the induced play in Γ may lie outside the test set, and $\hat{\sigma}_m$ may perform strictly worse, but only by an amount that is bounded. Aggregating across these possibilities, we will see $\hat{\sigma}_m$ performs strictly better overall as $t \to \infty$.

Since $\hat{\sigma}_m \in \Delta(S_m)$ weakly dominates σ_m in the test set, there exists some $\hat{\sigma}_{-m} \in T_m(\sigma)$ such that $\pi_m(\hat{\sigma}_m, \hat{\sigma}_{-m}) > \pi_m(\sigma_m, \hat{\sigma}_{-m})$. Further, there then exists a player $n \neq m$ such that $\hat{\sigma}_{-m}$ takes the form $(\hat{s}_n, \sigma_{-nm})$ for some $\hat{s}_n \in BR_n(\sigma_{-n})$. Let u be the recommendation information set of player n, and let $\hat{c}_u$ specify recommending $\hat{s}_n \in BR_n(\sigma_{-n})$. By Claim 3, b_n specifies approval with probability one at the review information set following a recommendation of $\hat{s}_n$. Hence, $\alpha_{-m}(b/\hat{c}_u) = \hat{\sigma}_{-m}$. In addition, define $\bar{\Delta}$ as the maximum difference between the payoffs in Γ from best and worst responses to any strategy profile:

$$\bar{\Delta} = \max_{n \in \mathcal{N}} \max_{s_{-n} \in S_{-n}} \left(\max_{s_n \in S_n} \pi_n(s_n, s_{-n}) - \min_{s_n \in S_n} \pi_n(s_n, s_{-n}) \right).$$

Combining all of the above, the difference in player m's expected payoff from $\hat{\sigma}_m$ and from

σ_m in Γ when the play of other players is distributed according to $\alpha_{-m}(d^{t,m})$ is

$$\begin{aligned}
&\pi_m(\hat{\sigma}_m, \alpha_{-m}(d^{t,m})) - \pi_m(\sigma_m, \alpha_{-m}(d^{t,m})) \\
&\quad \geq [\pi_m(\hat{\sigma}_m, \hat{\sigma}_{-m}) - \pi_m(\sigma_m, \hat{\sigma}_{-m})] (1-\varepsilon_t\delta_t)\varepsilon_t(1-\varepsilon_t)^{L-1}\tau^m_{nu}(\hat{c}_u) \\
&\qquad - \bar{\Delta}\left(\varepsilon_t\delta_t + (1-\varepsilon_t\delta_t)\left[1-(1-\varepsilon_t)^L - L(1-\varepsilon_t\delta_t)^L(1-\varepsilon_t)^{L-1}\varepsilon_t\right]\right).
\end{aligned}$$

The first term is positive and on the order of ε_t. The second term is negative, but only on the order of $\varepsilon_t^2 + \varepsilon_t\delta_t$. Thus, this difference is positive for sufficiently large values of t. This creates the desired contradiction because, as established in the proof of Claim 1, we must have $\sigma_m \in BR_m(\alpha_{-m}(d^{t,m}))$ for all t. □

Proof of Theorem 10 (Necessity). Let σ be a Nash equilibrium of Γ in undominated strategies that satisfies the test-set condition. Define the behavior strategy profile b as follows: (*i*) for any player n, if $u \in U_n$ is the recommendation information set of that player, then let b_{nu} be equivalent to σ_n; (*ii*) for any player n, if $u \in U_n$ is a review information set of that player that follows a recommendation to play some $s'_n \in BR_n(\sigma_{-n})$, then let b_{nu} specify approval; (*iii*) for any player n, if $u \in U_n$ is a review information set of that player that follows a recommendation to play some $s'_n \notin BR_n(\sigma_{-n})$, then let b_{nu} specify disapproval; (*iv*) for any player n, if $u \in U_n$ is a final decision information set of that player, then let b_{nu} be equivalent to σ_n. Because σ is a Nash equilibrium, for every player n, only best responses to σ_{-n} are recommended with positive probability. Consequently, disapproval never occurs on path, and hence $\alpha(b) = \sigma$.

If, for all players n, $BR_n(\sigma_{-n}) = S_n$, then let $\underline{\Delta}$ be any positive number. Otherwise, define $\underline{\Delta} > 0$ as the minimum difference between the payoffs in Γ from best and inferior responses to σ:

$$\underline{\Delta} = \min_{n\in\mathcal{N}:BR_n(\sigma_{-n})\neq S_n}\left(\max_{s_n\in S_n}\pi_n(s_n,\sigma_{-n}) - \max_{s_n\in S_n\setminus BR_n(\sigma_{-n})}\pi_n(s_n,\sigma_{-n})\right).$$

Fix any positive $c < \underline{\Delta}$. To complete the proof of sufficiency—with $\underline{\Delta}$ playing the role of $\bar{c}$—we argue through a series of claims that b, as defined above, is a quasi*-perfect equilibrium of $\bar{\Gamma}(c)$.

Claim 1: For any player m, there exists a full support probability distribution on S_{-m}, which we denote g^m, such that $\sigma_m \in BR_m(g^m)$.

Proof: The proof is standard. Let $L = \sum_{n\neq m}|S_n|$. Let $(s^1_{-m}, \ldots, s^L_{-m})$ be an enumeration of the elements of S_{-m}. For each $\sigma_m \in \Delta(S_m)$, define

$$x_m(\sigma_m) = \left(\pi_m(\sigma_m, s^1_{-m}), \ldots, \pi_m(\sigma_m, s^L_{-m})\right),$$

and define the set

$$X_m = \mathrm{Conv}\left(\{x_m(s_m) : s_m \in S_m\}\right).$$

Then X_m is closed, convex, and determined by a finite set. Because σ_m is not weakly dominated in Γ, the result of Arrow, Barankin and Blackwell (1953) implies the existence of a supporting hyperplane to the set X_m through the point $x_m(\sigma_m)$ whose normal has positive components. Dividing the components by their sum yields a normal vector whose components sum to one. The normalized components constitute the desired probability distribution g^m, which has full support on S_{-m}.

Claim 2: For any player m, define $T_m(\sigma) = \cup_{n \neq m} \{(s_n, \sigma_{-mn}) : s_n \in BR_n(\sigma_{-n})\}$. There exists a full support probability distribution on $T_m(\sigma)$, which we denote f^m, such that $\sigma_m \in BR_m(f^m)$.

Proof: The proof is similar to that of Claim 1. In this case, let $L = |T_m(\sigma)|$, and let $(\sigma^1_{-m}, \ldots, \sigma^L_{-m})$ be an enumeration of the elements of $T_m(\sigma)$. Additionally, for each $\sigma_m \in \Delta(S_m)$, define

$$x_m(\sigma_m) = \left(\pi_m(\sigma_m, \sigma^1_{-m}), \ldots, \pi_m(\sigma_m, \sigma^L_{-m})\right),$$

and define the set

$$X_m = \mathrm{Conv}\left(\{x_m(s_m) : s_m \in S_m\}\right).$$

Then X_m is closed, convex, and determined by a finite set. Because σ_m is not weakly dominated in $T_m(\sigma)$, the result of Arrow, Barankin and Blackwell (1953) implies the existence of a supporting hyperplane to the set X_m through the point $x_m(\sigma_m)$ whose normal has positive components. Dividing the components by their sum yields a normal vector whose components sum to one. The normalized components constitute the desired probability distribution f^m, which has full support on $T_m(\sigma)$.

Claim 3: For any player m and any full support distribution on S_{-m}, g^m, there exists a full support distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$, which we denote $\bar{g}^m$, such that $\alpha_{-m}(\bar{g}^m) = g^m$.

Proof: For each $s'_{-m} \in S_{-m}$, fix a full support distribution on $\{c_{-m} \in \prod_{n \neq m} \prod_{u \in U_n} C_u : \alpha_{-m}(c_{-m}) = s'_{-m}\}$. We refer to that distribution as "the vertex corresponding to s'_{-m}." Let Δ^* be the subset of distributions over $\prod_{n \neq m} \prod_{u \in U_n} C_u$ that is the $|S_{-m}|$-dimensional simplex defined by those vertices.

Abusing notation to let α_{-m} map a distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$ into the induced distribution on S_{-m}, notice that $\alpha_{-m} : \Delta^* \to \Delta(S_{-m})$ is a continuous function. Moreover, α_{-m} maps the vertex corresponding to s'_{-m} into the pure strategy profile s'_{-m}. Conversely, α_{-m} maps any point on the face of the simplex opposite the vertex corresponding to s'_{-m} to a distribution that puts zero probability on s'_{-m}. Finally, since g^m has full support on S_{-m}, $g^m(s'_{-m}) \in (0, 1)$. Thus, we can apply the generalization of the intermediate value theorem for simplices, due to Vrahatis (2016), to establish the existence of some $\bar{g}^m \in \mathrm{int}(\Delta^*)$ with $\alpha_{-m}(\bar{g}^m) = g^m$. By the construction of Δ^*, any such interior point is a full support distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$.

Claim 4: Let $L = \sum_{n \neq m} |U_n|$. For any player m and any full support distribution on the elements of $T_m(\sigma)$, f^m, there exists a completely mixed behavior strategy profile τ^m_{-m} such that

$$\alpha_{-m}\left(\frac{1}{L} \sum_{n \neq m} \sum_{u \in U_n} \sum_{c_u \in C_u} \tau^m_{nu}(c_u) \cdot b_{-m}/(b_n/c_u)\right) = \frac{2L-1}{2L} \sigma_{-m} + \frac{1}{2L} f^m.$$

Proof: We proceed by construction. First, for all players $n \neq m$ and all final decision information sets $u \in U_n$, let τ^m_{nu} be any local strategy with full support. Second, fix some probability $p \in (0, 1)$. Then for all players $n \neq m$ and all review information sets $u \in U_n$, let τ^m_{nu} be the local strategy that approves with probability p and disapproves with probability $1 - p$. Third, for all players $n \neq m$, at the recommendation information set $u \in U_n$, determine the local strategy τ^m_{nu} in the following way: (i) if $s'_n \in BR_n(\sigma_{-n})$, then let τ^m_{nu} tentatively assign one half of the probability to s'_n that f^m assigns to (s'_n, σ_{-nm}); (ii) if

$S_n \setminus BR_n(\sigma_{-n})$ is nonempty, then distribute the remaining probability in proportion to any full support distribution on the elements of that set; and (*iii*) if $S_n \setminus BR_n(\sigma_{-n})$ is empty, then distribute the remaining probability in proportion to σ_n. Because f^m has full support on $T_m(\sigma)$, the resulting behavior strategy profile τ^m_{-m} is completely mixed.

Letting τ^m_{-m} be as defined above, it can be checked that the distribution on pure behavior strategies in $\bar{\Gamma}$ given by

$$\frac{1}{L} \sum_{n \neq m} \sum_{u \in U_n} \sum_{c_u \in C_u} \tau^m_{nu}(c_u) \cdot b_{-m}/(b_n/c_u)$$

induces play in Γ that is distributed according to $\frac{2L-1}{2L}\sigma_{-m} + \frac{1}{2L} f^m$, as claimed.

Claim 5: b is a quasi*-perfect equilibrium of $\bar{\Gamma}(c)$.

Proof: Define $\{\varepsilon_t\}_{t=1}^\infty = \frac{1}{t^2}$ and $\{\delta_t\}_{t=1}^\infty = \frac{1}{t}$. For any player m, there exist distributions g^m, f^m, and $\bar{g}^m$, as well as a behavior strategy profile τ^m_{-m}, which satisfy the conditions of the previous claims. It will economize on notation to define the following distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$:

$$\bar{f}^m = \frac{1}{L} \sum_{n \neq m} \sum_{u \in U_n} \sum_{c_u \in C_u} \tau^m_{nu}(c_u) \cdot b_{-m}/(b_n/c_u),$$

where, by Claim 4, $\alpha_{-m}(\bar{f}^m) = \frac{2L-1}{2L}\sigma_{-m} + \frac{1}{2L} f^m$. We can then define the following distribution on $\prod_{n \neq m} \prod_{u \in U_n} C_u$:

$$\begin{aligned} d^{t,m}_{-m}(c_{-m}) = (1-\varepsilon_t\delta_t)(1-\varepsilon_t)^L b_{-m}(c_{-m}) + L(1-\varepsilon_t\delta_t)\varepsilon_t(1-\varepsilon_t)^{L-1}\bar{f}^m(c_{-m}) \\ + \left[1 - (1-\varepsilon_t\delta_t)(1-\varepsilon_t)^L - L(1-\varepsilon_t\delta_t)\varepsilon_t(1-\varepsilon_t)^{L-1}\right]\bar{g}^m(c_{-m}). \end{aligned}$$

In addition, let τ^m_m be any full support behavior strategy for player m, and define the following distribution on $\prod_{n=1}^N \prod_{u \in U_n} C_u$:

$$d^{t,m}(c) = d^{t,m}_{-m}(c_{-m}) \cdot \prod_{u \in U_m} \left[(1-\varepsilon_t) b_{mu}(c_u) + \varepsilon_t \tau^m_{mu}(c_u)\right].$$

By a simple change of variables, it suffices to show that $d^{t,m}$ meets the requirements of parts (ii) and (iii) of Definition 7 for all *sufficiently large* values of t, rather than for all t.

We first establish part (ii) of Definition 7, namely that for all sufficiently large values of t and for all $c \in \prod_{n \neq m} \prod_{u \in U_n} C_u$,

$$d^{t,m}(c) \geq (1-\varepsilon_t\delta_t) \prod_{n=1}^N \prod_{u \in U_n} \left[(1-\varepsilon_t) b_{nu}(c_u) + \varepsilon_t \tau^m_{nu}(c_u)\right].$$

To see that this is true, note that after dividing through by $\prod_{u \in U_m} [(1-\varepsilon_t) b_{mu}(c_u) + \varepsilon_t \tau^m_{mu}(c_u)]$ and then cancelling like terms, the left hand side is $\Theta(\varepsilon_t\delta_t) = \Theta(t^{-3})$, while the right hand side is $\Theta(\varepsilon_t^2) = \Theta(t^{-4})$.

Next, we note that by the construction of $d^{t,m}$ as the product of $d^{t,m}_{-m}$ and a separate distribution on the elements of $\prod_{u \in U_m} C_u$, the distribution of play by the players $n \neq m$ in $\bar{\Gamma}$, conditional on reaching any information set $v \in U_m$, is equivalent to the unconditional distribution, $d^{t,m}_{-m}$. Moreover, that distribution is a convex combination of b_{-m}, $\bar{f}^m$, and $\bar{g}^m$, which converges to b_{-m} as t grows large. Consequently, this induces play in Γ that is a convex combination of σ_{-m}, f^m, and g^m, which converges to σ_{-m} as t grows large. Moreover, note that σ_m is a best response to all three of those distributions, and therefore to any such convex combination. Using these observations, we establish part (iii) of Definition 7, namely that for all sufficiently large values of t and for all $v \in U_m$,

$$\bar{\pi}_{mv}(d^{t,m}/_v b_m) = \max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m).$$

To do so, we separately consider the recommendation information set, the review information sets, and the final decision information sets.

First, let v be the recommendation information set of player m. As argued above, σ_m is a best response to the distribution of play in Γ by the players $n \neq m$ for all t. Note that b_m specifies recommending according to σ_m at the recommendation information set v, and also specifies approval at all on-path review information sets. This is the best that player m can do, and so we indeed have $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$.

Second, let v be a final decision information set of player m. As argued above, σ_m is a best response to the distribution of play in Γ by the players $n \neq m$ for all t. Note that b_m specifies deciding according to σ_m at v. This is the best that player m can do conditional on v being reached, and so we indeed have $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$.

Third, let v be a review information set of player m. As argued above, σ_m is a best response to the distribution of play in Γ by the players $n \neq m$ for all t. Thus, conditional on information set v being reached, player m can do no better than the best of the following possibilities: (i) to disapprove the recommendation at information set v and to play according to σ_m at the subsequent final decision information set, or (ii) to approve the recommendation at information set v. Suppose first that v follows a recommendation to play some $s'_m \in BR_m(\sigma_{-m})$. In this case, b_m specifies approval at v, and so $\bar{\pi}_{mv}(d^{t,m}/_v b_m)$ converges to $\pi_m(s'_m, \sigma_{-m})$, which, since s'_m and σ_m are both best responses to σ_{-m} is equal to $\pi_m(\sigma'_m, \sigma_{-m})$. It suffices to compare this to an alternative, b'_m, which specifies disapproval at v and deciding according to σ_m at the subsequent information set, and so $\bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ converges to $\pi_m(\sigma_m, \sigma_{-m}) - c$. Thus, $\lim_{t\to\infty} \bar{\pi}_{mv}(d^{t,m}/_v b_m) - \bar{\pi}_{mv}(d^{t,m}/_v b'_m) = c > 0$. Consequently, we have $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ for all sufficiently large values of t. Suppose next that v follows a recommendation to play some $s'_m \notin BR_m(\sigma_{-m})$. In this case, b_m specifies disapproval at v and deciding according to σ_m at the subsequent information set, and so $\bar{\pi}_{mv}(d^{t,m}/_v b_m)$ converges to $\pi_m(\sigma_m, \sigma_{-m}) - c$. It suffices to compare this to an alternative, b'_m, which specifies approval at v, and so $\bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ converges to $\pi_m(s'_m, \sigma_{-m})$. Thus, $\lim_{t\to\infty} \bar{\pi}_{mv}(d^{t,m}/_v b_m) - \bar{\pi}_{mv}(d^{t,m}/_v b'_m) = \pi_m(\sigma_m, \sigma_{-m}) - c - \pi_m(s'_m, \sigma_{-m}) \geq \underline{\Delta} - c > 0$. Consequently, we have $b_m \in \arg\max_{b'_m \in B_m} \bar{\pi}_{mv}(d^{t,m}/_v b'_m)$ for all sufficiently large values of t. $\square$

Supplement to Equilibrium Selection in Auctions and High Stakes Games*

Paul Milgrom Joshua Mollner

September 13, 2017

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 an allocation $x \in X$ 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

*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 number 1525730.

[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 $\bigtimes_{n=1}^{N} X_n$.

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.

While Bernheim and Whinston (1986) demonstrate 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 allocation 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 allocation (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 allocation (w, l, l, l, l, l). This allocation is inefficient because it yields a total surplus of 29, whereas the allocation (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.

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.

Finally, in the equilibrium, no bidder is using a bid that is weakly dominated in the game, by an extension of Lemma 2. □

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, Athey, Babaioff and Grubb (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 $\{0, \frac{1}{m}, \ldots, \frac{m-1}{m}, 1\}$.

Part One: For all sufficiently small values of ε, the following completely mixed strategy profile is an ε-proper equilibrium: (i) for $k \in \{1, \ldots, m\}$, the low-type informed bidder places probability ε^k on k/m and places all remaining probability on 0, (ii) for $k \in \{0, \ldots, m-1\}$, the high-type informed bidder places probability ε^{2m-k} on k/m and places all remaining probability on 1, and (iii) for $k \in \{1, \ldots, m\}$, 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 $k \in \{1, \ldots, m\}$, the low-type informed bidder places probability

ε^{m+k} on k/m and places all remaining probability on 0, (ii) for $k \in \{0, \ldots, m-1\}$, the high-type informed bidder places probability ε^{m-k} on k/m and places all remaining probability on 1, and (iii) for $k \in \{0, \ldots, m-1\}$, 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 $b, b' \in \{0, \frac{1}{m}, \ldots, \frac{m-1}{m}, 1\}$ 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 $\sigma_0^t(b') \leq \varepsilon_t \sigma_0^t(b)$. 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 $\sigma_1^t(b) \leq \varepsilon_t \sigma_1^t(b')$.

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

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

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 $k \in \{2, \ldots, m-1\}$. Suppose to the contrary that it were. This implies $\pi_U^t\left(\frac{k}{m}\right) \geq \pi_U^t\left(\frac{k+1}{m}\right)$, or equivalently,

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

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

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

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

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

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

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

Together, (12) and (13) further imply

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

which is a contradiction for sufficiently large indices t. $\square$

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.

Claim (i): This follows immediately from Proposition 1 and the existence of proper equilibria in finite games.

Claim (ii): 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 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 (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 argument now 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.

- Suppose $BR_3(\sigma_1, \sigma_2) = \{a_3, b_3\}$. 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:

$$\begin{aligned}\pi_1(\hat{\sigma}_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) - \pi_1(\sigma_1, \sigma_2^\varepsilon, \sigma_3^\varepsilon) = {} & [\pi_1(\hat{\sigma}_1, a_2, \sigma_3^\varepsilon) - \pi_1(\sigma_1, a_2, \sigma_3^\varepsilon)]\, \sigma_2^\varepsilon(a_2) \\ & + [\pi_1(\hat{\sigma}_1, b_2, \sigma_3^\varepsilon) - \pi_1(\sigma_1, b_2, \sigma_3^\varepsilon)]\, \sigma_2^\varepsilon(b_2),\end{aligned}$$

  which is positive for sufficiently small values of ε. To see the last step:

  (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.

  This constitutes a contradiction to σ being an extended proper equilibrium.

- 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:

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

  which is positive for sufficiently small values of ε. To see the last step:

  (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 $\sigma_3 = a_3$, $\frac{\sigma_2^\varepsilon(b_2)}{\sigma_2^\varepsilon(a_2)}$ 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 one, this implies that the fourth term inside the braces converges to zero as well.

This constitutes a contradiction to σ being an extended proper equilibrium.

Claim (iii): 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 $s \in \bigtimes_{n=1}^N S_n$, $\pi_n(s) = P(s)$. It is known that any finite potential game possesses a pure strategy trembling-hand perfect equilibrium $s^* \in \arg\max_{s \in \bigtimes_{n=1}^N S_n} P(s)$ (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. □

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) establishes that, for almost all finite games in normal form, every Nash equilibrium is quasi-strict.[4] Similarly, Kreps and Wilson (1982) establish 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.

First, in such games, every Nash equilibrium is a trembling-hand perfect equilibrium, and therefore an equilibrium in undominated strategies. Second, in such games, every Nash equilibrium is 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. □

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

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

References

Abraham, Ittai, Susan Athey, Moshe Babaioff, and Michael Grubb, "Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information," *Working Paper*, 2016.

Bernheim, B. Douglas and Michael D. Whinston, "Menu Auctions, Resource Allocation, and Economic Influence," *The Quarterly Journal of Economics*, 1986, *101* (1), 1–32.

Carbonell-Nicolau, Oriol and Richard P. McLean, "Refinements of Nash Equilibrium in Potential Games," *Theoretical Economics*, 2014, *9* (3), 555–582.

Day, Robert and Paul Milgrom, "Core-Selecting Package Auctions," *International Journal of Game Theory*, 2008, *36* (3), 393–407.

Harsanyi, John C., "Oddness of the Number of Equilibrium Points: A New Proof," *International Journal of Game Theory*, 1973, *2* (1), 235–250.

Kreps, David M. and Robert Wilson, "Sequential Equilibria," *Econometrica*, 1982, *50* (4), 863–894.

Milgrom, Paul and Joshua Mollner, "Extended Proper Equilibrium," *Working Paper*, 2017. http://ssrn.com/abstract=3035565.

Monderer, Dov and Lloyd S. Shapley, "Potential Games," *Games and Economic Behavior*, 1996, *14* (1), 124–143.