Extended Proper Equilibrium*

Paul Milgrom† Joshua Mollner‡

August 20, 2020

Abstract

We introduce *extended proper equilibrium*, which refines proper equilibrium (Myerson, 1978) by adding across-player restrictions on trembles. This refinement coincides with proper equilibrium in games with two players but adds new restrictions in games with three or more players. One implication of these additional restrictions is that any tremble that is costless in equilibrium is regarded by all as more likely than any costly tremble, even one by a different player. At least one extended proper equilibrium exists in every finite game. The refinement can also be characterized in terms of a symmetric, meta-version of the game in which players originate from a common pool: if these players tremble symmetrically and in the way of proper equilibrium, then the induced play in the original game is an extended proper equilibrium.

Keywords: equilibrium refinement, trembles, extended proper equilibrium, proper equilibrium

1 Introduction

This paper introduces a refinement of proper equilibrium. Intuitively, a proper equilibrium is a trembling-hand perfect equilibrium in which each player is much less likely to make a more costly mistake than a less costly one. Proper equilibrium, however, still permits the possibility that one player has a fundamentally greater propensity to tremble than another, in the sense that each of the former's mistakes (no matter how costly) are more likely than each of the latter's mistakes (even costless "mistakes").

To illustrate this possibility, consider the three-player game in Figure 1, in which each player has two strategies. The "Geo" player picks the payoff matrix—*East* or *West*. For Row, the strategy *Up* weakly dominates *Down*: the latter pays either zero or one, while the former always pays one. Similarly, for Column, *Left* weakly dominates *Right*. Geo's decision is the focus of the analysis: Geo's best choice depends on what it believes the other two players will do. The undominated equilibria of this game are those in which Row plays its weakly dominant strategy of *Up*, Column plays its weakly dominant strategy of *Left*, and Geo mixes between *East* and *West* with any probabilities. Although all these equilibria are proper, all but (*Up*, *Left*, *West*) embed the idea that Geo believes it is at least as likely that Row will deviate to play *Down* as that Column will deviate to play *Right*,

*For helpful comments, we thank Gabriel Carroll, Piotr Dworczak, Drew Fudenberg, Fuhito Kojima, Kevin Li, Barton Lipman, Ellen Muir, Andrzej Skrzypacz, Jeroen Swinkels, Alex Wolitzky, seminar participants, and anonymous referees. Milgrom thanks the National Science Foundation for support under grant number 1525730.

†Department of Economics, Stanford University. E-mail: milgrom@stanford.edu.

‡Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University. E-mail: joshua.mollner@kellogg.northwestern.edu.

despite the fact that deviating from equilibrium to play *Down* would be a costly mistake, while deviating to *Right* would be costless.[1]

Figure 1: A three player game†

[*West*]	*Left*	*Right*
Up	1, 1, 0	1, 1, 1
Down	0, 1, 0	1, 0, 0

[*East*]	*Left*	*Right*
Up	1, 1, 0	1, 1, 0
Down	0, 1, 1	1, 0, 0

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

In contrast, our new refinement, which we call *extended proper equilibrium*, requires any tremble that is costless in equilibrium to be regarded as more likely than any tremble that is costly, with the consequence that (*Up*, *Left*, *West*) is the unique extended proper equilibrium of this game. The definition of extended proper equilibrium incorporates this requirement, along with others, by formalizing the idea that there should be some scaling of the players' payoffs such that more costly mistakes are less likely, whether made by the same player *or by another player*. It is the italicized condition that represents the extra restriction in this definition compared with proper equilibrium. And because the definition permits flexibility in the chosen scaling, the set of extended proper equilibria is invariant to positive affine transformations (henceforth, simply "affine transformations") of the payoff matrix. For the game in Figure 1, extended proper equilibrium requires Geo to treat Row's costly deviation to *Down* as less likely than Column's costless deviation to *Right*, because, for any scaling, the former deviation has zero cost and the latter has positive cost.

The idea that players should be comparable in their propensities to tremble would be especially appropriate if the players were to originate from a common pool of *ex ante* identical agents. This idea can be modeled formally using the same approach that has been taken in biological game theory to generalize the concept of evolutionary stability to asymmetric games. In that literature, given any asymmetric game Γ, one constructs a meta-game $\bar{\Gamma}$ in which Nature moves first, randomizing the roles of the participants. Because each player might be cast into any of the roles, a pure strategy in the meta-game corresponds to a pure strategy profile in the original game.

We define a symmetrically proper strategy of $\bar{\Gamma}$ to be the limit of a sequence of strategies $\bar{\sigma}^{\varepsilon}$ such that each profile $(\bar{\sigma}^{\varepsilon}, \ldots, \bar{\sigma}^{\varepsilon})$ is an ε-proper equilibrium of $\bar{\Gamma}$. Every symmetrically proper strategy induces some strategy profile for the original game.[2] We show that a strategy profile of Γ is an extended proper equilibrium if and only if it is induced by a symmetrically proper strategy of a game $\bar{\Gamma}'$, the meta-version of some game Γ' that is equivalent to Γ up to an affine transformation of the payoffs. In this way, extended proper equilibrium can be characterized as the implication of play according to symmetrically proper strategies in associated meta-games.

One interpretation of this result is that some proper equilibria innately depend upon play-

[1]For example, to see that (*Up*, *Left*, *East*) is a proper equilibrium, define

$$\sigma^t_{row} = \left(\frac{t}{t+1}, \frac{1}{t+1}\right), \quad \sigma^t_{col} = \left(\frac{t^2}{t^2+1}, \frac{1}{t^2+1}\right), \quad \sigma^k_{geo} = \left(\frac{1}{t+1}, \frac{t}{t+1}\right).$$

For all $t \in \mathbb{N}$, $(\sigma^t_{row}, \sigma^t_{col}, \sigma^t_{geo})$ is a $\frac{1}{t}$-proper equilibrium. Taking the limit as $t \to \infty$ establishes that (*Up*, *Left*, *East*) is a proper equilibrium.

[2]A mixed strategy $\bar{\sigma}^{\varepsilon}$ for the meta-game $\bar{\Gamma}$ is a distribution over the pure strategy profiles of the original game Γ. If $\bar{\sigma}^{\varepsilon}$ happens to be a product distribution, then its factors constitute a mixed strategy profile of the original game, which we call the induced strategy profile. Generalizing beyond the product case, the induced strategy profile for the original game is the profile of the component-wise marginal distributions of $\bar{\sigma}^{\varepsilon}$ (*cf.* Definition 5).

ers being fundamentally asymmetric in their propensities to tremble. We also show that perfect equilibrium and Nash equilibrium can be analogously characterized, respectively, in terms of symmetrically perfect strategies and symmetrically Nash strategies of these same meta-games. In that sense, both perfection and Nash are consistent with across-player symmetry in the propensity to tremble. It is only for properness that requiring such symmetry can lead to further restrictions, and when symmetry is required in the way that we model, extended proper equilibrium is what emerges.

In Section 3, we state a formal definition of extended proper equilibrium. We also show that for finite games, such equilibria always exist and are always proper equilibria. While the above example illustrates that proper equilibria may fail to be extended proper equilibria in games with three or more players, we establish that the concepts coincide when there are precisely two players.

In Section 4, we present our results on meta-games. In Section 5, we characterize extended proper equilibrium in terms of lexicographic probability systems (as in Blume, Brandenburger and Dekel, 1991), which can be interpreted as capturing beliefs about how the game will be played. This characterization is useful in establishing one of the main results described above and may also be of independent interest. Finally, in Section 6, we analyze an application for our new refinement: the generalized second-price auction, which has been prominently used for internet advertising. Section 7 concludes.

2 Related Refinements

Our new refinement is closely related to *proper equilibrium* (Myerson, 1978): our definition not only refines proper equilibrium but also connects to it through the meta-game model that we describe. Despite its various critiques,[3] proper equilibrium remains a prominent refinement, has proven its theoretical worth in many applications, and is perhaps the most classic refinement to restrict the relative likelihood of different trembles. However, its restrictions apply only when comparing trembles of the same player. In contrast, we are motivated to seek a refinement based on across-player tremble restrictions. In determining how to formulate such restrictions, it seems natural to use proper equilibrium as a starting point and exemplar. We propose a refinement that entails not only the within-player restrictions already encoded in proper equilibrium but also novel across-player restrictions that are in the same spirit.

A handful of other refinements of proper equilibrium have also been proposed elsewhere in the literature. *Finely settled equilibrium* (Myerson and Weibull, 2015) refines proper equilibrium by requiring that the support of the equilibrium be contained in a block (i.e., a set of strategies for each player) that is minimal with respect to a certain property. Their *fully settled equilibrium* is a further refinement that adds a similar minimality requirement. *Fall-back proper equilibrium* (Kleppe, Borm and Hendrickx, 2017) is defined by considering a class of games in which each action of a player n is blocked by Nature with some independent probability δ_n. The limit points of the projections of the Nash equilibria of these blocking games are the fall-back proper equilibria, which form a subset of the proper equilibria. *Strictly proper equilibrium* (van Damme, 1991) refines strictly perfect equilibrium (Okada, 1981) by adding a certain continuity assumption; it also refines proper equilibrium. *Truly proper equilibrium* and *more-than-proper equilibrium* (Neary, 2010) also refine proper equilibrium, yet are not as demanding as strictly proper equilibrium. *Strongly proper equilibrium* (García Jurado and Prada Sánchez, 1990) refines proper equilibrium by adding the

[3]Notably, van Damme (1991) has shown that it cannot be justified by a micro-foundation in which players pay control costs to reduce their tremble probabilities. However, in subsequent work, Myerson and Weibull (2015) do successfully provide a micro-foundation for proper equilibrium, based on consideration sets. (See their Remark 4.)

requirement that if two strategies are payoff-equivalent for a player, then they must be trembled to with the same probability. *Persistent equilibrium* (Kalai and Samet, 1984) requires a form of local stability. Although persistent equilibrium is not itself a refinement of proper equilibrium, every game contains an equilibrium that is both persistent and proper. Similarly, every *fully stable set* (Kohlberg and Mertens, 1986) contains a proper equilibrium.

Another related refinement is *test-set equilibrium* (Milgrom and Mollner, 2018), which also incorporates the idea that a tremble by one player to a strategy that is a best response to equilibrium play must be more likely than any costly tremble. However, test-set equilibrium incorporates an additional idea: any single tremble to a best response is more likely than any joint tremble by two or more players. For that reason, test-set equilibrium can be more demanding than extended proper equilibrium—and in fact can be so demanding that its existence is not guaranteed in finite games, whereas every such game contains at least one extended proper equilibrium. Another difference is that test-set equilibrium does not incorporate the idea that the trembles of different players are statistically independent. For that reason, test-set equilibrium can be less demanding than extended proper equilibrium (and similarly, even perfect and proper equilibrium), which do require such independence.[4]

Finally, several other solution concepts can be thought of as imposing across-player restrictions on the likelihood of deviations akin to those contemplated by extended proper equilibrium and test-set equilibrium. Many of the refinements for signaling games effectively impose across-type restrictions on the sender so as to discipline the beliefs of the receiver following an off-path message. Examples include the *intuitive criterion* (Cho and Kreps, 1987), *divine equilibrium* (Banks and Sobel, 1987), and the *compatibility criterion* (Fudenberg and He, 2018). Closely related to the last of these is *player-compatible equilibrium* (Fudenberg and He, 2020), which extends the compatibility criterion and is applicable in general finite normal-form games. Player-compatible equilibrium refines trembling-hand perfect equilibrium by imposing across-player restrictions on tremble probabilities, but it does not refine proper equilibrium. *Quantal response equilibrium* (McKelvey and Palfrey, 1995) refers to a class of solution concepts that is parametrized by the profile of quantal response functions, which govern for each player how a strategy's expected payoff maps into its probability of being played. Nevertheless, the most commonly-used specification is the one in which each player has the same logistic quantal response function, implying a form of across-player symmetry in the propensity to err. Likewise, many learning or evolutionary models (e.g., fictitious play, replicator dynamics) are typically specified in ways that imply across-player symmetry in the rate of learning or evolution.

3 Extended Proper Equilibrium

3.1 Notation

A game in normal form is denoted $\Gamma = (\mathcal{N}, S, \pi)$, where $\mathcal{N} = \{1, \ldots, N\}$ is a set of players, $S = (S_n)_{n\in\mathcal{N}}$ is a profile of pure strategy sets, and $\pi = (\pi_n)_{n\in\mathcal{N}}$ is a profile of payoff functions. Throughout, we restrict attention to *finite* games, in which for all players $n \in \mathcal{N}$, S_n is a finite set. We use $\bar{S}$ to denote $\prod_{n\in\mathcal{N}} S_n$, the set of pure strategy profiles.

We use Δ_n to denote the set of mixed strategies of player n and Δ_n^0 to denote its relative interior, which is the set of totally mixed strategies of player n. We embed S_n in Δ_n and extend

[4]To elaborate, any perfect equilibrium is a mixed strategy profile (so that the mixing of different players is statistically independent), *and the same is true for the associated ε-perfect equilibria.* A test-set equilibrium is also a strategy profile, but the implicitly associated trembles are allowed to be certain *correlated* strategy profiles. (Test-set equilibrium is not defined in terms of trembles, but it can be thought of in those terms—see Appendix C.1 for details.)

the utility functions π_n to the domain $\prod_{n \in \mathcal{N}} \Delta_n$ in the usual way. A mixed strategy profile is denoted $\sigma = (\sigma_1, \ldots, \sigma_N)$.

We use $BR_n(\sigma)$ for the set of player n's best responses to σ. We use σ/σ'_n for the strategy profile constructed from σ by replacing player n's strategy with σ'_n. We also define $L_n(\sigma)$ as the expected loss for player n from playing σ_n instead of a best response when others play according to some mixed strategy profile σ:

$$L_n(\sigma) = \max_{s_n \in S_n} \pi_n(\sigma/s_n) - \pi_n(\sigma).$$

This quantity is zero when σ_n is a best response to σ and positive otherwise.

3.2 Definition

Given a scaling vector $\alpha \in \mathbb{R}^N_{++}$, an (α, ε)-extended proper equilibrium is a profile of totally mixed strategies in which the following property holds: if a pure strategy s'_l has scaled loss $\alpha_l L_l(\sigma/s'_l)$ exceeding that of another strategy s''_m of the same or another player, then s'_l is played with probability at most ε times that of s''_m.

Definition 1. Let $\alpha \in \mathbb{R}^N_{++}$ and $\varepsilon > 0$. An (α, ε)-*extended proper equilibrium* is a profile of totally mixed strategies $\sigma \in \prod_{n \in \mathcal{N}} \Delta^0_n$ such that for all $l, m \in \mathcal{N}$, all $s'_l \in S_l$, and all $s''_m \in S_m$, if $\alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m)$, then $\sigma_l(s'_l) \leq \varepsilon \cdot \sigma_m(s''_m)$.

This definition is stronger than that of ε-proper equilibrium (Myerson, 1978). Indeed, ε-proper equilibrium is equivalent to the version of this definition in which—instead of quantifying over all pairs of pure strategies s'_l and s''_m—the quantification is over only all pairs for the same player.[5]

An extended proper equilibrium is a strategy profile that, for some scaling vector α, is a limit of (α, ε)-extended proper equilibria, as ε approaches zero.

Definition 2. A strategy profile $\sigma \in \prod_{n \in \mathcal{N}} \Delta_n$ is an *extended proper equilibrium* if and only if there exist $\alpha \in \mathbb{R}^N_{++}$ and sequences $(\varepsilon_t)^\infty_{t=1}$ and $(\sigma^t)^\infty_{t=1}$ such that:

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

(*ii*) each σ^t is an (α, ε_t)-extended proper equilibrium, and

(*iii*) $\lim_{t \to \infty} \sigma^t = \sigma$.

One restriction implied by extended proper equilibrium is that any tremble that is costless in equilibrium must be much more likely than any costly tremble, *even one by a different player*.[6] No such restriction—nor in fact any across-player restriction—is implied by proper equilibrium alone.

The scaling vector α used in the definitions above can play two roles. In traditional non-cooperative game theory, payoffs are expressed in terms of von Neumann-Morgenstern utilities,

[5] For purposes of comparison, we reiterate the following definitions from Myerson (1978). An ε-proper equilibrium is a profile of totally mixed strategies $\sigma \in \prod_{n \in \mathcal{N}} \Delta^0_n$ such that $(\forall n \in \mathcal{N})(\forall s'_n \in S_n)(\forall s''_n \in S_n) : \pi_n(\sigma/s'_n) < \pi_n(\sigma/s''_n) \implies \sigma_n(s'_n) \leq \varepsilon \cdot \sigma_n(s''_n)$. Note that for any $\alpha_n > 0$, the antecedent inequality is equivalent to $\alpha_n L_n(\sigma/s'_n) > \alpha_n L_n(\sigma/s''_n)$. Similarly, an ε-perfect equilibrium is a profile of totally mixed strategies $\sigma \in \prod_{n \in \mathcal{N}} \Delta^0_n$ such that $(\forall n \in \mathcal{N})(\forall s'_n \in S_n)(\forall s''_n \in S_n) : \pi_n(\sigma/s'_n) < \pi_n(\sigma/s''_n) \implies \sigma_n(s'_n) \leq \varepsilon$. A strategy profile is a proper equilibrium (perfect equilibrium) if and only if it is the limit of ε-proper equilibria (ε-perfect equilibria) as $\varepsilon \to 0$, in the same sense as Definition 2.

[6] To see this, suppose σ is an extended proper equilibrium, so that it is the limit of a sequence $(\sigma^t)^\infty_{t=1}$ of (α, ε_t)-extended proper equilibria. Suppose also that there are two players $l \neq m$, as well as pure strategies $s'_l \notin BR_l(\sigma)$ and $s''_m \in BR_m(\sigma)$, which implies $\alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m)$. By continuity, we also have $\alpha_l L_l(\sigma^t/s'_l) > \alpha_m L_m(\sigma^t/s''_m)$—and therefore $\sigma^t_l(s'_l) \leq \varepsilon_t \sigma^t_m(s''_m)$—for sufficiently large values of t.

which are unique only up to affine transformations. With that interpretation, game-theoretic predictions should not depend on the particular payoff representation that is chosen. Thus, the first role of the scaling vector α in the above definition is to ensure that the set of extended proper equilibria is invariant to affine transformations of the utility functions of the players. In some models, however, payoffs are understood to be normalized in some way to facilitate interpersonal welfare comparisons. For example, with quasilinear preferences, payoffs may be normalized relative to a transferable numeraire good. In those settings, α may instead be interpreted as representing the propensities to tremble of the various players, and the definition as it is written allows for the possibility that players may differ somewhat in these propensities—albeit by only a finite multiple.[7]

3.3 Three Basic Facts About Extended Proper Equilibrium

The first three results record some basic facts about extended proper equilibrium. First, it refines the existing tremble-based refinements of proper equilibrium and perfect equilibrium. Second, its existence is guaranteed in finite games. Third, it coincides with proper equilibrium in two-player games.

Proposition 1. *For any finite normal-form game, the extended proper equilibria form a subset of the proper equilibria, which form a subset of the perfect equilibria, which form a subset of the Nash equilibria. All three inclusions may be strict.*

Proof. It is well known that the perfect equilibria form a subset of the Nash equilibria (e.g., Selten, 1975). For the two other inclusions, it is immediate from their definitions that any (α, ε)-extended proper equilibrium is an ε-proper equilibrium, which is an ε-perfect equilibrium. Consequently, an extended proper equilibrium, as the limit of (α, ε)-extended proper equilibria, is also the limit of ε-proper equilibria and therefore also proper. Likewise, a proper equilibrium is also perfect. That the first inclusion may be strict is demonstrated by the game in Figure 1 in the introduction. That the other inclusions may be strict is well known (e.g., Myerson, 1978). □

The proposition implies that any extended proper equilibrium inherits the properties that are possessed by every proper equilibrium, including the celebrated relationship between proper equilibria of the normal form and quasi-perfect equilibria of the extensive form, as established by van Damme (1984): an extended proper equilibrium of a normal-form game induces a quasi-perfect equilibrium (van Damme, 1984)—and hence a sequential equilibrium (Kreps and Wilson, 1982)—in every extensive-form game having this normal form.

Although extended proper equilibrium places restrictions beyond those required by proper equilibrium, the next result states that these additional restrictions are not so strong as to be incompatible with existence in finite games.

Theorem 2. *Every finite normal-form game has at least one extended proper equilibrium.*

[7] An alternative definition would be to require $\alpha = 1$ in the definition of extended proper equilibrium. This corresponds to the definition of *perfectly proper equilibrium* (García Jurado, 1989b). Under this definition, the selection fails to be invariant to affine transformations of the players' utility functions. But it may be viable for settings where payoffs have already been normalized. That definition is more restrictive than extended proper equilibrium, although not so restrictive as to be incompatible with existence in finite games.

Also related is *normalized perfectly proper equilibrium* (García Jurado, 1989a), which is equivalent to requiring $\alpha_n = [\max_{s \in \bar{S}} \pi_n(s) - \min_{s \in \bar{S}} \pi_n(s)]^{-1}$, and is also more restrictive than extended proper equilibrium. In contrast to his perfectly proper equilibrium, this definition is invariant to affine transformations of the players' utility functions; however, it may be quite sensitive to payoffs under strategy profiles that are unrelated to equilibrium play.

The proof of Theorem 2 is deferred to Appendix B, as are the proofs of most subsequent results. The proof is similar to the one used by Myerson (1978) to establish the existence of proper equilibrium. In the first step, a fixed point argument is used to establish existence of an (α, ε)-extended proper equilibrium for any $\alpha \in \mathbb{R}^N_{++}$ and every $\varepsilon > 0$. Thus, for any sequence $(\varepsilon_t)_{t=1}^{\infty}$ of positive numbers converging to zero, there is a corresponding sequence of (α, ε_t)-extended proper equilibria. In the second step, we appeal to the compactness of $\prod_{n \in \mathcal{N}} \Delta_n$ to establish existence of a convergent subsequence, the limit of which is therefore an extended proper equilibrium.

Although the extended proper equilibria may form a strict subset of the proper equilibria in games with three or more players (e.g., the game in Figure 1), the concepts coincide in games with just two players.

Theorem 3. *In two-player games, the sets of proper equilibria and extended proper equilibria coincide.*

The motivation for extended proper equilibrium comes from questions about how one player ought to form beliefs about the likelihood of deviations from different opponents. Extended proper equilibrium refines proper equilibrium by imposing additional structure on these across-opponent comparisons, but one interpretation of Theorem 3 is that extended proper equilibrium does not impose much more. Indeed, with two players, each player has only a single opponent (hence, none of these across-opponent comparisons exist to be made), and the result says that extended proper equilibrium reduces to proper equilibrium in that case.

Theorem 3 is, however, more subtle than the simple observation that there are no across-opponent comparisons to make in two-player games. Indeed, a given ε-proper equilibrium might fail to be (α, ε)-extended proper, even with just two players. (We illustrate with an example in the next section.) To prove the result, we instead show that for any convergent sequence of ε-proper equilibria, one can construct a sequence of (α, ε)-extended proper equilibria with the same limit. Our construction relies on lexicographic probability systems, which are discussed in Section 5. The details of the proof therefore wait until then.

3.4 Examples

To illustrate the above definitions and results, we next consider some examples.

Figure 1. We first return to the game in Figure 1 in the introduction, for which $(Up, Left, West)$ is the unique extended proper equilibrium. We now establish this formally. Suppose $\alpha \in \mathbb{R}^3_{++}$, $(\varepsilon_t)_{t=1}^{\infty}$, and $(\sigma^t)_{t=1}^{\infty}$ together satisfy the conditions of Definition 2. Because *Down* and *Right* are both weakly dominated strategies, we must have $\sigma^t_{row}(Down) \to 0$ and $\sigma^t_{col}(Right) \to 0$. Thus, when play is according to σ^t, Row's scaled loss from playing *Down* exceeds Column's scaled loss from playing *Right* for all sufficiently large t:

$$\alpha_{row} L_{row}(\sigma^t / Down) = \alpha_{row} \sigma^t_{col}(Left) \to \alpha_{row}$$
$$\alpha_{col} L_{col}(\sigma^t / Right) = \alpha_{col} \sigma^t_{col}(Down) \to 0$$

Thus, for all such t, (α, ε_t)-extended proper equilibrium requires $\sigma^t_{col}(Down) \le \varepsilon_t \sigma^t_{col}(Right)$. In particular, for all sufficiently large t, $\sigma^t_{col}(Down) < \sigma^t_{col}(Right)$. Turning our attention now to Geo, *East* therefore entails greater loss than *West* when opponents play according to σ^t. Thus, for all such t, (α, ε_t)-extended proper equilibrium requires $\sigma^t_{geo}(East) \le \varepsilon_t \sigma^t_{geo}(West)$. In particular, this implies $\sigma^t_{geo}(East) \to 0$. In conclusion, $(Up, Left, West)$ must be the limit of $(\sigma^t)_{t=1}^{\infty}$.

Of the solution concepts surveyed in Section 2, only test-set equilibrium (Milgrom and Mollner, 2018) also uniquely selects (*Up*, *Left*, *West*) in this game.[8] Indeed, both extended proper equilibrium and test-set equilibrium can be interpreted as imposing across-player restrictions on the likelihood of trembles.[9] But in general, neither concept implies the other, as subsequent examples illustrate: Figure 2(b) is a game in which extended proper equilibrium is more demanding than test-set equilibrium, and Figure 3 is an example of the reverse.

Figure 2. Next, consider the three-player extensive-form game with the centipede-like structure depicted in Figure 2(a). For each player $n \in \{1, 2, 3\}$, the continuing strategy C_n is dominated by the stopping strategy S_n. In consequence, (S_1, S_2, S_3) is the unique perfect, proper, extended proper, and test-set equilibrium of this game. But what we wish to emphasize with the game in Figure 2(a) is the relative probability of trembles by players 2 and 3.

Earlier discussion has focused on across-player comparisons of tremble probabilities between strategy pairs consisting of a best response to the equilibrium for one player and an inferior response to the equilibrium for another. This example illustrates that extended proper equilibrium can also imply meaningful restrictions even when it is two best responses to the equilibrium that are being compared. As Figure 2(a) suggests, such situations can naturally arise in situations based on certain extensive-form games.

Although both C_2 and C_3 are both best responses to the equilibrium (S_1, S_2, S_3), extended proper equilibrium requires player 2's tremble to be less likely than player 3's. To see this, suppose $\alpha \in \mathbb{R}^3_{++}$, $(\varepsilon_t)_{t=1}^{\infty}$, and $(\sigma^t)_{t=1}^{\infty}$ together satisfy the conditions of Definition 2. When play is according to σ^t, player 2's scaled loss from playing C_2 exceeds player 3's scaled loss from playing C_3 for all sufficiently large t:

$$\frac{\alpha_3 L_3(\sigma^t/C_3)}{\alpha_2 L_2(\sigma^t/C_2)} = \frac{\alpha_3 \sigma_1^t(C_1)\sigma_2^t(C_2)}{\alpha_2 \sigma_1^t(C_1)} = \frac{\alpha_3}{\alpha_2}\sigma_2^t(C_2) \to 0$$

Thus, for all such t, (α, ε_t)-extended proper equilibrium requires $\sigma_2^t(S_2) \leq \varepsilon_t \sigma_3^t(S_3)$. In contrast, proper equilibrium is consistent with either player 2 or 3 being more likely to tremble. Similarly, the restrictions on trembles that are implicit in the definition of test-set equilibrium are also consistent with either possibility—this is because C_2 and C_3 are both best responses to the equilibrium (S_1, S_2, S_3).

These differing restrictions on relative tremble probabilities (between proper equilibrium and test-set equilibrium on one hand and extended proper equilibrium on the other) do not translate into differing predictions for equilibrium play in the game in Figure 2(a) itself. However, if player 1's payoffs were instead such that its best strategy depended on the relative probabilities of trembles by players 2 and 3, then we would in fact obtain different predictions for equilibrium play. Figure 2(b) depicts such a game. From S_1, player 1 obtains a payoff of zero, as before. But from C_1, player 1's payoff now depends on the strategies of players 2 and 3. In particular, player 1 obtains a bonus of $+1$ if player 3 matches by playing C_3, but also a penalty of -1 if player 2 matches by playing

[8] Like proper equilibrium, the following concepts are consistent with any mixture between *East* and *West*: finely settled equilibrium, fully settled equilibrium, fall-back proper equilibrium, truly proper equilibrium, more-than-proper equilibrium, strongly proper equilibrium, persistent equilibrium, and player-compatible equilibrium. Both *East* and *West* (but not mixtures between them) are consistent with full stability. Like extended proper equilibrium, limiting logit quantal response equilibrium is consistent with only *West*; however, it also requires Column to play a uniform mixture between *Left* and *Right*, which is a weakly dominated strategy. Strictly proper equilibrium fails to exist in this game. Finally, as this game is not equivalent to a signaling game, the intuitive criterion, divine equilibrium, and the compatibility criterion are not defined.

[9] See Appendix C for a formal definition of test-set equilibrium, more detailed explanations of what it selects in the examples discussed in this section, and a further discussion of its relationship with extended proper equilibrium.

Figure 2: Two three-player games

(a) An extensive-form game

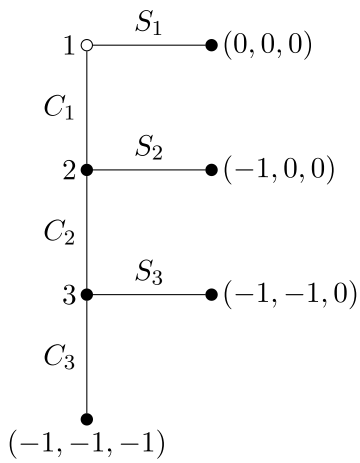

(b) A normal-form game

$[S_1]$	S_3	C_3
S_2	$0, 0, 0$	$0, 0, 0$
C_2	$0, 0, 0$	$0, 0, 0$

$[C_1]$	S_3	C_3
S_2	$0, 0, 0$	$1, 0, 0$
C_2	$-1, -1, 0$	$0, -1, -1$

C_2. Payoffs for players 2 and 3 are as before. For the reasons above, extended proper equilibrium requires C_2 to be less likely than C_3, which induces player 1 to select C_1 in equilibrium. In contrast, proper equilibrium and test-set equilibrium are both consistent with the opposite: C_2 being more likely than C_3, and hence an equilibrium choice of S_1. This is therefore a game in which extended proper equilibrium is stronger than the intersection of proper and test-set equilibrium.

Figure 3. Theorem 2 states that every finite game has at least one extended proper equilibrium, but the same conclusion is not true for test-set equilibrium. Indeed, in Milgrom and Mollner (2018), we present an example of a game—reproduced below as Figure 3—in which no test-set equilibrium exists. (*Up*, *Left*, *West*) is the unique Nash equilibrium of this game, but it is not a test-set equilibrium.

Nevertheless, (*Up*, *Left*, *West*) is an extended proper equilibrium of the game. Indeed, for any $\varepsilon > 0$, the following strategy profile is an (α, ε)-extended proper equilibrium with $\alpha = (1, 1, 1)$ and any sufficiently large choice of t:[10]

$$\sigma^t_{row} = \left(\frac{2t^2+1}{2t^2+3t+1}, \frac{3t}{2t^2+3t+1} \right)$$
$$\sigma^t_{col} = \left(\frac{2t-2}{2t+1}, \frac{3}{2t^2+3t+1}, \frac{3t}{2t^2+3t+1} \right)$$
$$\sigma^t_{geo} = \left(\frac{2t-2}{2t+1}, \frac{3}{2t+1} \right)$$

[10] To see that this profile is indeed an (α, ε)-extended proper equilibrium, we compute the following payoffs, from which one can easily check that all required restrictions are met:

$$\pi_{row}(\sigma^t / Up) = \pi_{col}(\sigma^t / Left) = \pi_{geo}(\sigma^t / West) = 0$$
$$\pi_{row}(\sigma^t / Down) = \pi_{col}(\sigma^t / Right) = \pi_{geo}(\sigma^t / East) = -\frac{3(t-1)}{(t+1)(2t+1)^2}$$
$$\pi_{col}(\sigma^t / Center) = -\frac{3(5t+1)}{(t+1)(2t+1)^2}$$

Moreover, $\lim_{t\to\infty} \sigma^t = (\mathit{Up}, \mathit{Left}, \mathit{West})$.

Note that as $t \to \infty$, Column's tremble to *Center* becomes infinitely less likely than its tremble to *Right*, despite the fact that both are best responses to the equilibrium $(\mathit{Up}, \mathit{Left}, \mathit{West})$. In that sense, these trembles might be regarded as extreme. Owing to this extreme trembling, the unilateral deviation to *Center* is less likely than the joint deviation to $(\mathit{Down}, \mathit{Right})$, which runs contrary to the logic embedded in the definition of test-set equilibrium—but which is nevertheless consistent with extended proper equilibrium.

Figure 3: 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.

Figure 4. Finally, consider the two-player game in Figure 4. For Row, *Up* strictly dominates *Down*. For Column, *Left* weakly dominates *Right*. Thus, $(\mathit{Up}, \mathit{Left})$ is the game's unique perfect equilibrium. It follows that the sets of proper and extended proper equilibria coincide: both are the singleton $\{(\mathit{Up}, \mathit{Left})\}$. Although it is therefore clear that the claim of Theorem 3 holds in the case of this game, it is useful to consider the trembles that justify $(\mathit{Up}, \mathit{Left})$ as proper and extended proper.

Figure 4: A two player game†

	Left	*Right*
Up	$1,1$	$1,1$
Down	$0,1$	$0,0$

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

In particular, consider the following class of strategy profiles: for $t \in \mathbb{N}$, let σ^t_{row} specify *Down* with probability $\frac{1}{t}$ (and *Up* otherwise), and let σ^t_{col} similarly specify *Right* with probability $\frac{1}{t}$ (and *Left* otherwise). For any $\varepsilon > 0$, it is straightforward to show that $\sigma^t = (\sigma^t_{row}, \sigma^t_{col})$ is an ε-proper equilibrium when t is sufficiently large. However, an analogous statement is not true for (α, ε)-extended proper equilibria. Rather, for any $\alpha \in \mathbb{R}^2_{++}$ and any $\varepsilon \in (0,1)$, σ^t fails to be an (α, ε)-extended proper equilibrium when t is large. To see this, observe that, when play is according to σ^t, Row's scaled loss from *Down* exceeds Column's scaled loss from *Right* for all sufficiently large t:

$$\alpha_{row} L_{row}(\sigma^t/\mathit{Down}) = \alpha_{row}$$
$$\alpha_{col} L_{col}(\sigma^t/\mathit{Right}) = \frac{\alpha_{col}}{t} \to 0$$

If σ^t were an (α, ε)-extended proper equilibrium, then for such t we would have $\frac{1}{t} = \sigma^t_{row}(\mathit{Down}) \leq \varepsilon \cdot \sigma^t_{col}(\mathit{Right}) = \frac{\varepsilon}{t}$, a contradiction.

The key to Theorem 3 lies in showing that, when there are two players, the trembles σ^t can be perturbed so as to satisfy the requirements of extended proper equilibrium. In this case, it is enough to reduce the probability of Row's trembles: let $\hat{\sigma}^t_{row}$ specify *Down* with probability $\frac{1}{t^2}$ (and

Up otherwise). This change does not disturb the within-player restrictions that had already been satisfied: just as σ^t is an ε-proper equilibrium when t is sufficiently large, so is $(\hat{\sigma}^t_{row}, \sigma^t_{col})$. More significantly, the change accommodates the across-player restrictions. Indeed, the following does now hold when t is sufficiently large: $\frac{1}{t^2} = \hat{\sigma}^t_{row}(Down) \leq \varepsilon \cdot \sigma^t_{col}(Right) = \frac{\varepsilon}{t}$.

After introducing lexicographic probability systems in Section 5, we sketch the proof of Theorem 3, which entails showing how such perturbations can be constructed in general.

4 Equilibrium Refinement in Meta-Games

For some applications, it may be natural to expect players to be symmetric in their propensities to tremble, but the definitions of standard tremble-based refinements do not entail any across-player restrictions designed to encode such symmetry. This section proposes a means of encoding symmetry into the definitions of the standard refinements, investigates the implications, and demonstrates a way in which extended proper equilibrium emerges from this analysis.

Although there might be many reasons for expecting players to have similar trembling propensities, one could be that they originate from a single pool of *ex ante* identical agents. This idea can be modeled by embedding a game Γ into a symmetric meta-game $\bar{\Gamma}$ in which Γ is played only after Nature randomly assigns players to roles.

In this section, we consider the implications of the standard tremble-based refinements when two forms of across-player symmetry are enforced. First, we apply refinements to a class of meta-games that includes the meta-version of Γ itself (i.e., $\bar{\Gamma}$) as well as the meta-versions of games that are equivalent to Γ up to affine transformations of the payoffs. Second, we focus on symmetric versions of the refinements in which symmetry is required not only of the equilibrium strategy profile but also of the trembles. The predictions of Nash equilibrium and perfect equilibrium are unaffected by embedding symmetry in this way: applying Nash equilibrium or this form of perfection to these meta-games generates predictions for observed play in Γ that coincide with the predictions derived by applying Nash equilibrium or perfect equilibrium directly to Γ. However, a corresponding result is not true for proper equilibrium. Rather, applying this form of properness to these meta-games generates predictions for observed play in Γ that coincide with the predictions derived by applying *extended* proper equilibrium directly to Γ.

4.1 Preliminaries

Given an N-player game Γ, we let $\bar{\Gamma}$ represent the interaction among N agents in which they play Γ after being randomly assigned to their roles. Players in $\bar{\Gamma}$ can be thought of as choosing their strategies for Γ behind a "veil of ignorance," that is, before they discover their role assignment. Thus, pure strategies in $\bar{\Gamma}$ are isomorphic to pure strategy *profiles* in Γ.[11] Likewise, mixed strategies in $\bar{\Gamma}$ are isomorphic to distributions over pure strategy profiles in Γ.

Definition 3. The *meta-game* associated with a game $\Gamma = (\mathcal{N}, S, \pi)$ is the symmetric game $\bar{\Gamma}$ consisting of the same set of players $\mathcal{N}$, and where each player $n \in \mathcal{N}$ has strategy set $\bar{S} = \prod_{n \in \mathcal{N}} S_n$ and payoff function

$$\bar{\pi}_n(\bar{s}_1, \ldots, \bar{s}_N) = \frac{1}{N!} \sum_{\tau \in \mathrm{Sym}(\mathcal{N})} \pi_{\tau(n)}\Big(\mathrm{proj}_1\left(\bar{s}_{\tau^{-1}(1)}\right), \ldots, \mathrm{proj}_N\left(\bar{s}_{\tau^{-1}(N)}\right)\Big).^{12}$$

[11] Thus, the elements of $\bar{S}$ can be interpreted either as a strategy profile in Γ or as a strategy in $\bar{\Gamma}$. We will typically denote such an element by s when the former interpretation is intended and by $\bar{s}$ when the latter interpretation is intended.

[12] In this sum, $\mathrm{Sym}(\mathcal{N})$ denotes the set of all permutations of $\mathcal{N}$. Under a permutation τ, the role of player m in

In the biological game theory literature, this meta-game construction is used to apply the concept of an evolutionarily stable strategy—which was originally defined only for symmetric games—to asymmetric games (e.g., Maynard Smith and Parker, 1976; Selten, 1980, 1983; van Damme, 1991). Similarly, we exploit the symmetric structure of $\bar{\Gamma}$ to apply symmetric versions of the standard tremble-based refinements to that game. Note that the following definitions require symmetry not only of the equilibrium strategies but also of the trembles (as in e.g., Nachbar, 1990).

Definition 4. In a finite symmetric game Γ,

(*i*) a strategy σ is a *symmetrically Nash strategy* if $(\sigma, \dots, \sigma)$ is a Nash equilibrium of Γ;

(*ii*) a strategy σ is a *symmetrically perfect strategy* if there exists a sequence of positive numbers $(\varepsilon_t)_{t=1}^{\infty}$ converging to zero and a sequence of strategies $(\sigma^t)_{t=1}^{\infty}$ converging to σ such that for all t, $(\sigma^t, \dots, \sigma^t)$ is an ε_t-perfect equilibrium of Γ; and

(*iii*) a strategy σ is a *symmetrically proper strategy* if there exists a sequence of positive numbers $(\varepsilon_t)_{t=1}^{\infty}$ converging to zero and a sequence of strategies $(\sigma^t)_{t=1}^{\infty}$ converging to σ such that for all t, $(\sigma^t, \dots, \sigma^t)$ is an ε_t-proper equilibrium of Γ.

4.2 Results

A goal of this section is to apply the concepts of symmetrically Nash, perfect, and proper strategies to the meta-game $\bar{\Gamma}$ and to characterize the implications of such strategies for observed play in Γ. To do so, we observe that every mixed strategy in $\bar{\Gamma}$ implies a mixed strategy profile in Γ in the following way.

Recall that every mixed strategy $\bar{\sigma}$ in $\bar{\Gamma}$ is equivalent to a distribution over pure strategy profiles in Γ. Unless this distribution has a product structure, it may imply some across-role correlation of behavior. Nevertheless, because each player assumes only one role at once in $\bar{\Gamma}$, these correlations are irrelevant to what could be observed as the induced behavior in Γ. Rather, if play in $\bar{\Gamma}$ is according to the symmetric mixed strategy profile $(\bar{\sigma}, \dots, \bar{\sigma})$, then what would be observed as the induced behavior in Γ is the profile consisting of the component-wise marginals of $\bar{\sigma}$, which we define below as the *projection* of $\bar{\sigma}$.

Definition 5. Given a game Γ and a distribution $\bar{\sigma}$ over the pure strategy profiles of Γ, the *projection* of $\bar{\sigma}$ is the strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$ in Γ, where for all $n \in \mathcal{N}$, σ_n is the marginal of $\bar{\sigma}$ on S_n.

We can then state our result, one direction of which implies that (*i*) symmetrically Nash strategies of $\bar{\Gamma}$ induce Nash equilibrium play in Γ, (*ii*) symmetrically perfect strategies of $\bar{\Gamma}$ induce perfect equilibrium play in Γ, and (*iii*) symmetrically proper strategies of $\bar{\Gamma}$ induce *extended* proper equilibrium play in Γ. Furthermore, the same conclusions also hold for symmetrically Nash, perfect, and proper strategies of the meta-versions of games that are equivalent to Γ up to affine transformations of the payoffs. The other direction of the result establishes a set of converses.

Theorem 4. *A strategy profile σ in a finite game Γ is a Nash (or perfect or extended proper, respectively) equilibrium if and only if there exists a distribution $\bar{\sigma}$ over the strategy profiles of Γ, which has σ as its projection, and there exists a game Γ', which is equivalent to Γ up to affine transformations of the payoffs, such that $\bar{\sigma}$ is a symmetrically Nash (or perfect or proper, respectively) strategy of the meta-game $\bar{\Gamma}'$.*

the game Γ is assigned to player $\tau^{-1}(m)$ in the meta-game $\bar{\Gamma}$, so the strategy played by role m in Γ is $\mathrm{proj}_m(\bar{s}_{\tau^{-1}(m)})$. By averaging the payoffs to player n in $\bar{\Gamma}$ over all permutations, we obtain the given form for $\bar{\pi}_n$.

In other words, Nash equilibrium remains the prediction if symmetry is embedded into the Nash equilibrium logic in the way described above. Likewise, perfect equilibrium remains the prediction if symmetry is embedded into the perfect equilibrium logic. However, an analogous result is not true for proper equilibrium. Rather, if players originate from the same population before being cast into the different roles of a game Γ and if such players tremble symmetrically and in the way of proper equilibrium, then some additional restrictions on trembles in Γ are implied. These restrictions are precisely those implied by extended proper equilibrium when applied directly to Γ.[13]

The parts of the result pertaining to Nash equilibrium are related to the equivalence between ex ante and interim Bayesian Nash equilibrium in games of incomplete information, and the intuition is similar. The parts pertaining to perfect equilibrium are also straightforward to establish. For those parts pertaining to extended proper equilibrium, we argue as follows:

Necessity. Suppose $\bar{\sigma}$ is a symmetrically proper strategy of $\bar{\Gamma}$. Let $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)_{t=1}^{\infty}$ be an associated sequence of ε_t-proper equilibria of $\bar{\Gamma}$. Let σ^t be the projection of $\bar{\sigma}^t$. It can be checked that σ^t is an $(\alpha, |\bar{S}|\varepsilon_t)$-extended proper equilibrium of Γ with $\alpha = (1, \ldots, 1)$. It follows that the projection of $\bar{\sigma}$ is an extended proper equilibrium of Γ. If we had instead started with a symmetrically proper strategy of the meta-version of a game Γ' that is equivalent to Γ up to affine transformations of the payoffs, then we would simply use a correspondingly different choice for α.

Sufficiency. Suppose σ is an extended proper equilibrium of Γ. Let $(\sigma^t)_{t=1}^{\infty}$ be an associated sequence of $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibria. Let $\phi^t = \prod_{n \in \mathcal{N}} \sigma_n^t$ be the distribution over $\bar{S}$ induced by σ^t. We would be done if we could show that $(\phi^t, \ldots, \phi^t)$ were an ε_t-proper equilibrium of $\bar{\Gamma}$ (or the meta-version of a game that is equivalent to Γ up to an affine transformation of the payoffs). But this is not guaranteed to be the case. The bulk of the proof lies in using ϕ^t to construct another distribution over $\bar{S}$ that also has σ^t as its projection and that does constitute such an ε_t-proper equilibrium.[14]

5 Lexicographic Characterization

5.1 Framework

In this section, we employ the framework of Blume, Brandenburger and Dekel (1991)—henceforth, BBD—to characterize extended proper equilibrium directly in terms of players' beliefs. Toward that end, let $\rho = (p^1, \ldots, p^K)$ be a sequence of probability distributions over $\bar{S}$, the set of pure strategy profiles. BBD refer to such a sequence as a *lexicographic probability system* (LPS).[15] An LPS can be interpreted as follows: p^1 represents the primary theory about how the game will be played, p^2 represents beliefs about how the game will be played in the zero-probability event that the primary theory is incorrect, and so on.[16] To economize on notation, we define

[13]In the proof of Theorem 4, we additionally establish that extended proper equilibrium remains the prediction if symmetry is embedded into its logic (in the same way that symmetry changes neither Nash nor perfect equilibrium). Only for proper equilibrium does the solution change.

[14]In Appendix D, we provide further intuition for the proof of this part of the result by illustrating the construction of this alternative distribution in the context of a particular example game.

[15]Our approach and notation differ slightly from BBD. One difference is that we directly assume that players' beliefs about the strategies of their opponents are the marginals of a "common prior LPS," whereas BBD add that assumption only when extending their focus beyond two-player games. Because extended proper equilibrium incorporates additional restrictions only in games with three or more players, it is natural for us to assume this common prior from the outset. There are also small notational differences (e.g., while they designate players with superscripts and levels of an LPS with subscripts, we do the reverse).

[16]This interpretation is valid even if the supports of $(p^1, \ldots, p^K)$ overlap. BBD point out that it is necessary to allow for overlapping supports.

$\pi_n(s_n|p^k) = \sum_{s\in\bar{S}} p^k(s)\pi_n(s/s_n)$. In words, this is player n's payoff from s_n when others' play is distributed according to p^k (i.e., the kth level of the LPS).

We define the best response set of player $n \in \mathcal{N}$ to the LPS ρ as follows, where we use the symbol $\geq_L$ to represent the lexicographic ordering.[17]

$$BR_n(\rho) = \left\{ s'_n \in S_n \middle| \, \forall s''_n \in S_n : \left[\pi_n(s'_n|p^k)\right]_{k=1}^{K} \geq_L \left[\pi_n(s''_n|p^k)\right]_{k=1}^{K} \right\}.$$

As additional notation, given any $J \subseteq \mathcal{N}$, let p^k_J be the marginal distribution of p^k on $\prod_{n\in J} S_j$. An LPS ρ gives rise to a partial order on $\bigcup_{J\subseteq\mathcal{N}} \prod_{n\in J} S_j$, which is the set of pure strategy profiles for subsets of $\mathcal{N}$. If $s_J \in \prod_{n\in J} S_j$ and $s_I \in \prod_{n\in I} S_i$, then we say that $s_J >_\rho s_I$, read as "s_J is infinitely more likely than s_I according to the LPS ρ," if $\min\{k : p^k_J(s_J) > 0\} < \min\{k : p^k_I(s_I) > 0\}$. We write $s_J \geq_\rho s_I$ to mean it is not the case that $s_I >_\rho s_J$.

Next, we repeat a definition from BBD. A lexicographic Nash equilibrium is a pair: an LPS that captures beliefs about how the game will be played and a strategy profile.

Definition 6. A pair (ρ, σ) is a *lexicographic Nash equilibrium* if

(i) for all $n \in \mathcal{N}$ and all $s_n \in S_n$, $p^1_n(s_n) > 0$ implies $s_n \in BR_n(\rho)$, and

(ii) for all $s \in \bar{S}$, $p^1(s) = \prod_{n\in\mathcal{N}} \sigma_n(s_n)$.

Condition (i) requires a form of rationality: under the primary theory, player n assigns zero probability to strategies that are not best responses to the beliefs. Condition (ii) requires a form of consistency: the primary theory for how the game will be played coincides with the strategies.

We next introduce two properties that an LPS may possess, which are progressively more demanding. The first, "respects within-player preferences," is equivalent to what BBD define as "respects preferences." We use this different terminology to accentuate the distinction between this property and "respects within-and-across-player preferences," which is also defined below.

Definition 7. An LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ *respects within-player preferences* if for all $n \in \mathcal{N}$ and all $s'_n, s''_n \in S_n$, $s'_n \geq_\rho s''_n$ implies

$$\left[\pi_n(s'_n|p^k)\right]_{k=1}^{K} \geq_L \left[\pi_n(s''_n|p^k)\right]_{k=1}^{K}.$$

Definition 8. An LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ *respects within-and-across-player preferences* if there exists some $\alpha \in \mathbb{R}^N_{++}$ such that for all $l, m \in \mathcal{N}$, $s^*_l \in BR_l(\rho)$, $s^*_m \in BR_m(\rho)$, and all $s'_l \in S_l$, $s''_m \in S_m$, $s'_l \geq_\rho s''_m$ implies

$$\left[\alpha_l \left(\pi_l(s^*_l|p^k) - \pi_l(s'_l|p^k)\right)\right]_{k=1}^{K} \leq_L \left[\alpha_m \left(\pi_m(s^*_m|p^k) - \pi_m(s''_m|p^k)\right)\right]_{k=1}^{K}.$$

Definition 7 requires that under the beliefs, player n is infinitely less likely to use a strategy that is a "worse response" to the beliefs than a strategy that is a "better response." This strengthens condition (i) of Definition 6, which imposed this requirement only when the latter strategy in the comparison was a best response. Definition 8 strengthens Definition 7 by adding similar requirements for across-player comparisons. In particular, Definition 8 requires the existence of a scaling vector α such that whenever the scaled loss (relative to a best response to the beliefs) of a strategy s'_l exceeds the scaled loss of another strategy s''_m, whether of the same or another player, then s'_l must be infinitely less likely than s''_m under the beliefs.

Finally, we repeat two other definitions from BBD. These are generalizations for LPSs of what it means for a probability distribution to have full support or to be a product distribution.

[17] Formally, for $a, b \in \mathbb{R}^K$, $a \geq_L b$ if and only if whenever $b^k > a^k$, there exists a $j < k$ such that $a^j > b^j$.

Definition 9. An LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ has *full support* if for each $s \in \bar{S}$, $p^k(s) > 0$ for some $k \in \{1, \ldots, K\}$.

Definition 10. An LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ satisfies *strong independence* if there is an equivalent $\mathbb{F}$-valued probability distribution that is a product distribution, where $\mathbb{F}$ is some non-Archimedean ordered field that is a proper extension of $\mathbb{R}$.[18]

To clarify the latter definition and in anticipation of subsequent discussion, it is useful to define the following notation. If $r = (r^1, \ldots, r^{K-1})$ is a vector and $\rho = (p^1, \ldots, p^K)$ is an LPS, then let $r\Box\rho$ denote the probability distribution

$$\begin{aligned} r\Box\rho = (1 - r^1)p^1 + r^1 \left[(1 - r^2)p^2 + r^2 \left[(1 - r^3)p^3\right.\right. \\ \left.\left. + \ r^3 \left[\cdots + r^{K-2} \left[(1 - r^{K-1})p^{K-1} + r^{K-1}p^K\right] \cdots\right]\right]\right]. \end{aligned} \tag{1}$$

As BBD show, ρ satisfies strong independence if and only if there exists a sequence $r(t) \in (0, 1)^{K-1}$ with $r(t) \to 0$ such that $r(t)\Box\rho$ is a product distribution for all t.

5.2 Characterizations

Using these definitions, Proposition 8 provides a characterization of extended proper equilibrium. For comparison, Propositions 5–7 are characterizations of Nash equilibrium, perfect equilibrium, and proper equilibrium, which are due to BBD.[19] These characterizations are useful for two reasons. First, they simplify some proofs (particularly that of Theorem 3), since LPSs are easier to work with than sequences of trembles. Second, compared to sequences of trembles, LPSs correspond more closely to intuitive statements about some actions being "infinitely more likely" than others.

Proposition 5. *A strategy profile σ is a Nash equilibrium if and only if there exists some LPS ρ on $\bar{S}$ that satisfies strong independence for which (ρ, σ) is a lexicographic Nash equilibrium.*

Proposition 6. *A strategy profile σ is a perfect equilibrium if and only if there exists some LPS ρ on $\bar{S}$ that satisfies strong independence and has full support for which (ρ, σ) is a lexicographic Nash equilibrium.*

Proposition 7. *A strategy profile σ is a proper equilibrium if and only if there exists some LPS ρ on $\bar{S}$ that satisfies strong independence, has full support, and respects within-player preferences for which (ρ, σ) is a lexicographic Nash equilibrium.*

Proposition 8. *A strategy profile σ is an extended proper equilibrium if and only if there exists some LPS ρ on $\bar{S}$ that satisfies strong independence, has full support, and respects within-and-across-player preferences for which (ρ, σ) is a lexicographic Nash equilibrium.*

[18] An ordered field $\mathbb{F}$ is non-Archimedean if it contains a non-zero infinitesimal (i.e., an element $\varepsilon \neq 0$ such that $|\varepsilon| < \frac{1}{n}$ for all $n \in \mathbb{N}$). An LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ and an $\mathbb{F}$-valued probability distribution $\hat{p}$ over $\bar{S}$, are said to be *equivalent* if there exists a vector of positive infinitesimals $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_{K-1}) \in \mathbb{F}^{K-1}$ such that $\hat{p}(s) = \varepsilon\Box\rho(s)$ for all $s \in \bar{S}$, where $\varepsilon\Box\rho$ is defined as in equation (1).

[19] Propositions 6 and 7 restate Propositions 7 and 8 of BBD. While Proposition 5 is in the same spirit as their Proposition 3, which also characterizes Nash equilibrium in terms of lexicographic probability systems, it is not an identical result. We therefore provide a separate proof. The differences are: (*i*) their result does not require players' beliefs about the strategies of their opponents to be the marginals of a "common prior LPS," and (*ii*) their result does not require the LPS to satisfy strong independence. However, the same result holds both with and without these requirements. As discussed in footnote 15, it is natural for us to assume this common prior from the outset. We also require strong independence for more continuity with the other characterizations.

For intuition into one direction of these characterizations, suppose $\rho = (p^1, \ldots, p^K)$ is an LPS on $\bar{S}$ that satisfies strong independence. Hence, there exists a sequence $r(t) \in (0,1)^{K-1}$ with $r(t) \to 0$ such that $r(t)\Box\rho$ is a product distribution for all t, where we let σ^t denote its profile of component-wise marginals.

Now suppose that σ is a strategy profile and that (ρ, σ) is a lexicographic Nash equilibrium. We have $r(t)\Box\rho \to p^1$, so that by condition *(ii)* in the definition of lexicographic Nash equilibrium, $\sigma^t \to \sigma$. By condition *(i)* in the definition of lexicographic Nash equilibrium, $\sigma_n(s_n) > 0$ implies $s_n \in BR_n(\rho)$. This can be shown to imply $s_n \in BR_n(\sigma^t)$ for all sufficiently large t.[20] By continuity, it follows that $s_n \in BR_n(\sigma)$, so that σ is a Nash equilibrium, establishing one direction of Proposition 5.

Straightforward arguments establish the same direction of the other characterizations. Suppose ρ also has full support. It can be shown that for any $\varepsilon > 0$, σ^t is an ε-perfect equilibrium for all sufficiently large t. (For instance, full support of ρ implies that σ^t is totally mixed.) Analogously, it can be shown that respecting within-player preferences (respectively, respecting within-and-across-player preferences) implies that for any $\varepsilon > 0$, σ^t is an ε-proper equilibrium (respectively, an (α, ε)-extended proper equilibrium) for all sufficiently large t. Thus, in each case, the limiting strategy profile σ satisfies the claimed refinement.

5.3 Proof Sketch of Theorem 3

We are now in a position to sketch the proof of Theorem 3, which states that proper equilibrium and extended proper equilibrium coincide in two-player games. Let σ be a proper equilibrium of such a game. Let ρ be a corresponding LPS with the properties guaranteed by Proposition 7. To show that σ is an extended proper equilibrium, we construct trembles of the form $((r_1\Box\rho)_1, (r_2\Box\rho)_2)$ with $r_1, r_2 \in (0,1)^{K-1}$.

Within-player restrictions. All sufficiently small choices of (r_1, r_2) guarantee that the resulting strategy profile is an ε-proper equilibrium. To see this, focus on the perspective of player 1. The distribution of the opponent's play induced by $((r_1\Box\rho)_1, (r_2\Box\rho)_2)$ coincides with that of $r_2\Box\rho$.[21] Suppose s_1' delivers a strictly lower payoff than s_1'' against such play, where r_2 is small. Because ρ respects within-player preferences, it can be shown that $s_1' <_\rho s_1''$. And this implies that s_1' is much less likely than s_1'' under $(r_1\Box\rho)_1$ when r_1 is small, as required.

Across-player restrictions. To complete the proof, it suffices to demonstrate that for any $\alpha \in \mathbb{R}^2_{++}$, there exists a choice of (r_1, r_2) under which the across-player restrictions are also satisfied, so that the resulting strategy profile is an (α, ε)-extended proper equilibrium. This can be done using a corollary of the Eilenberg and Montgomery (1946) fixed point result.[22]

To illustrate further, return to the two-player game in Figure 4. As mentioned, (*Up*, *Left*) is a proper equilibrium of that game. Proposition 7 therefore guarantees the existence of an LPS with certain properties. One such LPS is $\rho = (p^1, p^2, p^3)$ in which p^1 puts full weight on the

[20] In fact, it can be shown that *all* payoff rankings induced by ρ coincide with those induced by σ^t when t is sufficiently large. See Proposition 1 of BBD or Lemma 12 in the appendix.

[21] This is where the presence of only two players becomes relevant. With three players, the distribution of play for players 2 and 3 induced by $((r_1\Box\rho)_1, (r_2\Box\rho)_2, (r_3\Box\rho)_3)$ is guaranteed to coincide with that of $r\Box\rho$ for some choice of r only if $r_2 = r_3$. Symmetric observations would limit the search to the case of $r_1 = r_2 = r_3$, which removes the flexibility necessary for satisfying across-player restrictions.

[22] Let us briefly comment on the barrier to a more direct proof (i.e., not in terms of the LPS machinery). The proof strategy relies on formulating a set a strategy profiles such that (*i*) every element of the set is close to σ, (*ii*) every element of the set is an ε-proper equilibrium, and (*iii*) the set is sufficiently rich to accommodate the fixed-point argument used to satisfy the additional across-player restrictions. The LPS machinery permits a tractable formulation of such a set.

equilibrium (*Up*, *Left*); p^2 puts equal weight on the two scenarios in which a single player trembles, (*Up*, *Right*) and (*Down*, *Left*); and p^3 puts full weight on the scenario in which both players tremble, (*Down*, *Right*).[23] To show that the equilibrium is also extended proper, let $\alpha = (1, 1)$ and use ρ to construct a sequence of (α, ε)-extended proper equilibria using the approach described above. In this case, the aforementioned fixed point construction might deliver the pair of sequences

$$r_{row}(t) = \left(\frac{2t-1}{t^3}, \frac{1}{2t-1}\right) \quad r_{col}(t) = \left(\frac{2t-1}{t^2}, \frac{1}{2t-1}\right).$$

Then for each player n, let $\hat{\sigma}_n^t$ be the marginal of $r_n(t)\Box\rho$ on S_n. Computing, we find that this is the same sequence of (α, ε)-extended proper equilibria previously mentioned in conjunction with this game: $\hat{\sigma}_{row}^t$ specifies *Down* with probability $\frac{1}{t^2}$ (and *Up* otherwise), and $\hat{\sigma}_{col}^t$ specifies *Right* with probability $\frac{1}{t}$ (and *Left* otherwise).

6 Application: Generalized Second-Price Auction

Extended proper equilibrium can be productively applied to analyze the generalized second-price (GSP) auction, which is an auction format that has been widely used to sell search advertising on the Internet. This auction was modeled and studied by Edelman, Ostrovsky and Schwarz (2007) who also proposed *locally envy-free equilibrium* as a refinement of the auction's many equilibria in settings of complete information.[24] That refinement, however, is defined directly in terms of the GSP auction game and does not apply to general non-cooperative games, rendering it difficult to assess the logic of the refinement separately from the game. In contrast, extended proper equilibrium is a general refinement that, as we argue below, can do much of the same work.

6.1 Framework

There are I ad positions and N bidders. The ith position has a click rate $\kappa_i > 0$. Bidder n has a per-click value $v_n > 0$. Bidder n's payoff from being in position i is therefore $\kappa_i v_n$ minus its payments to the auctioneer. We label positions and bidders so that click rates and bidder values are strictly in descending order: $\kappa_1 > \cdots > \kappa_I$ and $v_1 > \cdots > v_N$. For simplicity, we also assume $N \leq I + 1$.

Bids in a GSP auction are submitted simultaneously. Let $b^{(i)}$ denote the ith highest bid, and for convenience, define $b^{(N+1)} = 0$. After breaking ties uniformly at random, position i is allocated to the ith highest bidder at a per-click price $b^{(i+1)}$, for a total payment $\kappa_i b^{(i+1)}$.

In the original analysis, allowable bids are the nonnegative reals. However, the refinements of proper equilibrium and extended proper equilibrium are defined only for finite games. Consequently, our analysis restricts bids to the finite set $\{0, \frac{1}{m}, \frac{2}{m}, \ldots, \frac{mM}{m}\}$ for positive integers m and M. Altogether, this defines a game of complete information among the bidders.

We assume M is large enough to be non-binding (e.g., $M > v_1$ suffices). It can be shown that if m is also sufficiently large, then ties do not occur in the pure Nash equilibria;[25] we assume that

[23] In fact, this LPS corresponds to the sequence of ε-proper equilibria previously mentioned in conjunction with this game. As before, let σ_{row}^t specify *Down* with probability $\frac{1}{t}$ (and *Up* otherwise), and let σ_{col}^t similarly specify *Right* with probability $\frac{1}{t}$ (and *Left* otherwise). The probability distribution induced by σ^t coincides with $r(t)\Box\rho$, where $r(t) = (\frac{2t-1}{t^2}, \frac{1}{2t-1})$.

[24] Varian (2007) studies the same model and makes the same selection, calling these *symmetric equilibria*.

[25] Lemma 17 in the appendix shows this is the case if $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$, where $\Delta_\kappa = \min_{i\in\{1,\ldots,I-1\}} \kappa_i - \kappa_{i+1}$ and $\Delta_v = \min_{n\in\{1,\ldots,N-1\}} v_n - v_{n+1}$. Note that the earlier assumptions imply $\Delta_\kappa > 0$ and $\Delta_v > 0$.

m is large enough to do this. Thus, given a pure Nash equilibrium, it is well defined to let $g(i)$ denote the identity of the ith highest bidder.

Edelman, Ostrovsky and Schwarz (2007) defined an equilibrium to be *locally envy-free* if no bidder would prefer to exchange bids with the bidder that is one position higher. (See their paper for an intuitive argument for why such a refinement might be reasonable.) Thus, unlike our concept, the definition is in terms of the GSP auction itself and not in terms of trembles. To modify that refinement for this setting with a discrete bid set, what seems most in keeping with their original motivation is to slightly weaken it by instead requiring that no bidder would prefer to exchange bids with the bidder that is one position higher *and then to have the latter increase its bid by one increment*. Note that as m grows large (so that the discretization of the bid set becomes progressively fine), this definition converges to their original definition for a continuous bid set.

Definition 11. A pure equilibrium of the GSP auction $b = (b_1, \ldots, b_N)$ is a *locally envy-free** equilibrium if for all $i \in \{2, \ldots, N\}$, $\kappa_i[v_{g(i)} - b^{(i+1)}] \geq \kappa_{i-1}[v_{g(i)} - (b^{(i)} + \frac{1}{m})]$.

Under the aforementioned assumptions (e.g., that M and m are sufficiently large), we can then state the main result of this section:

Proposition 9. *Every pure extended proper equilibrium of the GSP auction is a locally envy-free* equilibrium.*[26]

What is more, as we argue below, this conclusion could not be obtained from proper equilibrium alone.

6.2 An Example

To illustrate and to sketch the proof of Proposition 9, consider the following specific example. There are three bidders, with per-click values $(v_1, v_2, v_3) = (4, 2, 1)$. There are three ad positions, with click rates $(\kappa_1, \kappa_2, \kappa_3) = (4, 2, 1)$. For the discretized bid set, it suffices to let $m = 6$ and $M = 3$.

The equilibrium $b^* = (\frac{5}{2}, 1, \frac{1}{2})$ is not locally envy-free*. This is because bidder 2 "envies" bidder 1 in the sense described above: bidder 2's equilibrium payoff is $\pi_2(\frac{5}{2}, 1, \frac{1}{2}) = 3$, but if bidder 2 exchanged bids with bidder 1 and bidder 1 then raised its bid by one increment, bidder 2 would obtain the higher payoff $\pi_2(\frac{7}{6}, \frac{5}{2}, \frac{1}{2}) = \frac{10}{3}$.

Nevertheless, b^* is a proper equilibrium of this example. Indeed, in Milgrom and Mollner (2018, Section 4.1) we argue that it is supported as a proper equilibrium by trembles in which 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 bidder 2 is much less likely to tremble than bidder 3. These trembles are, however, inconsistent with extended proper equilibrium because they require it to be more likely that bidder 3 trembles to $\frac{7}{6}$ (one increment above bidder 2's equilibrium bid of 1), which is not a best response to b^*, than that bidder 1 makes the same tremble, which is a best response.

What is more, b^* is not an extended proper equilibrium at all, because it cannot be supported by *any* trembles satisfying the necessary restrictions. Building on Proposition 8, suppose to the

[26] Milgrom and Mollner (2018, Section 3.2) show that essentially the same conclusion can be obtained from test-set equilibrium. Focusing on the case in which the allowable bids are the nonnegative reals, Theorem 6 in that paper states that every pure test-set equilibrium of the GSP auction is a locally envy-free equilibrium. However, there are examples of GSP auction games in which extended proper equilibrium rules out strategy profiles that test-set equilibrium does not (*cf.* Appendix C.3). This is, of course, in addition to the differences between the two solution concepts illustrated by the examples in Section 3.4.

contrary that ρ is an LPS that satisfies strong independence, has full support, and respects within-and-across-player preferences for which (ρ, b^*) is a lexicographic Nash equilibrium. To see that this produces a contradiction, suppose that bidder 2 contemplates raising its bid from $b_2^* = 1$ to $\frac{7}{6}$. This change makes a difference only if at least one of its opponents trembles within the set $\{1, \frac{7}{6}\}$. Thus, the relevant profiles can be partitioned into four sets:

- *Case 1:* $b_1 = \frac{7}{6}$ and $b_3 = \frac{1}{2}$
- *Case 2(a):* $b_1 = 1$ and $b_3 = \frac{1}{2}$
- *Case 2(b):* $b_1 = \frac{5}{2}$ and $b_3 \in \{1, \frac{7}{6}\}$
- *Case 3:* $b_1 \neq \frac{5}{2}$ and $b_3 \neq \frac{1}{2}$, with $b_1 \in \{1, \frac{7}{6}\}$ or $b_3 \in \{1, \frac{7}{6}\}$ or both

Against the profile in case 1, bidder 2 is strictly better off deviating to $\frac{7}{6}$:

$$\pi_2(\tfrac{7}{6}, \tfrac{7}{6}, \tfrac{1}{2}) = \tfrac{19}{6} > 3 = \pi_2(\tfrac{7}{6}, 1, \tfrac{1}{2}).$$

Moreover, each profile in the other cases is infinitely less likely under ρ than the profile in case 1. Indeed, for any profile in case 2(a) or 2(b), the deviating bidder is playing an inferior response to b^*.[27] In case 1, in contrast, only bidder 1 is deviating from b^*, and the deviation is to a best response to b^*. It therefore follows from ρ satisfying strong independence and respecting within-and-across-player preferences that for any (b_1, b_3) in case 2(a) or 2(b), $(\frac{7}{6}, \frac{1}{2}) >_\rho (b_1, b_3)$. Similarly, each profile in case 3 involves both bidders deviating from their equilibrium bids. For each such profile, there is a corresponding profile in one of the other cases that involves only a single bidder deviating, which, by strong independence, must be infinitely more likely under ρ. Altogether, we conclude that $\frac{7}{6}$ is a profitable deviation for bidder 2 against ρ, which contradicts (ρ, b^*) being a lexicographic Nash equilibrium. Moreover, the preceding argument can be generalized to prove Proposition 9.

7 Conclusion

We introduce a new refinement for games in normal form: extended proper equilibrium. Every finite game possesses at least one extended proper equilibrium. Like proper equilibrium, our new refinement is defined using trembles and imposes the following within-player restrictions on those trembles: in the limit, a player should be overwhelmingly less likely to tremble to a strategy that is a more costly mistake than to a different strategy that would be a less costly mistake. However, extended proper equilibrium strengthens that concept by adding further restrictions to the allowable trembles, which are precisely the consequences of combining proper equilibrium with a certain form of across-player symmetry in the propensity to tremble—namely, extended proper equilibrium describes the play induced by symmetrically proper strategies in the meta-game model in which players originate from a common pool. We also show that these definitions in terms of trembles have a corresponding characterization in terms of beliefs using lexicographic probability systems. In that formulation, extended proper equilibrium requires that, for some scaling of payoffs, costly mistakes should be regarded as infinitely less likely than less costly ones.

These additional restrictions have bite in games with three or more players, restricting what a player may believe about the relative likelihood of mistakes by two different opponents. In games

[27] If bidder 1 deviates from $b_1^* = \frac{5}{2}$ to $b_1 = 1$ as in case 2(a), then, assuming the others continue to play their equilibrium bids of $b_2^* = 1$ and $b_3^* = \frac{1}{2}$, bidder 1's allocation changes from always winning the highest position to winning the highest position only half the time. It can be shown that this is not a best response to b^*. For similar reasons, it is not a best response for bidder 3 to change the allocation by deviating to $b_3 \in \{1, \frac{7}{6}\}$ as in case 2(b).

with only two players, each player has just a single opponent, and so there are no such beliefs to be formed. Consistent with this intuition, extended proper equilibrium coincides with proper equilibrium in two-player games. Nevertheless, many games of economic interest involve more than two players, and in those settings, extended proper equilibrium may be a useful refinement of proper equilibrium. The GSP auction model (Edelman, Ostrovsky and Schwarz, 2007; Varian, 2007) is one such application.

References

Banks, Jeffrey S. and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 1987, *55* (3), 647–661.

Blume, Lawrence, Adam Brandenburger, and Eddie Dekel, "Lexicographic Probabilities and Equilibrium Refinements," *Econometrica*, 1991, *59* (1), 81–98.

Cho, In-Koo and David M. Kreps, "Signaling Games and Stable Equilibria," *The Quarterly Journal of Economics*, 1987, *102* (2), 179–221.

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.

Eilenberg, Samuel and Deane Montgomery, "Fixed Point Theorems for Multi-Valued Transformations," *American Journal of Mathematics*, 1946, *68* (2), 214–222.

Fudenberg, Drew and Kevin He, "Learning and Type Compatibility in Signalling Games," *Econometrica*, 2018, *86* (4), 1215–1255.

_ **and** _ , "Player-Compatible Equilibrium," *Working Paper*, 2020. https://arxiv.org/abs/1712.08954.

Jurado, I. García, "Nuevos Equilibrios Estables ante Tendencias al Error en Juegos no Cooperativos." PhD dissertation, Universidade de Santiago de Compostela, 1989. http://eio.usc.es/pub/io/xogos/MATERIALES/ThesisIgnacio.pdf.

_ , "Un Refinamiento del Concepto de Equilibrio Propio de Myerson," *Trabajos de Investigación Operativa*, 1989, *4* (1), 11–19.

_ **and J. M. Prada Sánchez**, "A Remark on Myerson's Concept of Proper Equilibrium," *International Journal of Game Theory*, 1990, *19* (1), 11–16.

Kakutani, Shizuo, "A Generalization of Brouwer's Fixed Point Theorem," *Duke Mathematical Journal*, 1941, *8* (3), 416–427.

Kalai, Ehud and Dov Samet, "Persistent Equilibria in Strategic Games," *International Journal of Game Theory*, 1984, *13* (3), 129–144.

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.

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

Maynard Smith, John and Geoffrey A. Parker, "The Logic of Asymmetric Contests," *Animal Behaviour*, 1976, *24* (1), 159–175.

McKelvey, Richard D. and Thomas R. Palfrey, "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, 1995, *10* (1), 6–38.

Milgrom, Paul and Joshua Mollner, "Equilibrium Selection in Auctions and High Stakes Games," *Econometrica*, 2018, *86* (1), 219–261.

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.

Nachbar, John H., "'Evolutionary' Selection Dynamics in Games: Convergence and Limit Properties," *International Journal of Game Theory*, 1990, *19* (1), 59–89.

Neary, Philip R., "The Class of Reasonable Equilibria," *Working Paper*, 2010. https://sites.google.com/site/prneary/docs/Reasonable.pdf.

Okada, Akira, "On Stability of Perfect Equilibrium Points," *International Journal of Game Theory*, 1981, *10* (2), 67–73.

Reny, Philip J., "On the Existence of Monotone Pure-Strategy Equilibria in Bayesian Games," *Econometrica*, 2011, *79* (2), 499–553.

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

_ , "A Note on Evolutionarily Stable Strategies in Asymmetric Animal Conflicts," *Journal of Theoretical Biology*, 1980, *84* (1), 93–101.

_ , "Evolutionary Stability in Extensive Two-Person Games," *Mathematical Social Sciences*, 1983, *5* (3), 269–363.

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.

_ , *Stability and Perfection of Nash Equilibria*, Springer-Verlag, 1991.

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

A Lemmas

A.1 Fixed Point Theorem

The proof of Theorem 3 requires a more powerful fixed point theorem than that of Kakutani (1941). We use the following corollary of the Eilenberg and Montgomery (1946) fixed point result.[28]

Lemma 10 (Eilenberg-Montgomery)**.** *Suppose that X is a nonempty, compact, and convex subset of a Euclidean space and that $F : X \rightrightarrows X$ is an upper-hemicontinuous, nonempty-valued, contractible-valued correspondence. Then F has a fixed point.*

A.2 Lexicographic Probability Systems

Here, we state and prove some technical lemmas concerning LPSs. The first such lemma is an immediate consequence of the definition of the "square operator" given by equation (1):

Lemma 11. *Suppose $\rho = (p^1, \dots, p^K)$ is an LPS on $\bar{S}$. Define $p_0 = \min_{k \in \{1,\dots,K\}} \min_{\bar{s} \in supp(p^k)} p^k(\bar{s})$. For any $J \subseteq \mathcal{N}$ and $s'_J \in \prod_{n \in J} S_n$, define $k' = \min\{k : p^k_J(s'_J) > 0\}$. For any $r \in (0,1)^{K-1}$,*

$$r^1 \cdots r^{k'-1}(1 - r^{k'})p_0 \leq (r\Box\rho)_J(s'_J) \leq r^1 \cdots r^{k'-1}.$$

Proof of Lemma 11. Because $(r\Box\rho)_J(s'_J) = \sum_{s \in \bar{S}: s_J = s'_J} r\Box\rho(s)$, it can be expressed as

$$\begin{aligned}
&(1-r^1)\sum_{s\in\bar{S}:s_J=s'_J} p^1(s) + r^1\left[(1-r^2)\sum_{s\in\bar{S}:s_J=s'_J} p^2(s) + r^2\left[(1-r^3)\sum_{s\in\bar{S}:s_J=s'_J} p^3(s)\right.\right.\\
&\quad\left.\left. + \; r^3\left[\cdots + r^{K-2}\left[(1-r^{K-1})\sum_{s\in\bar{S}:s_J=s'_J} p^{K-1}(s) + r^{K-1}\sum_{s\in\bar{S}:s_J=s'_J} p^K(s)\right]\cdots\right]\right]\right]\\
&= (1-r^1)p^1_J(s'_J) + r^1\left[(1-r^2)p^2_J(s'_J) + r^2\left[(1-r^3)p^3_J(s'_J)\right.\right.\\
&\quad\left.\left. + \; r^3\left[\cdots + r^{K-2}\left[(1-r^{K-1})p^{K-1}_J(s'_J) + r^{K-1}p^K_J(s'_J)\right]\cdots\right]\right]\right].
\end{aligned}$$

Using $p^k_J(s'_J) = 0$ for all $k < k'$, this simplifies to

$$r^1 \cdots r^{k'-1}\left[(1-r^{k'})p^{k'}_J(s'_J) + r^{k'}\left[\cdots + r^{K-2}\left[(1-r^{K-1})p^{K-1}_J(s'_J) + r^{K-1}p^K_J(s'_J)\right]\cdots\right]\right].$$

On one hand, we have for all $k \geq k'$ that $p^k_J(s'_J) \leq 1$. Plugging this in, we obtain $(r\Box\rho)_J(s'_J) \leq r^1 \cdots r^{k'-1}$. On the other hand, we have $p^{k'}_J(s'_J) \geq p_0$ and for all $k > k'$ that $p^k_J(s'_J) \geq 0$. Plugging these in, we obtain $(r\Box\rho)_J(s'_J) \geq r^1 \cdots r^{k'-1}(1-r^{k'})p_0$. $\square$

Lemma 12.[29] *Given any LPS $\rho = (p^1, \dots, p^K)$ on $\bar{S}$ and any $\alpha \in \mathbb{R}^N_{++}$, there exists a $\delta > 0$ such that if $r \in (0,\delta)^{K-1}$, then $(\forall l \in \mathcal{N})(\forall m \in \mathcal{N})(\forall s'_l \in S_l)(\forall s''_m \in S_m)(\forall s^*_l \in BR_l(\rho))(\forall s^*_m \in BR_m(\rho))$*

$$\alpha_l \max_{\hat{s}_l \in S_l} \sum_{s \in \bar{S}} (r\Box\rho)(s)[\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] > \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s \in \bar{S}} (r\Box\rho)(s)[\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)] \tag{2}$$

[28] See Reny (2011) for another economic application of the Eilenberg and Montgomery (1946) result.

[29] This result is related to Proposition 1 of BBD. The primary difference is that Lemma 12 expands focus from within-player comparisons to within-and-across-player comparisons.

if and only if

$$\left[\alpha_l\left(\pi_l(s_l^*|p^k)-\pi_l(s_l'|p^k)\right)\right]_{k=1}^K >_L \left[\alpha_m\left(\pi_m(s_m^*|p^k)-\pi_m(s_m''|p^k)\right)\right]_{k=1}^K. \tag{3}$$

Proof of Lemma 12. Fix an LPS $\rho=(p^1,\dots,p^K)$ and a scaling vector $\alpha\in\mathbb{R}_{++}^N$. Select for every $n\in\mathcal{N}$ some $s_n^*\in BR_n(\rho)$. It does not matter which element is chosen—this is because all elements of $BR_n(\rho)$ must yield identical expected payoffs against every LPS level p^k, and so the specific choice does not affect equation (3). We then define for each $k\in\{1,\dots,K\}$, $l,m\in\mathcal{N}$ (possibly equal), $s_l'\in S_l$, and $s_m''\in S_m$,

$$D_{lm}^k(s_l',s_m'')=\alpha_l\left[\pi_l(s_l^*|p^k)-\pi_l(s_l'|p^k)\right]-\alpha_m\left[\pi_m(s_m^*|p^k)-\pi_m(s_m''|p^k)\right].$$

If $D_{lm}^k(s_l',s_m'')=0$ for all (l,m,k,s_l',s_m''), then neither (2) nor (3) ever applies, and the statement holds vacuously with any δ. We therefore assume henceforth that this is not the case. Next, define

$$\begin{aligned}\overline{D}&=\max_{\substack{k\in\{1,\dots,K\}\\ l,m\in\mathcal{N}\\ s_l'\in S_l,s_m''\in S_m}}|D_{lm}^k(s_l',s_m'')|\\ \underline{D}&=\min_{\substack{k\in\{1,\dots,K\}\\ l,m\in\mathcal{N}\\ s_l'\in S_l,s_m''\in S_m}}\left\{D_{lm}^k(s_l',s_m''):D_{lm}^k(s_l',s_m'')>0\right\}.\end{aligned}$$

By the previous assumption, $\underline{D}$ is a well-defined positive number. And clearly, $\overline{D}\geq 0$. Set $\delta=\frac{\underline{D}}{\underline{D}+\overline{D}}>0$, and we show that the statement holds with this choice of δ. Suppose $r\in(0,\delta)^{K-1}$. Below, part 1 establishes the statement for the case of $l=m$, and part 2 proceeds to the general case.

Part 1: We begin by observing that for all $n\in\mathcal{N}$, $s_n'\in S_n$, and $s_n''\in S_n$,

$$\begin{aligned}&\alpha_n\max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n)-\pi_n(s/s_n')]-\alpha_n\max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n)-\pi_n(s/s_n'')]\\ &=-\alpha_n\sum_{s\in\bar{S}}(r\Box\rho)(s)\pi_n(s/s_n')+\alpha_n\sum_{s\in\bar{S}}(r\Box\rho)(s)\pi_n(s/s_n'')\\ &=\alpha_n\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/s_n^*)-\pi_n(s/s_n')]-\alpha_n\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/s_n^*)-\pi_n(s/s_n'')]\\ &=\sum_{s\in\bar{S}}(r\Box\rho)(s)\left(\alpha_n[\pi_n(s/s_n^*)-\pi_n(s/s_n')]-\alpha_n[\pi_n(s/s_n^*)-\pi_n(s/s_n'')]\right).\end{aligned}$$

This observation will be used in the two cases we now consider. For the first case, suppose that for some $n\in\mathcal{N}$, $s_n'\in S_n$, and $s_n''\in S_n$,

$$\left[\alpha_n\left(\pi_n(s_n^*|p^k)-\pi_n(s_n'|p^k)\right)\right]_{k=1}^K >_L \left[\alpha_n\left(\pi_n(s_n^*|p^k)-\pi_n(s_n''|p^k)\right)\right]_{k=1}^K.$$

Then there is some k such that both $D_{nn}^k(s_n',s_n'')\geq\underline{D}>0$ and for all $j<k$, $D_{nn}^j(s_n',s_n'')=0$. For

this k, let $\hat{\rho} = (p^{k+1}, \ldots, p^K)$ and $\hat{r} = (r^{k+1}, \ldots, r^{K-1})$. Then

$$\begin{aligned}
&\sum_{s\in\bar{S}}(r\Box\rho)(s)\left(\alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n')] - \alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n'')]\right) \\
&\quad = r^1 \cdots r^{k-1}\Bigg[(1-r^k)D_{nn}^k(s_n', s_n'') \\
&\qquad + r^k \sum_{s\in\bar{S}}(\hat{r}\Box\hat{\rho})(s)\left(\alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n')] - \alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n'')]\right)\Bigg] \\
&\quad \geq r^1 \cdots r^{k-1}\left[(1-r^k)\underline{D} - r^k\overline{D}\right] \\
&\quad > r^1 \cdots r^{k-1}\left[\underline{D} - \delta(\underline{D} + \overline{D})\right],
\end{aligned}$$

which is equal to zero because $\delta = \frac{\underline{D}}{\underline{D}+\overline{D}}$. We conclude that

$$\alpha_n \max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n) - \pi_n(s/s_n')] > \alpha_n \max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n) - \pi_n(s/s_n'')].$$

For the second case, suppose instead that for some $n \in \mathcal{N}$, $s_n' \in S_n$, and $s_n'' \in S_n$,

$$\left[\alpha_n\left(\pi_n(s_n^*|p^k) - \pi_n(s_n'|p^k)\right)\right]_{k=1}^K =_L \left[\alpha_n\left(\pi_n(s_n^*|p^k) - \pi_n(s_n''|p^k)\right)\right]_{k=1}^K.$$

Then $D_{nn}^k(s_n', s_n'') = 0$ for all k. Thus,

$$\sum_{s\in\bar{S}}(r\Box\rho)(s)\left(\alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n')] - \alpha_n[\pi_n(s/s_n^*) - \pi_n(s/s_n'')]\right),$$

as a weighted average of $\{D_{nn}^k(s_n', s_n'')\}_{k=1}^K$, is also equal to zero. We conclude that

$$\alpha_n \max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n) - \pi_n(s/s_n')] = \alpha_n \max_{\hat{s}_n\in S_n}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_n(s/\hat{s}_n) - \pi_n(s/s_n'')].$$

The first case establishes (3) $\Longrightarrow$ (2). By exchanging the roles of s_n' and s_n'', the two cases together establish the contrapositive of (2) $\Longrightarrow$ (3).

Part 2: We begin by observing that for all $l, m \in \mathcal{N}$, $s_l' \in S_l$, and $s_m'' \in S_m$,

$$\begin{aligned}
&\alpha_l \max_{\hat{s}_l\in S_l}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_l(s/\hat{s}_l) - \pi_l(s/s_l')] - \alpha_m \max_{\hat{s}_m\in S_m}\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_m(s/\hat{s}_m) - \pi_m(s/s_m'')] \\
&\quad = \alpha_l\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_l(s/s_l^*) - \pi_l(s/s_l')] - \alpha_m\sum_{s\in\bar{S}}(r\Box\rho)(s)[\pi_m(s/s_m^*) - \pi_m(s/s_m'')] \\
&\quad = \sum_{s\in\bar{S}}(r\Box\rho)(s)\left(\alpha_l[\pi_l(s/s_l^*) - \pi_l(s/s_l')] - \alpha_m[\pi_m(s/s_m^*) - \pi_m(s/s_m'')]\right),
\end{aligned}$$

where the first step follows from part 1 of this proof, which implies that any maximizers $\hat{s}_l$ and $\hat{s}_m$ must be elements of $BR_l(\rho)$ and $BR_m(\rho)$, respectively. This observation will be used in the two cases we now consider. For the first case, suppose that for some $l, m \in \mathcal{N}$, $s_l' \in S_l$, and $s_m'' \in S_m$,

$$\left[\alpha_l\left(\pi_l(s_l^*|p^k) - \pi_l(s_l'|p^k)\right)\right]_{k=1}^K >_L \left[\alpha_m\left(\pi_m(s_m^*|p^k) - \pi_m(s_m''|p^k)\right)\right]_{k=1}^K.$$

Then there is some k such that both $D^k_{lm}(s'_l, s''_m) \geq \underline{D} > 0$ and for all $j < k$, $D^j_{lm}(s'_l, s''_m) = 0$. For this k, let $\hat{\rho} = (p^{k+1}, \ldots, p^K)$ and $\hat{r} = (r^{k+1}, \ldots, r^{K-1})$. Then we have that

$$\begin{aligned}
&\sum_{s \in \bar{S}} (r \square \rho)(s) \left(\alpha_l [\pi_l(s/s^*_l) - \pi_l(s/s'_l)] - \alpha_m [\pi_m(s/s^*_m) - \pi_m(s/s''_m)] \right) \\
&\quad = r^1 \cdots r^{k-1} \Bigg[(1 - r^k) D^k_{lm}(s'_l, s''_m) \\
&\qquad + r^k \sum_{s \in \bar{S}} (\hat{r} \square \hat{\rho})(s) \left(\alpha_l [\pi_l(s/s^*_l) - \pi_l(s/s'_l)] - \alpha_m [\pi_m(s/s^*_m) - \pi_m(s/s''_m)] \right) \Bigg] \\
&\quad \geq r^1 \cdots r^{k-1} \left[(1 - r^k) \underline{D} - r^k \overline{D} \right] \\
&\quad > r^1 \cdots r^{k-1} \left[\underline{D} - \delta(\underline{D} + \overline{D}) \right],
\end{aligned}$$

which is equal to zero because $\delta = \frac{\underline{D}}{\underline{D} + \overline{D}}$. We conclude that

$$\alpha_l \max_{\hat{s}_l \in S_l} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] > \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)].$$

For the second case, suppose instead that for some $l, m \in \mathcal{N}$, $s'_l \in S_l$, and $s''_m \in S_m$,

$$\left[\alpha_l \left(\pi_l(s^*_l | p^k) - \pi_l(s'_l | p^k) \right) \right]_{k=1}^K =_L \left[\alpha_m \left(\pi_m(s^*_m | p^k) - \pi_m(s''_m | p^k) \right) \right]_{k=1}^K .$$

Then $D^k_{lm}(s'_l, s''_m) = 0$ for all k. Thus,

$$\sum_{s \in \bar{S}} (r \square \rho)(s) \left(\alpha_l [\pi_l(s/s^*_l) - \pi_l(s/s'_l)] - \alpha_m [\pi_m(s/s^*_m) - \pi_m(s/s''_m)] \right),$$

as a weighted average of $\{D^k_{lm}(s'_l, s''_m)\}_{k=1}^K$, is also equal to zero. We conclude that

$$\alpha_l \max_{\hat{s}_l \in S_l} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] = \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)].$$

The first case establishes (3) $\Longrightarrow$ (2). By exchanging the roles of s'_l and s''_m, the two cases together establish the contrapositive of (2) $\Longrightarrow$ (3). □

Lemma 13. *Suppose an LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ that respects within-and-across-player preferences and a strategy profile σ are such that (ρ, σ) is a lexicographic Nash equilibrium. For all $l, m \in \mathcal{N}$, all $s'_l \in S_l$, and all $s''_m \in S_m$, if $s'_l \notin BR_l(\sigma)$ and $s''_m \in BR_m(\sigma)$, then $s''_m >_\rho s'_l$.*

Proof of Lemma 13. Fix any $\alpha \in \mathbb{R}^N_{++}$, any $s^*_l \in BR_l(\rho)$, and any $s^*_m \in BR_m(\rho)$. Suppose $s'_l \in S_l$ and $s''_m \in S_m$ are such that $s'_l \notin BR_l(\sigma)$ and $s''_m \in BR_m(\sigma)$. Because (ρ, σ) is a lexicographic Nash equilibrium, $p^1(s) = \prod_{n \in \mathcal{N}} \sigma_n(s)$. Thus, $s'_l \notin BR_l(\sigma)$ implies that $\alpha_l \sum_{s \in \bar{S}} p^1(s) [\pi_l(s/s^*_l) - \pi_l(s/s'_l)] > 0$. On the other hand, $s''_m \in BR_m(\sigma)$ implies that $\alpha_m \sum_{s \in \bar{S}} p^1(s) [\pi_m(s/s^*_m) - \pi_m(s/s'_m)] = 0$. This implies

$$\left[\alpha_l \left(\pi_l(s^*_l | p^k) - \pi_l(s'_l | p^k) \right) \right]_{k=1}^K >_L \left[\alpha_m \left(\pi_m(s^*_m | p^k) - \pi_m(s''_m | p^k) \right) \right]_{k=1}^K .$$

Because ρ respects within-and-across-player preferences, $s''_m >_\rho s'_l$. □

Lemma 14. *Suppose an LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ and a strategy profile σ are such that (ρ, σ) is a lexicographic Nash equilibrium. For all $J \subseteq \mathcal{N}$, all $I \subseteq \mathcal{N}$, all $s'_J \in \prod_{n \in J} S_n$, and all $s''_I \in \prod_{n \in I} S_n$, if $\sigma_n(s''_n) > 0$ for all $n \in I$, then*

(i) $s''_I \geq_\rho s'_J$; and

(ii) if in addition, $\sigma_m(s'_m) = 0$ for some $m \in J$, then $s''_I >_\rho s'_J$.

Proof of Lemma 14. *Claim (i).* Suppose $s''_I \in \prod_{n \in I} S_n$ is such that $\sigma_n(s''_n) > 0$ for all $n \in I$. Because (ρ, σ) is a lexicographic Nash equilibrium, $p^1_I(s''_I) = \prod_{n \in I} \sigma_n(s''_n) > 0$. Therefore, $s''_I \geq_\rho s'_J$, regardless of the choice of $s'_J \in \prod_{n \in J} S_n$.

Claim (ii). Suppose in addition that $s'_J \in \prod_{n \in J} S_n$ is such that $\sigma_m(s'_m) = 0$ for some $m \in J$. Thus, $p^1_J(s'_J) = \prod_{n \in J} \sigma_n(s'_n) = 0$. On the other hand, we had $p^1_I(s''_I) > 0$. Therefore, $s''_I >_\rho s'_J$. $\square$

Lemma 15. *Suppose an LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ is equivalent to the $\mathbb{F}$-valued probability distribution $\hat{p}$ on $\bar{S}$. For all $J \subseteq \mathcal{N}$, $I \subseteq \mathcal{N}$, $s_J \in \prod_{n \in J} S_j$, and $s_I \in \prod_{n \in I} S_i$: $s_J \geq_\rho s_I$ iff there exists an $N \in \mathbb{N}$ such that $\frac{\hat{p}_J(s_J)}{\hat{p}_I(s_I)} \geq \frac{1}{N}$.*

Proof of Lemma 15. Define

$$p_0 = \min_{k \in \{1, \ldots, K\}} \min_{\bar{s} \in \operatorname{supp}(p^k)} p^k(\bar{s}).$$

Choose $N_0 \in \mathbb{N}$ to be such that $p_0 > \frac{1}{N_0}$. By what it means for ρ and $\hat{p}$ to be equivalent, there exists a vector of infinitesimals $\varepsilon = (\varepsilon_1, \ldots \varepsilon_{K-1}) \in \mathbb{F}^{K-1}$ such that $\hat{p} = \varepsilon \square \rho$. Suppose $s_J \geq_\rho s_I$, so that $k := \min\{j : p^j_J(s_J) > 0\} \leq \min\{j : p^j_I(s_I) > 0\}$. Then Lemma 11 implies both $\hat{p}_J(s_J) \geq \varepsilon_1 \cdots \varepsilon_{k-1}(1 - \varepsilon_k) p_0$, and $\hat{p}_I(s_I) \leq \varepsilon_1 \cdots \varepsilon_{k-1}(1 - \varepsilon_k)$. Therefore $\frac{\hat{p}_J(s_J)}{\hat{p}_I(s_I)} \geq p_0 > \frac{1}{N_0}$.

Suppose instead that $s_J <_\rho s_I$, so that $k := \min\{j : p^j_J(s_J) > 0\} > \min\{j : p^j_I(s_I) > 0\}$. Then $\hat{p}_J(s_J) \leq \varepsilon_1 \cdots \varepsilon_{k-1}(1 - \varepsilon_k)$ and $\hat{p}_I(s_I) \geq \varepsilon_1 \cdots \varepsilon_{k-2}(1 - \varepsilon_{k-1}) p_0$. Consequently,

$$\frac{\hat{p}_J(s_J)}{\hat{p}_I(s_I)} \leq \frac{\varepsilon_{k-1}(1 - \varepsilon_k)}{(1 - \varepsilon_{k-1}) p_0},$$

which is an infinitesimal and less than $\frac{1}{N}$ for all $N \in \mathbb{N}$. $\square$

Lemma 16. *Suppose an LPS $\rho = (p^1, \ldots, p^K)$ satisfies strong independence. Let I and J be disjoint subsets of $\mathcal{N}$. Similarly, let I' and J' be disjoint subsets of $\mathcal{N}$. For all $s_I \in \prod_{n \in I} S_n$, $s_J \in \prod_{n \in J} S_n$, $s_{I'} \in \prod_{n \in I'} S_n$, and $s_{J'} \in \prod_{n \in j'} S_n$, if $s_I >_\rho s_{I'}$ and $s_J \geq_\rho s_{J'}$ then $(s_I, s_J) >_\rho (s_{I'}, s_{J'})$.*

Proof of Lemma 16. Because ρ satisfies strong independence, it is equivalent to some $\mathbb{F}$-valued probability distribution $\hat{p}$ on $\bar{S}$ that is a product distribution. Suppose $s_I >_\rho s_{I'}$ and $s_J \geq_\rho s_{J'}$. Then by Lemma 15, we have that $\forall N \in \mathbb{N}$, $\frac{\hat{p}_{I'}(s_{I'})}{\hat{p}_I(s_I)} < \frac{1}{N}$. Similarly, $\exists N_0 \in \mathbb{N}$ such that $\frac{\hat{p}_J(s_J)}{\hat{p}_{J'}(s_{J'})} > \frac{1}{N_0}$. Equivalently, $\frac{\hat{p}_{J'}(s_{J'})}{\hat{p}_J(s_J)} < N_0$. Because $\hat{p}$ is a product distribution,

$$\frac{\hat{p}_{I',J'}(s_{I'}, s_{J'})}{\hat{p}_{I,J}(s_I, s_J)} = \frac{\hat{p}_{I'}(s_{I'})}{\hat{p}_I(s_I)} \frac{\hat{p}_{J'}(s_{J'})}{\hat{p}_J(s_J)} < \frac{N_0}{N} \text{ for all } N \in \mathbb{N},$$

which implies that $\frac{\hat{p}_{I',J'}(s_{I'}, s_{J'})}{\hat{p}_{I,J}(s_I, s_J)} < \frac{1}{N}$ for all $N \in \mathbb{N}$. Another application of Lemma 15 gives the desired result, $(s_I, s_J) >_\rho (s_{I'}, s_{J'})$. $\square$

A.3 GSP Auction

Lemma 17. *Let $\Delta_\kappa = \min_{i\in\{1,\dots,I-1\}} \kappa_i - \kappa_{i+1}$ and $\Delta_v = \min_{n\in\{1,\dots,N-1\}} v_n - v_{n+1}$. Assume $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$. If $b = (b_1, \dots, b_N)$ is a pure Nash equilibrium of the GSP auction, then for all $i \in \{1, \dots, N-1\}$, $b^{(i)} > b^{(i+1)}$.*

Proof of Lemma 17. Suppose, by way of contradiction, that $i' \in \{1, \dots, N-1\}$ and $i'' > i'$ are such that $b^{(i')} = \dots = b^{(i'')} = b^*$ is a maximal tie. First, note that any bidder with per-click value v who is part of the tie earns the payoff

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

If such a bidder raises its bid to $b^* + \frac{1}{m}$, that leads to a payoff of at least

$$\kappa_{i'}\left[v - \left(b^* + \frac{1}{m}\right)\right].$$ [30,31]

The rest of the proof consists of two parts. First, we argue that $b^{(i''+1)} < b^* - \frac{1}{m}$. Second, we argue that at least one of the bidders who are part of the tie has a profitable deviation.

Part 1: If $b^{(i''+1)} \geq b^* - \frac{1}{m}$,[32] then a bidder with per-click value v who is part of the tie earns a payoff that is at most

$$\frac{1}{i''-i'+1}\left[\frac{\kappa_{i''}}{m} + \sum_{i=i'}^{i''} \kappa_i (v - b^*)\right].$$

Every such bidder must have $v \geq b^* - \frac{1}{m}$, or else its payoff would be negative, and it could profitably deviate to a bid of zero. There therefore exists a bidder that is part of the tie with a per-click value $v \geq b^* - \frac{1}{m} + \Delta_v$. We argue that it would be profitable for the latter bidder to deviate to $b^* + \frac{1}{m}$:

$$\begin{aligned}
&\kappa_{i'}\left[v - \left(b^* + \frac{1}{m}\right)\right] - \frac{1}{i''-i'+1}\left[\frac{\kappa_{i''}}{m} + \sum_{i=i'}^{i''} \kappa_i (v - b^*)\right] \\
&= \frac{1}{i''-i'+1}\left[\sum_{i=i'+1}^{i''} (\kappa_{i'} - \kappa_i)(v - b^*)\right] - \frac{1}{m}\left(\kappa_{i'} + \frac{\kappa_{i''}}{i''-i'+1}\right) \\
&\geq \frac{1}{i''-i'+1}\left[(i''-i')\Delta_\kappa\left(\Delta_v - \frac{1}{m}\right)\right] - \frac{\kappa_1}{m}\left(\frac{i''-i'+2}{i''-i'+1}\right) \\
&= \frac{1}{i''-i'+1}\left[(i''-i')\Delta_\kappa\Delta_v - \frac{(i''-i')\Delta_\kappa + \kappa_1(i''-i'+2)}{m}\right] \\
&\geq \frac{i''-i'}{i''-i'+1}\left[\Delta_\kappa\Delta_v - \frac{4\kappa_1}{m}\right],
\end{aligned}$$

which is strictly positive under the assumption that $m > \frac{4\kappa_1}{\Delta_\kappa\Delta_v}$. In the above, the first and third steps are algebra. The second step uses $\kappa_{i'} - \kappa_i \geq \Delta_\kappa$ for all $i > i'$, $v \geq b^* - \frac{1}{m} + \Delta_v$, $\kappa_{i''} \leq \kappa_1$, and $\kappa_{i'} \leq \kappa_1$. The fourth step uses $\Delta_\kappa \leq \kappa_1$ and $2i'' - 2i' + 2 \leq 4(i'' - i')$.

[30] If $b^{(i'-1)} > b^* + \frac{1}{m}$, then the payoff is exactly $\kappa_{i'}(v - b^*)$. If $b^{(i'-1)} = b^* + \frac{1}{m}$, then the payoff is a lottery between that and terms of the form $\kappa_i\left[v - \left(b^* + \frac{1}{m}\right)\right]$ for $i < i'$. In either case, the bound applies.

[31] It is true that this deviation to $b^* + \frac{1}{m}$ is not possible if $b^* = M$. But the deviation is necessary for the subsequent arguments only when it is profitable. And it cannot be profitable if $b^* = M$, by the assumption that $M > v_1$.

[32] We cannot have $b^{(i''+1)} > b^*$. But we can have $b^{(i''+1)} = b^*$ in the case where $i'' = N$ and $b^* = 0$. (Recall that we have defined $b^{(N+1)} = 0$.)

Part 2: From the previous part of the proof, $b^{(i''+1)} < b^* - \frac{1}{m}$. Therefore, a bidder with per-click value v who is part of the tie can reduce its bid to $b^* - \frac{1}{m}$ and earn the payoff

$$\kappa_{i''}(v - b^{(i''+1)}).$$

This would be a profitable deviation for any such bidder with per-click value $v < b^*$. We therefore assume henceforth that $v \geq b^*$ for all such bidders. For the following arguments, we use the fact that the following is an upper bound on such a bidder's payoff from not deviating:

$$\frac{1}{i'' - i' + 1}\left[\kappa_{i''}(v - b^{(i''+1)}) + (i'' - i')\kappa_{i'}(v - b^*) - (i'' - i' - 1)\Delta_\kappa(v - b^*)\right].$$

Let $v' > v''$ be the per-click values of two of the tied bidders. Because the bidder with value v' does not find it profitable to deviate to $b^* + \frac{1}{m}$,

$$\frac{1}{i'' - i' + 1}\left[\kappa_{i''}(v' - b^{(i''+1)}) + (i'' - i')\kappa_{i'}(v' - b^*) - (i'' - i' - 1)\Delta_\kappa(v' - b^*)\right] \geq \kappa_{i'}\left[v - \left(b^* + \frac{1}{m}\right)\right]$$
$$\implies \kappa_{i''}(v' - b^{(i''+1)}) \geq \kappa_{i'}(v' - b^*) + (i'' - i' - 1)\Delta_\kappa(v' - b^*) - (i'' - i' + 1)\frac{\kappa_{i'}}{m},$$

and therefore,

$$\begin{aligned}
\kappa_{i''}(v' - b^{(i''+1)}) &\geq \kappa_{i'}(v' - b^*) + (i'' - i' - 1)\Delta_\kappa\Delta_v - (i'' - i' + 1)\frac{\kappa_1}{m} \\
&= \kappa_{i'}(v' - b^*) + (i'' - i' - 1)\left(\Delta_\kappa\Delta_v - \frac{\kappa_1}{m}\right) - \frac{2\kappa_1}{m} \\
&\geq \kappa_{i'}(v' - b^*) - \frac{2\kappa_1}{m}, \qquad (4)
\end{aligned}$$

where the first step uses $\kappa_{i'} \leq \kappa_1$ and $v' - b^* \geq \Delta_v$, which follows from $v' \geq v'' \geq b^*$; the second step is algebra; and the third step uses $m > \frac{4\kappa_1}{\Delta_\kappa\Delta_v}$. Because the bidder with value v'' does not find it profitable to deviate to $b^* - \frac{1}{m}$,

$$\frac{1}{i'' - i' + 1}\left[\kappa_{i''}(v'' - b^{(i''+1)}) + (i'' - i')\kappa_{i'}(v'' - b^*) - (i'' - i' - 1)\Delta_\kappa(v'' - b^*)\right] \geq \kappa_{i''}(v'' - b^{(i''+1)})$$
$$\implies \kappa_{i'}(v'' - b^*) \geq \kappa_{i''}(v'' - b^{(i''+1)}) + \frac{(i'' - i' - 1)}{i'' - i'}\Delta_\kappa(v'' - b^*),$$

and therefore,

$$\kappa_{i'}(v'' - b^*) \geq \kappa_{i''}(v'' - b^{(i''+1)}), \qquad (5)$$

because $v'' \geq b^*$. Adding (4) and (5), then canceling like terms, we obtain

$$(\kappa_{i'} - \kappa_{i''})(v'' - v') \geq -\frac{2\kappa_1}{m},$$

but this is a contradiction because $\kappa_{i'} - \kappa_{i''} \geq \Delta_\kappa$, $v'' - v' \leq -\Delta_v$, and $m > \frac{4\kappa_1}{\Delta_\kappa\Delta_v}$. □

B Omitted Proofs

B.1 Proofs Corresponding to Section 3

Proof of Theorem 2. Fix any $\alpha \in \mathbb{R}^N_{++}$ and any $\varepsilon \in (0,1)$. Let $M = \sum_{n\in\mathcal{N}} |S_n|$ and $\delta = \frac{\varepsilon^M}{M}$. For each player $n \in \mathcal{N}$, define $\Delta^*_n = \{\sigma_n \in \Delta_n \,|\, (\forall s_n \in S_n)\, \sigma_n(s_n) \geq \delta\}$, which is a nonempty and compact subset of Δ^0_n. We also define the correspondence $F : \prod_{n\in\mathcal{N}} \Delta^*_n \rightrightarrows \prod_{n\in\mathcal{N}} \Delta^*_n$ as

$$F(\sigma) = \left\{ \hat{\sigma} \in \prod_{n\in\mathcal{N}} \Delta^*_n \;\middle|\; \begin{array}{l} (\forall l \in \mathcal{N})(\forall m \in \mathcal{N})(\forall s'_l \in S_l)(\forall s''_m \in S''_m) \\ \text{if } \alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m) \\ \text{then } \hat{\sigma}_l(s'_l) \leq \varepsilon \cdot \hat{\sigma}_m(s''_m) \end{array} \right\}.$$

Claim: For all $\sigma \in \prod_{n\in\mathcal{N}} \Delta^*_n$, $F(\sigma)$ is closed, convex, and nonempty.

Proof of Claim: The points in $F(\sigma)$ are those that satisfy a finite collection of linear inequalities, so $F(\sigma)$ is a closed, convex set. We next demonstrate that $F(\sigma)$ is nonempty. For all $n \in \mathcal{N}$ and all $s_n \in S_n$, define

$$\phi_n(s_n) = \#\{m \in \mathcal{N}, s'_m \in S_m \mid \alpha_n L_n(\sigma/s_n) > \alpha_m L_m(\sigma/s'_m)\}.$$

Define also $\Phi^0_n = \#\{s_n \in S_n \,|\, \phi_n(s_n) = 0\}$ and

$$\bar{\sigma}_n(s_n) = \begin{cases} \dfrac{\varepsilon^{\phi_n(s_n)}}{M} & \text{if } \phi_n(s_n) > 0 \\ \dfrac{1}{\Phi^0_n}\left(1 - \displaystyle\sum_{\hat{s}_n:\phi_n(\hat{s}_n)>0} \dfrac{\varepsilon^{\phi_n(\hat{s}_n)}}{M}\right) & \text{if } \phi_n(s_n) = 0 \end{cases}$$

To verify that $\bar{\sigma} \in \prod_{n\in\mathcal{N}} \Delta^*_n$, let $n \in \mathcal{N}$ and $s_n \in S_n$. If $\phi_n(s_n) > 0$, then

$$\bar{\sigma}_n(s_n) = \frac{\varepsilon^{\phi_n(s_n)}}{M} \geq \frac{\varepsilon^M}{M} = \delta.$$

And if $\phi_n(s_n) = 0$, then we also have

$$\bar{\sigma}_n(s_n) = \frac{1}{\Phi^0_n}\left(1 - \sum_{\hat{s}_n:\phi_n(\hat{s}_n)>0} \frac{\varepsilon^{\phi_n(s_n)}}{M}\right) \geq \frac{1}{\Phi^0_n}\left(1 - (M - \Phi^0_n)\frac{1}{M}\right) = \frac{1}{M} \geq \delta.$$

To verify that $\bar{\sigma} \in F(\sigma)$, suppose $l, m \in \mathcal{N}$, $s'_l \in S_l$, and $s''_m \in S_m$ are such that $\alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m)$. This implies that $\phi_l(s'_l) \geq \phi_m(s''_m) + 1$. If $\phi_m(s''_m) > 0$, then

$$\bar{\sigma}_l(s'_l) = \frac{\varepsilon^{\phi_l(s'_l)}}{M} \leq \frac{\varepsilon^{\phi_m(s''_m)+1}}{M} = \varepsilon \cdot \bar{\sigma}_m(s''_m).$$

And if $\phi_m(s''_m) = 0$, then we also have

$$\bar{\sigma}_l(s'_l) = \frac{\varepsilon^{\phi_l(s'_l)}}{M} \leq \frac{\varepsilon}{M} \leq \varepsilon \cdot \bar{\sigma}_m(s''_m),$$

where the last step uses the fact, established above, that $\phi_m(s''_m) = 0$ implies $\bar{\sigma}_m(s''_m) \geq \frac{1}{M}$.

Claim: F is upper-hemicontinuous.

Proof of Claim: This follows from continuity of $L_n(\cdot)$ for each $n \in \mathcal{N}$.

Claim: There exists an (α, ε)-extended proper equilibrium.

Proof of Claim: By the above claims, F satisfies all the conditions of the Kakutani Fixed Point Theorem (Kakutani, 1941), so there exists some $\sigma^\varepsilon \in \prod_{n\in\mathcal{N}} \Delta_n^*$ such that $\sigma^\varepsilon \in F(\sigma^\varepsilon)$. It follows from the definition of F that this fixed point is an (α, ε)-extended proper equilibrium.

Claim: There exists an extended proper equilibrium.

Proof of Claim: For all $t \in \mathbb{N}$, define $\varepsilon_t = \frac{1}{t+1}$. Applying the conclusion of the previous claim, there exists a sequence $(\sigma^t)_{t=1}^\infty$, where each σ^t is an (α, ε_t)-extended proper equilibrium. Because $\prod_{n\in\mathcal{N}} \Delta_n$ is a compact set, there exists a convergent subsequence, the limit of which is an extended proper equilibrium. $\square$

Proof of Theorem 3. Let $\hat{\sigma}$ be a proper equilibrium. Using the characterization of proper equilibrium given by Proposition 8 of BBD (and restated as Proposition 7 of this paper), there exists some LPS $\rho = (p^1, \ldots, p^K)$ on $S_1 \times S_2$ that satisfies strong independence, has full support, respects within-player preferences, and for which $(\rho, \hat{\sigma})$ is a lexicographic Nash equilibrium. For $n \in \{1, 2\}$ and $k \in \{1, \ldots, K\}$, define

$$\zeta_n^k = \{s_n \in S_n : p_n^k(s_n) > 0 \text{ and } p_n^j(s_n) = 0 \text{ for } j < k\}.$$

Because ρ has full support, $\{\zeta_n^1, \ldots, \zeta_n^K\}$ is a partition of S_n. Because ρ respects within-player preferences, $(\forall n \in \{1, 2\})(\forall k \in \{1, \ldots, K\})(\forall s_n' \in \zeta_n^k)(\forall s_n'' \in \cup_{j=k}^K \zeta_n^j)$:

$$\left[\pi_n(s_n'|p^k)\right]_{k=1}^K \geq_L \left[\pi_n(s_n''|p^k)\right]_{k=1}^K.$$

That is, each $s_n' \in \zeta_n^k$ is optimal for player n against ρ among strategies in $\cup_{j=k}^K \zeta_n^j$.

Let δ be as in Lemma 12. Fix any $\varepsilon \in (0, \delta)$ and any $\alpha \in \mathbb{R}_{++}^2$. Define $R = [\varepsilon^{2K^2-3K+2}, \varepsilon]^{K-1}$. Let $r \in R$. Because $\varepsilon < \delta$, we have, by Lemma 12, that $(\forall n \in \{1, 2\})(\forall k \in \{1, \ldots, K\})(\forall s_n' \in \zeta_n^k)(\forall s_n'' \in \cup_{j=k}^K \zeta_n^j)$:

$$\sum_{s\in\bar{S}} (r\Box\rho)(s)\pi_n(s/s_n') \geq \sum_{s\in\bar{S}} (r\Box\rho)(s)\pi_n(s/s_n'').$$

In particular, if $s_n', s_n'' \in \zeta_n^k$, then by symmetry we must have

$$\sum_{s\in\bar{S}} (r\Box\rho)(s)\pi_n(s/s_n') = \sum_{s\in\bar{S}} (r\Box\rho)(s)\pi_n(s/s_n'').$$

A consequence of the above is that there exist, for each player $n \in \{1, 2\}$, a nonincreasing sequence $\pi_n^1(r) \geq \cdots \geq \pi_n^K(r)$ such that $(\forall k \in \{1, \ldots, K\})(\forall s_n \in \zeta_n^k) : \sum_{s\in\bar{S}} (r\Box\rho)(s)\pi_n(s/s_n') = \pi_n^k(r)$. We then define $L_n^k(r) = \alpha_n[\pi_n^1(r) - \pi_n^k(r)]$, so that $0 = L_n^1(r) \leq \cdots \leq L_n^K(r)$. The interpretation of $L_n^k(r)$ is player n's scaled loss (scaled according to α_n) from using a strategy in ζ_n^k when the opponent's play is according to $r\Box\rho$ rather than a best response (i.e., a strategy in ζ_n^1).

Using this notation, we define the correspondence $F : R \times R \rightrightarrows R \times R$ by

$$F(r_1, r_2) = \left\{ (\hat{r}_1, \hat{r}_2) \in R \times R \;\middle|\; \begin{array}{l} (\forall k' \in \{2, \ldots, K\})(\forall k'' \in \{1, \ldots, K\}) : \\ \text{if } L_1^{k'}(r_2) > L_2^{k''}(r_1) \text{ then } \hat{r}_1^1 \cdots \hat{r}_1^{k'-1} \leq \varepsilon \cdot \hat{r}_2^1 \cdots \hat{r}_2^{k''-1} \\ \text{if } L_2^{k'}(r_1) > L_1^{k''}(r_2) \text{ then } \hat{r}_2^1 \cdots \hat{r}_2^{k'-1} \leq \varepsilon \cdot \hat{r}_1^1 \cdots \hat{r}_1^{k''-1} \end{array} \right\}.$$

To accommodate the case of $k''=1$ in this definition, we use the convention that the empty product $\hat{r}_1^1 \cdots \hat{r}_1^{k''-1}$ equals one.

Claim: For all $(r_1, r_2) \in R \times R$, $F(r_1, r_2)$ is nonempty.

Proof of Claim: Fix (r_1, r_2). For all $n \in \{1,2\}$ and all $k \in \{1, \dots, K\}$, define

$$\phi_n^k = \#\{m \in \{1,2\}, j \in \{1, \dots, K\} \mid L_n^k(r_{-n}) > L_m^j(r_{-m})\}.$$[33]

Note that, by construction, $L_n^1(r_{-n}) = 0$, and so $\phi_n^1 = 0$. Next, for all $n \in \{1,2\}$ and all $k \in \{1, \dots, K-1\}$, define

$$\bar{r}_n^k = \varepsilon^{1+(K-1)(\phi_n^{k+1} - \phi_n^k)}.$$

Because $L_n^{k+1}(r_{-n}) \geq L_n^k(r_{-n})$, we have $\phi_n^{k+1} \geq \phi_n^k$, and so $\bar{r}_n^k \leq \varepsilon$. We also have $\phi_n^{k+1} \leq 2K-1$, and so $\bar{r}_n^k = \varepsilon^{1+(K-1)(\phi_n^{k+1}-\phi_n^k)} \geq \varepsilon^{1+(K-1)(2K-1)} = \varepsilon^{2K^2-3K+2}$. From the previous two inequalities, we conclude $(\bar{r}_1, \bar{r}_2) \in R \times R$. Next, we verify that $(\bar{r}_1, \bar{r}_2) \in F(r_1, r_2)$. To that end, it is without loss of generality to suppose $k' \in \{2, \dots, K\}$ and $k'' \in \{1, \dots, K\}$ are such that $L_1^{k'}(r_2) > L_2^{k''}(r_1)$. This implies $\phi_1^{k'} \geq \phi_2^{k''} + 1$. We then have the following:

$$\begin{aligned}
\bar{r}_1^1 \cdots \bar{r}_1^{k'-1} &= \prod_{k=1}^{k'-1} \varepsilon^{1+(K-1)(\phi_1^{k+1}-\phi_1^k)} \\
&= \varepsilon^{k'-1+(K-1)\phi_1^{k'}} \\
&\leq \varepsilon^{k'-1+(K-1)(\phi_2^{k''}+1)} \\
&= \varepsilon \cdot \varepsilon^{[k'-2+K-k'']+[k''-1+(K-1)\phi_2^{k''}]} \\
&\leq \varepsilon \cdot \varepsilon^{k''-1+(K-1)\phi_2^{k''}} \\
&= \varepsilon \cdot \prod_{k=1}^{k''-1} \varepsilon^{1+(K-1)(\phi_2^{k+1}-\phi_2^k)} \\
&= \varepsilon \cdot \bar{r}_2^1 \cdots \bar{r}_2^{k''-1},
\end{aligned}$$

as required. In the above, the first and seventh (i.e., last) steps follow from the definition of $\bar{r}_n^k$. The second and sixth steps use $\phi_n^1 = 0$. The third step uses $\phi_1^{k'} \geq \phi_2^{k''} + 1$. The fourth step is algebra. The fifth step uses $k' \geq 2$ and $k'' \leq K$.

Claim: For all $(r_1, r_2) \in R \times R$, $F(r_1, r_2)$ is contractible.

Proof of Claim: Fix (r_1, r_2). Fix some $(\bar{r}_1, \bar{r}_2) \in F(r_1, r_2)$. Define $H : F(r_1, r_2) \times [0,1] \to F(r_1, r_2)$ in the following way. For any $\hat{r} = (\hat{r}_1, \hat{r}_2) \in F(r_1, r_2)$, define

$$H(\hat{r}, t) = \left(\left((\hat{r}_1^k)^{1-t} (\bar{r}_1^k)^t \right)_{k=1}^{K-1}, \left((\hat{r}_2^k)^{1-t} (\bar{r}_2^k)^t \right)_{k=1}^{K-1} \right).$$

It follows from this definition that H is continuous, that $H(\hat{r}, 0) = \hat{r}$, and that $H(\hat{r}, 1) = \bar{r}$. To establish that H is a contracting homotopy and that $F(r_1, r_2)$ is therefore contractible, it only remains to show that $H(\hat{r}, t) \in F(r_1, r_2)$.

Fix any $\hat{r} \in F(r_1, r_2)$ and $t \in [0,1]$. Because $\hat{r}, \bar{r} \in R \times R$, we have $\hat{r}_n^k, \bar{r}_n^k \in [\varepsilon^{2K^2-3K+2}, \varepsilon]$ for all $n \in \{1,2\}$ and $k \in \{1, \dots, K-1\}$. This implies $(\hat{r}_n^k)^{1-t}(\bar{r}_n^k)^t \in [\varepsilon^{2K^2-3K+2}, \varepsilon]$, and therefore

[33] Throughout, r_{-n} denotes the following: when $n=1$ it denotes r_2, and when $n=2$ it denotes r_1.

$H(\hat{r},t) \in R \times R$. To finish the argument, it is without loss of generality to suppose $k' \in \{2,\ldots,K\}$ and $k'' \in \{1,\ldots,K\}$ are such that $L_1^{k'}(r_2) > L_2^{k''}(r_1)$. Because $\hat{r}, \bar{r} \in F(r_1,r_2)$, we have $\hat{r}_1^1 \cdots \hat{r}_1^{k'-1} \leq \varepsilon \cdot \hat{r}_2^1 \cdots \hat{r}_2^{k''-1}$ and $\bar{r}_1^1 \cdots \bar{r}_1^{k'-1} \leq \varepsilon \cdot \bar{r}_2^1 \cdots \bar{r}_2^{k''-1}$ We then have the following:

$$\begin{aligned}
\big((\hat{r}_1^1)^{1-t}(\bar{r}_1^1)^t\big) \cdots \big((\hat{r}_1^{k'-1})^{1-t}(\bar{r}_1^{k'-1})^t\big) &= \big(\hat{r}_1^1 \cdots \hat{r}_1^{k'-1}\big)^{1-t}\big(\bar{r}_1^1 \cdots \bar{r}_1^{k'-1}\big)^t \\
&\leq \big(\varepsilon \cdot \hat{r}_2^1 \cdots \hat{r}_2^{k''-1}\big)^{1-t}\big(\varepsilon \cdot \bar{r}_2^1 \cdots \bar{r}_2^{k''-1}\big)^t \\
&= \varepsilon \cdot \big((\hat{r}_2^1)^{1-t}(\bar{r}_2^1)^t\big) \cdots \big((\hat{r}_2^{k''-1})^{1-t}(\bar{r}_2^{k''-1})^t\big),
\end{aligned}$$

as required.

Claim: F is upper-hemicontinuous.

Proof of Claim: This follows from continuity of $L_n^k(\cdot)$ for each $n \in \{1,2\}$ and $k \in \{1,\ldots,K\}$. Formally, let $r(t)$ and $\hat{r}(t)$ be convergent sequences in $R \times R$ with $\lim_{t\to\infty} r(t) = r$, $\lim_{t\to\infty} \hat{r}(t) = \hat{r}$, and $\hat{r}(t) \in F(r(t))$ for all t. We demonstrate that $\hat{r} \in F(r)$. It is without loss of generality to suppose $k' \in \{2,\ldots,K\}$ and $k'' \in \{1,\ldots,K\}$ are such that $L_1^{k'}(r_2) > L_2^{k''}(r_1)$. By continuity, we have $L_1^{k'}(r_2(t)) > L_2^{k''}(r_1(t))$ for all sufficiently large t. Because $\hat{r}(t) \in F(r(t))$, we therefore have $\hat{r}_1^1(t) \cdots \hat{r}_1^{k'-1}(t) \leq \varepsilon \cdot \hat{r}_2^1(t) \cdots \hat{r}_2^{k''-1}(t)$ for all such t. Because $\lim_{t\to\infty} \hat{r}(t) = \hat{r}$, we obtain $\hat{r}_1^1 \cdots \hat{r}_1^{k'-1} \leq \varepsilon \cdot \hat{r}_2^1 \cdots \hat{r}_2^{k''-1}$, as required.

Claim: There exists some $(r_1,r_2) \in R \times R$ such that $(r_1,r_2) \in F(r_1,r_2)$.

Proof of Claim: $R \times R$ is a nonempty, compact, and convex subset of a Euclidean space. Moreover, the above claims establish that F satisfies all the conditions of Lemma 10, so F has a fixed point.

For purposes of the following claims, we define the following quantity in terms of the LPS $\rho = (p^1,\ldots,p^K)$:

$$p_0 = \min_{k\in\{1,\ldots,K\}} \min_{\bar{s}\in \mathrm{supp}(p^k)} p^k(\bar{s})$$

and define $\gamma = \dfrac{\varepsilon}{(1-\varepsilon)p_0}$.

Claim: If $(r_1,r_2) \in F(r_1,r_2)$, then the strategy profile $\sigma = ((r_1\Box\rho)_1,(r_2\Box\rho)_2)$ is an (α,γ)-extended proper equilibrium.

Proof of Claim: That σ is totally mixed follows from the definition of R and the fact that ρ has full support. We next make a useful observation. Fix any $n \in \{1,2\}$ and any $s_n \in S_n$. Let k be such that $s_n \in \zeta_n^k$. Then, observe that because this is a two-player game, player n has just a single opponent, so that both $\sigma = ((r_1\Box\rho)_1,(r_2\Box\rho)_2)$ and $r_{-n}\Box\rho$ induce the same distribution of play over the opponent's strategies. Then, given the above definition of $L_n^k(\cdot)$, it follows that $\alpha_n L_n(\sigma/s_n) = L_n^k(r_{-n})$.

To complete the argument, we consider the within-player and across-player restrictions separately. For the within-player case, it is without loss of generality to focus on player 1. First, suppose $s_1' \in S_1$ and $s_1'' \in S_1$ are such that $\alpha_1 L_1(\sigma/s_1') > \alpha_1 L_1(\sigma/s_1'')$. Let k' and k'' be such that $s_1' \in \zeta_1^{k'}$ and $s_1'' \in \zeta_1^{k''}$. We then have $L_1^{k'}(r_2) > L_1^{k''}(r_2)$. This requires $k' > k''$. Therefore, we have the following:

$$\begin{aligned}
\sigma_1(s_1') &= (r_1\Box\rho)_1(s_1') \\
&\leq r_1^1 \cdots r_1^{k'-1} \\
&\leq r_1^1 \cdots r_1^{k''-1} r_1^{k'-1}
\end{aligned}$$

$$
\begin{aligned}
&\le \varepsilon \cdot r_1^1 \cdots r_1^{k''-1} \\
&= \gamma \cdot r_1^1 \cdots r_1^{k''-1}(1-\varepsilon)p_0 \\
&\le \gamma \cdot r_1^1 \cdots r_1^{k''-1}(1-r_1^{k''})p_0 \\
&\le \gamma \cdot (r_1 \square \rho)_1(s_1'') \\
&= \gamma \cdot \sigma_1(s_1''),
\end{aligned}
$$

as required. In the above, the first and eighth (i.e., last) steps follow from the definition of σ. The second and seventh steps follow from Lemma 11 and the definition of p_0. The third step uses $k' > k''$. The fourth step uses $r_1^{k'-1} \le \varepsilon$. The fifth step uses the definition of γ. The sixth step uses $r_1^{k''} \le \varepsilon$.

For the across-player case, it is without loss of generality to suppose $s_1' \in S_1$ and $s_2'' \in S_2$ are such that $\alpha_1 L_1(\sigma/s_1') > \alpha_2 L_2(\sigma/s_2'')$. Let k' and k'' be such that $s_1' \in \zeta_1^{k'}$ and $s_2'' \in \zeta_2^{k''}$. We then have $L_1^{k'}(r_2) > L_2^{k''}(r_1)$. Note that this precludes the case of $k' = 1$. Because $(r_1, r_2) \in F(r_1, r_2)$, we then have $r_1^1 \cdots r_1^{k'-1} \le \varepsilon \cdot r_2^1 \cdots r_2^{k''-1}$. Therefore, we have the following (where to accommodate the accommodate the case of $k'' = K$, the argument remains valid if we define $r_2^K = 0$):

$$
\begin{aligned}
\sigma_1(s_1') &= (r_1 \square \rho)_1(s_1') \\
&\le r_1^1 \cdots r_1^{k'-1} \\
&\le \varepsilon \cdot r_2^1 \cdots r_2^{k''-1} \\
&= \gamma \cdot r_2^1 \cdots r_2^{k''-1}(1-\varepsilon)p_0 \\
&\le \gamma \cdot r_2^1 \cdots r_2^{k''-1}(1-r_2^{k''})p_0 \\
&\le \gamma \cdot (r_2 \square \rho)_2(s_2'') \\
&= \gamma \cdot \sigma_2(s_2''),
\end{aligned}
$$

as required. The primary departure from the within-player case lies in the third step, where we now use $r_1^1 \cdots r_1^{k'-1} \le \varepsilon \cdot r_2^1 \cdots r_2^{k''-1}$. As pointed out above, this is where the status of (r_1, r_2) as a fixed point of F is relevant.

Claim: $\hat{\sigma}$ is an extended proper equilibrium.

Proof of Claim: For all $t \in \mathbb{N}$, define $\varepsilon_t = \frac{\delta}{t+1}$ and $\gamma_t = \frac{\varepsilon_t}{(1-\varepsilon_t)p_0}$. Note that $\lim_{t\to\infty} \varepsilon_t = \lim_{t\to\infty} \gamma_t = 0$. By the previous claims, there exists a sequence $(\sigma^t)_{t=1}^{\infty}$, where each σ^t is an (α, γ_t)-extended proper equilibrium that can be written as $\sigma^t = ((r_1(t) \square \rho)_1, (r_2(t) \square \rho)_2)$ for some $(r_1(t), r_2(t)) \in [\varepsilon_t^{2K^2-3K+2}, \varepsilon_t]^{K-1} \times [\varepsilon_t^{2K^2-3K+2}, \varepsilon_t]^{K-1}$. For all $n \in \{1, 2\}$ and all $s_n \in S_n$, we therefore have

$$
\sigma_n^t(s_n) = (r_n(t) \square \rho)_n(s_n) \ge (1 - r_n^1(t))p_n^1(s_n) \ge (1-\varepsilon_t)p_n^1(s_n).
$$

We conclude that $\lim_{t\to\infty} \sigma^t = (p_1^1, p_2^1)$, which by condition *(ii)* of Definition 6 must be $\hat{\sigma}$. Consequently, $\hat{\sigma}$ is an extended proper equilibrium. $\square$

B.2 Proofs Corresponding to Section 4

Proof of Theorem 4. For the symmetric game $\bar{\Gamma}$, we use $\bar{\pi}(\bar{\sigma}'; \bar{\sigma})$ throughout this proof to denote the expected payoff received from $\bar{\sigma}'$ when all opponents play $\bar{\sigma}$. Thus, letting σ' and σ be, respectively, the projections of $\bar{\sigma}'$ and $\bar{\sigma}$, the expression for the payoff simplifies to

$$
\bar{\pi}(\bar{\sigma}'; \bar{\sigma}) = \frac{1}{N} \sum_{n \in \mathcal{N}} \pi_n(\sigma/\sigma_n').
$$

Given a strategy profile σ in Γ, we refer to $\prod_{n\in\mathcal{N}}\sigma_n$ as the distribution over $\bar{S}$ *induced* by σ.

Necessity of Nash equilibrium. Suppose Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose $\bar{\sigma}$ is a symmetrically Nash strategy of the meta-game $\bar{\Gamma}'$. Let σ be the projection of $\bar{\sigma}$, and suppose, by way of contradiction, that σ is not a Nash equilibrium of Γ. Then σ is not a Nash equilibrium of Γ' either. Thus, there exists some $n \in \mathcal{N}$ and some $s_n \in S_n$ such that $\pi'_n(\sigma/s_n) > \pi'_n(\sigma)$. Letting $\bar{\sigma}'$ denote the distribution over $\bar{S}$ induced by σ/s_n, this means that

$$\bar{\pi}'(\bar{\sigma}';\bar{\sigma}) - \bar{\pi}'(\bar{\sigma};\bar{\sigma}) = \frac{1}{N}\left[\pi'_n(\sigma/s_n) - \pi'_n(\sigma)\right] > 0,$$

which contradicts $(\bar{\sigma},\ldots,\bar{\sigma})$ having been a Nash equilibrium of $\bar{\Gamma}'$. We conclude that σ is a Nash equilibrium of Γ.

Sufficiency of Nash equilibrium. Let σ be a Nash equilibrium of Γ. Let $\bar{\sigma}$ be the distribution over $\bar{S}$ induced by σ. It suffices to show that $\bar{\sigma}$ is a symmetrically Nash strategy of $\bar{\Gamma}$, the meta-version of Γ itself. For any $\bar{s} = (s_1,\ldots,s_N)$,

$$\bar{\pi}(\bar{s};\bar{\sigma}) - \bar{\pi}(\bar{\sigma};\bar{\sigma}) = \frac{1}{N}\sum_{n\in\mathcal{N}}\left[\pi_n(\sigma/s_n) - \pi_n(\sigma)\right].$$

Because σ is a Nash equilibrium of Γ, every element of the sum on the righthand side is nonpositive. We conclude that $\bar{\pi}(\bar{s};\bar{\sigma}) \le \bar{\pi}(\bar{\sigma};\bar{\sigma})$. Because the same argument can be made for all $\bar{s}$, we conclude that $(\bar{\sigma},\ldots,\bar{\sigma})$ is a Nash equilibrium of $\bar{\Gamma}$, so that $\bar{\sigma}$ is a symmetrically Nash strategy of $\bar{\Gamma}$.

Necessity of perfect equilibrium. Suppose Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose $\bar{\sigma}$ is a symmetrically perfect strategy of $\bar{\Gamma}'$. Equivalently, there exists a sequence of positive numbers $(\varepsilon_t)_{t=1}^{\infty}$ converging to zero and a sequence of totally mixed strategies $(\bar{\sigma}^t)_{t=1}^{\infty}$ converging to $\bar{\sigma}$ such that for all t, $(\bar{\sigma}^t,\ldots,\bar{\sigma}^t)$ is an ε_t-perfect equilibrium of $\bar{\Gamma}'$. Let, for all t, σ^t be the projection of $\bar{\sigma}^t$, and let σ be the projection of $\bar{\sigma}$. Then $(\sigma^t)_{t=1}^{\infty}$ is a sequence of totally mixed strategy profiles in Γ' converging to σ. We complete the proof by showing that, for all t, σ^t is an $(|\bar{S}|\varepsilon_t)$-perfect equilibrium of Γ. To see this, suppose that $n \in \mathcal{N}$ and $s'_n, s''_n \in S_n$ are such that $\pi_n(\sigma^t/s'_n) < \pi_n(\sigma^t/s''_n)$. A corresponding inequality must hold in the equivalent game Γ': $\pi'_n(\sigma^t/s'_n) < \pi'_n(\sigma^t/s''_n)$. For any $\bar{s} \in \bar{S}$, we then have

$$\bar{\pi}'(\bar{s}/s''_n;\bar{\sigma}^t) - \bar{\pi}'(\bar{s}/s'_n;\bar{\sigma}^t) = \frac{1}{N}\left[\pi'_n(\sigma^t/s''_n) - \pi'_n(\sigma^t/s'_n)\right] > 0.$$

Because $(\bar{\sigma}^t,\ldots,\bar{\sigma}^t)$ is an ε_t-perfect equilibrium of $\bar{\Gamma}'$, we therefore have $\bar{\sigma}^t(\bar{s}/s'_n) \le \varepsilon_t$. Since the same argument can be made for all $\bar{s} \in \bar{S}$, and since σ^t_n is the marginal of $\bar{\sigma}^t$ on S_n, we conclude that $\sigma^t_n(s'_n) \le |\bar{S}|\varepsilon_t$, as required.

Sufficiency of perfect equilibrium. Let σ be a perfect equilibrium of Γ. Equivalently, there exists a sequence of positive numbers $(\varepsilon_t)_{t=1}^{\infty}$ converging to zero and a sequence of totally mixed strategy profiles $(\sigma^t)_{t=1}^{\infty}$ converging to σ such that for all t, σ^t is an ε_t-perfect equilibrium of Γ. Let, for all t, $\bar{\sigma}^t$ be the distribution over $\bar{S}$ induced by σ^t, and let $\bar{\sigma}$ be the distribution over $\bar{S}$ induced by σ. Then $(\bar{\sigma}^t)_{t=1}^{\infty}$ is a sequence of full support distributions converging to $\bar{\sigma}$. It suffices to show that $\bar{\sigma}$ is a symmetrically perfect strategy of $\bar{\Gamma}$, the meta-version of Γ itself. This can be done by showing that for all t, $(\bar{\sigma}^t,\ldots,\bar{\sigma}^t)$ is an ε_t-perfect equilibrium of $\bar{\Gamma}$. Let $\bar{s}' = (s'_1,\ldots,s'_N)$ and $\bar{s}'' = (s''_1,\ldots,s''_N)$ be such that

$$0 < \bar{\pi}(\bar{s}'';\bar{\sigma}^t) - \bar{\pi}(\bar{s}';\bar{\sigma}^t) = \frac{1}{N}\sum_{n\in\mathcal{N}}\left[\pi_n(\sigma^t/s''_n) - \pi_n(\sigma^t/s'_n)\right].$$

This requires there to exist some $n \in \mathcal{N}$ such that $\pi_n(\sigma^t/s'_n) < \pi_n(\sigma^t/s''_n)$. Because σ^t is an ε_t-perfect equilibrium of Γ, we have $\sigma^t_n(s'_n) \le \varepsilon_t$. Because σ^t_n is the marginal of $\bar{\sigma}^t$ on S_n, we have $\bar{\sigma}^t(\bar{s}') \le \sigma^t_n(s'_n)$, so that $\bar{\sigma}^t(\bar{s}') \le \varepsilon_t$, as required.

Necessity of extended proper equilibrium. Suppose Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose $\bar{\sigma}$ is a symmetrically proper strategy of $\bar{\Gamma}'$. Equivalently, there exists a sequence of positive numbers $(\varepsilon_t)_{t=1}^{\infty}$ converging to zero and a sequence of totally mixed strategies $(\bar{\sigma}^t)_{t=1}^{\infty}$ converging to $\bar{\sigma}$ such that for all t, $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an ε_t-proper equilibrium of $\bar{\Gamma}'$. Let, for all t, σ^t be the projection of $\bar{\sigma}^t$, and let σ be the projection of $\bar{\sigma}$. Then $(\sigma^t)_{t=1}^{\infty}$ is a sequence of totally mixed strategy profiles in Γ' converging to σ. Let α be such that Γ' can be constructed from Γ by multiplying the utility function of each player n by α_n and adding some constant. We complete the proof by showing that, for all t, σ^t is an $\big(\alpha, |\bar{S}|\varepsilon_t\big)$-extended proper equilibrium of Γ. To see this, suppose that for some players l and m, $s'_l \in S_l$ and $s''_m \in S_m$ are such that $\alpha_l L_l(\sigma^t/s'_l) > \alpha_m L_m(\sigma^t/s''_m)$. For all $n \in \mathcal{N}$, select some $s^*_n \in BR_n(\sigma^t)$. Define $s'' = s^*/s''_m$. Then for all $\hat{s} \in \bar{S}$ for which $\hat{s}_l = s'_l$, it is the case that

$$\sum_{n\in\mathcal{N}} \alpha_n L_n(\sigma^t/\hat{s}_n) \ge \alpha_l L_l(\sigma^t/s'_l) > \alpha_m L_m(\sigma^t/s''_m) = \sum_{n\in\mathcal{N}} \alpha_n L_n(\sigma^t/s''_n).$$

Equivalently, $\bar{\pi}'(\hat{s};\bar{\sigma}^t) < \bar{\pi}'(s'';\bar{\sigma}^t)$. Because $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an ε_t-proper equilibrium of $\bar{\Gamma}'$, we therefore have $\bar{\sigma}^t(\hat{s}) \le \varepsilon_t \bar{\sigma}^t(s'')$. Because σ^t_m is the marginal of $\bar{\sigma}^t$ on S_m, we have $\bar{\sigma}^t(\bar{s}'') \le \sigma^t_m(s''_m)$, so that we obtain

$$\bar{\sigma}^t(\hat{s}) \le \varepsilon_t \sigma^t_m(s''_m).$$

Moreover, since we can make the above argument for each strategy profile $\hat{s} \in \bar{S}$ for which $\hat{s}_l = s'_l$, and since σ^t_l is the marginal of $\bar{\sigma}^t$ on S_l, we conclude $\sigma^t_l(s'_l) \le |\bar{S}|\varepsilon_t \sigma^t_m(s''_m)$, as required.

Sufficiency of extended proper equilibrium. Let σ be an extended proper equilibrium of Γ. Equivalently, there exists an $\alpha \in \mathbb{R}^N_{++}$, a sequence of numbers $(\varepsilon_t)_{t=1}^{\infty}$ in the open unit interval converging to zero, and a sequence of totally mixed strategy profiles $(\sigma^t)_{t=1}^{\infty}$ converging to σ such that for all t, σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium of Γ. Let, for all t, ϕ^t be the distribution over $\bar{S}$ induced by σ^t. For all t, we construct another distribution over $\bar{S}$, $\bar{\sigma}^t$, in the way described below. To foreshadow, we will subsequently establish (i) that $\bar{\sigma}^t$ has σ^t as its projection, and (ii) that for all sufficiently large values of t, $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an ε_t-proper equilibrium of the meta-version of a game that is equivalent to Γ up to an affine transformation of the payoffs.

To begin, let $\succ$ be the partial order over $\bar{S}$ defined as follows (where its dependence on t is suppressed in the notation):

$$s' \succ s'' \iff \sum_{n\in\mathcal{N}} \alpha_n \pi_n(\sigma^t/s'_n) > \sum_{n\in\mathcal{N}} \alpha_n \pi_n(\sigma^t/s''_n).$$

We say that $s' \succcurlyeq s''$ if it is not the case that $s'' \succ s'$. Next, select for all $n \in \mathcal{N}$ some $s^*_n \in \arg\max_{s_n \in S_n} \sigma^t_n(s_n)$. Note that because σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium with $\varepsilon_t \in (0,1)$, it must be the case that $s^*_n \in BR_n(\sigma^t)$. It follows that s^* is a maximal element under $\succcurlyeq$ (although perhaps not the unique maximal element).

Next, we partition the set of strategy profiles $\bar{S}$ into (i) s^* itself, (ii) $\mathcal{S}^1$, which we define as the set of unilateral deviations from s^*, (iii) $\mathcal{S}^2$, which we define as the set of joint deviations from s^* that are played relatively infrequently under σ^t, and (iv) $\mathcal{S}^3$, which we define as the set of joint deviations from s^* that are played relatively frequently under σ^t. Formally,

$$\begin{aligned}\mathcal{S}^1 =&\{s^*/s_n : n \in \mathcal{N}, s_n \in S_n \setminus \{s_n^*\}\}\\ \mathcal{S}^2 =&\{s \in \bar{S} \setminus (\{s^*\} \cup \mathcal{S}^1) : \phi^t(s) < \varepsilon_t \phi^t(s^*)\}\\ \mathcal{S}^3 =&\{s \in \bar{S} \setminus (\{s^*\} \cup \mathcal{S}^1) : \phi^t(s) \geq \varepsilon_t \phi^t(s^*)\}\end{aligned}$$

Note that because σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium, a strategy profile in which at least one player uses an inferior response to σ^t can be played with probability at most $\varepsilon_t^{|\bar{S}|+2}\phi^t(s^*)$. Hence, the elements of $\mathcal{S}^3$ must all be maximal under $\succcurlyeq$.

To construct $\bar{\sigma}^t$, the desired distribution over $\bar{S}$, initialize $\bar{\sigma}^t(s) = \phi^t(s)$. Then loop over the elements $s' \in \mathcal{S}^2$, at each step modifying $\bar{\sigma}^t$ in the following way:

- *[Modify s' itself:]* At the step for $s' \in \mathcal{S}^2$, select some $\hat{s} \in \arg\min_{s \in \{s^*\} \cup \mathcal{S}^1 : s \succcurlyeq s'} \phi^t(s)$ and let $\ell = \#\{s \in \mathcal{S}^2 : s \succ s'\}$. Set $\bar{\sigma}^t(s') \leftarrow \varepsilon_t^{2+\ell}\phi^t(\hat{s})$.
- *[Modify corresponding elements of $\mathcal{S}^1$:]* For each $n \in \mathcal{N}$ for which $s'_n \neq s_n^*$, set $\bar{\sigma}^t(s^*/s'_n) \leftarrow \bar{\sigma}^t(s^*/s'_n) + \phi^t(s') - \bar{\sigma}^t(s')$.
- *[Modify s^*:]* Set $\bar{\sigma}^t(s^*) \leftarrow \bar{\sigma}^t(s^*) - \big(\#\{n \in \mathcal{N} : s'_n \neq s_n^*\} - 1\big)\big(\phi^t(s') - \bar{\sigma}^t(s')\big)$.

In words, the step for $s' \in \mathcal{S}^2$ alters the probability put on s'; but so as to keep the marginals unchanged, the step also alters the probability on corresponding elements of $\mathcal{S}^1$ and on s^*. The elements of $\mathcal{S}^3$ remain untouched throughout. Henceforth, let $\bar{\sigma}^t$ refer to its value after all steps of the loop have been completed, unless specified otherwise.

Claim: σ^t is the projection of $\bar{\sigma}^t$.

Proof of Claim: Because σ^t is the projection of ϕ^t, it suffices to show that the step of the loop corresponding to each $s' \in \mathcal{S}^2$ does not alter the marginal distributions. For any $m \in \mathcal{N}$ and any $s_m \in S_m$, we verify that $\bar{\sigma}^t$ assigns the same total probability to the strategy profiles involving s_m both before and after the step. To do so, we consider separately the two cases in which $s'_m = s_m^*$ and $s'_m \neq s_m^*$.

Suppose $m \in \mathcal{N}$ is such that $s'_m = s_m^*$. The step modifies $\bar{\sigma}^t$ by assigning $\phi^t(s') - \bar{\sigma}^t(s')$ less weight to one strategy profile involving s_m^* (namely s'), $\big(\#\{n \in \mathcal{N} : s'_n \neq s_n^*\} - 1\big)\big(\phi^t(s') - \bar{\sigma}^t(s')\big)$ less weight to another strategy profile involving s_m^* (namely s^*), as well as $\phi^t(s') - \bar{\sigma}^t(s')$ more weight to $\#\{n \in \mathcal{N} : s'_n \neq s_n^*\}$ strategy profiles involving s_m^*. Thus, the total weight on strategy profiles involving s_m^* remains unchanged. Moreover, the step does not adjust the weight that $\bar{\sigma}^t$ places on any strategy profile not involving s_m^*.

Suppose $m \in \mathcal{N}$ is such that $s'_m \neq s_m^*$. The step modifies $\bar{\sigma}^t$ by assigning $\big(\#\{n \in \mathcal{N} : s'_n \neq s_n^*\} - 1\big)\big(\phi^t(s') - \bar{\sigma}^t(s')\big)$ less weight to one strategy profile involving s_m^* (namely s^*), as well as $\phi^t(s') - \bar{\sigma}^t(s')$ more weight to $\#\{n \in \mathcal{N} : s'_n \neq s_n^*\} - 1$ strategy profiles involving s_m^*. Thus, the total weight on strategy profiles involving s_m^* remains unchanged. It also modifies $\bar{\sigma}^t$ by assigning $\phi^t(s') - \bar{\sigma}^t(s')$ less weight to one strategy profile involving s'_m (namely s'), as well as $\phi^t(s') - \bar{\sigma}^t(s')$ more weight to another strategy profile involving s'_m (namely s^*/s'_m). Thus, the total weight on strategy profiles involving s'_m remains unchanged. Moreover, the step does not adjust the weight that $\bar{\sigma}^t$ places on any strategy profile not involving either s_m^* or s'_m.

Claim: For all $s \in \mathcal{S}^1$, $\bar{\sigma}^t(s) \leq |\bar{S}|\phi^t(s)$.

Proof of Claim: To see this, note that every $s \in \mathcal{S}^1$ can be written as $s = s^*/s_m$ for some $m \in \mathcal{N}$ and some $s_m \in S_m$. Moreover, $\bar{\sigma}^t(s)$ is updated only at steps in the loop corresponding to $s' \in \mathcal{S}^2$

where $s'_m = s_m$. At each such step, $\bar{\sigma}^t(s)$ is raised by at most $\phi^t(s')$. Based on the previous definitions,

$$\phi^t(s') = \sigma^t_m(s_m) \prod_{n \neq m} \sigma^t_n(s'_n) \leq \sigma^t_m(s_m) \prod_{n \neq m} \sigma^t_n(s^*_n) = \phi^t(s).$$

Because $\bar{\sigma}^t(s)$ is initialized at $\phi^t(s)$ and because there are no more than $|\bar{S}| - 1$ steps, we conclude that $\bar{\sigma}^t(s) \leq |\bar{S}|\phi^t(s)$, as desired.

Claim: For all $s \in \mathcal{S}^1$, $\bar{\sigma}^t(s) \geq (1 - |\bar{S}|\varepsilon_t^2)\phi^t(s)$.

Proof of Claim: To see this, note that every $s \in \mathcal{S}^1$ can be written as $s = s^*/s_m$ for some $m \in \mathcal{N}$ and some $s_m \in S_m$. Moreover, $\bar{\sigma}^t(s)$ is updated only at steps in the loop corresponding to $s' \in \mathcal{S}^2$ where $s'_m = s_m$. As previously observed, we have $s^*_n \in BR_n(\sigma^t)$ for all $n \in \mathcal{N}$. Thus,

$$\alpha_m \pi_m(\sigma^t/s_m) + \sum_{n \neq m} \alpha_n \pi_n(\sigma^t/s^*_n) \geq \alpha_m \pi_m(\sigma^t/s_m) + \sum_{n \neq m} \alpha_n \pi_n(\sigma^t/s'_n),$$

so that $s \succcurlyeq s'$. Then by construction, $\bar{\sigma}^t(s') \leq \varepsilon_t^2 \phi^t(s)$. At each such step, $\bar{\sigma}^t(s)$ is reduced by at most this amount. Because $\bar{\sigma}^t(s)$ is initialized at $\phi^t(s)$ and because there are no more than $|\bar{S}|$ steps, we conclude that $\bar{\sigma}^t(s) \geq (1 - |\bar{S}|\varepsilon_t^2)\phi^t(s)$, as desired.

Claim: $\bar{\sigma}^t(s^*) \geq (1 - N|\bar{S}|\varepsilon_t)\phi^t(s^*)$.

Proof of Claim: For all $s' \in \mathcal{S}^2$, $\phi^t(s') \leq \varepsilon_t \phi^t(s^*)$, by the definition of $\mathcal{S}^2$. Next, note that the step of the loop corresponding to $s' \in \mathcal{S}^2$ reduces $\bar{\sigma}^t(s^*)$ by no more than $N\phi^t(s')$. Because there are no more than $|\bar{S}|$ steps, the conclusion follows.

Next, define T such that for all $t \geq T$, both of the following hold. First, $\varepsilon_t \leq \frac{1}{N|\bar{S}|}$, which ensures that $\bar{\sigma}^t(s) \geq 0$ for all $s \in \bar{S}$, so that $\bar{\sigma}^t$ is a well-defined probability distribution over $\bar{S}$. Second,

$$\varepsilon_t \geq \max\left\{ \frac{|\bar{S}|^2 \varepsilon_t^{|\bar{S}|+2}}{1 - N|\bar{S}|\varepsilon_t}, \frac{\varepsilon_t^2}{1 - N|\bar{S}|\varepsilon_t}, |\bar{S}|^2 \varepsilon_t^2 \right\},$$

which we use in the arguments below. Furthermore, let Γ' be the game that is constructed from Γ by multiplying the utility function of each player n by α_n. Let $\bar{\Gamma}'$ be the meta-version of Γ'. Because the projection of $\bar{\sigma}^t$ is σ^t, the payoff of some strategy $\bar{s} = (s_1, \ldots, s_N)$ in $\bar{\Gamma}'$ when all other players play according to $\bar{\sigma}^t$ is

$$\bar{\pi}'(\bar{s}; \bar{\sigma}^t) = \frac{1}{N} \sum_{n \in \mathcal{N}} \alpha_n \pi_n(\sigma^t/s_n).$$

Claim: For all $t \geq T$, $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an ε_t-proper equilibrium of $\bar{\Gamma}'$.[34]

Proof of Claim: Suppose that $t \geq T$. Suppose that $\bar{s}' = (s'_1, \ldots, s'_N)$ and $\bar{s}'' = (s''_1, \ldots, s''_N)$ are such that $\bar{\pi}'(\bar{s}'; \bar{\sigma}^t) > \bar{\pi}'(\bar{s}''; \bar{\sigma}^t)$. Equivalently, $\frac{1}{N}\sum_{n \in \mathcal{N}} \alpha_n \pi_n(\sigma^t/s'_n) > \frac{1}{N}\sum_{n \in \mathcal{N}} \alpha_n \pi_n(\sigma^t/s''_n)$, which is also equivalent to $\bar{s}' \succ \bar{s}''$. We will show that $\bar{\sigma}^t(\bar{s}'') \leq \varepsilon_t \bar{\sigma}^t(\bar{s}')$, which in turn implies that $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an ε_t-proper equilibrium of $\bar{\Gamma}'$. To establish this, we consider eight cases, which are distinguished by how $\bar{s}'$ and $\bar{s}''$ fit into the partition $\{\{s^*\}, \mathcal{S}^1, \mathcal{S}^2, \mathcal{S}^3\}$. Note that because s^* and the elements of $\mathcal{S}^3$ are maximal under $\succcurlyeq$, we can have only $\bar{s}'' \in \mathcal{S}^2 \cup \mathcal{S}^3$, which eliminates some potential cases.

[34] Because $\bar{\Gamma}'$ is a symmetric game and because the strategy profile in question is symmetric, we will have also shown that $(\bar{\sigma}^t, \ldots, \bar{\sigma}^t)$ is an (α, ε_t)-extended proper equilibrium with $\alpha = (1, \ldots, 1)$.

Case 1: $s' = s^*$ and $s'' \in \mathcal{S}^1$. Note that for some player m, we can write $s'' = s^*/s''_m$. Thus, $s^* \succ s''$ implies $\alpha_m L_m(\sigma^t/s^*_m) < \alpha_m L_m(\sigma^t/s''_m)$. Because σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium of Γ, we have that $\sigma^t_m(s''_m) \leq \varepsilon_t^{|\bar{S}|+2}\sigma^t_m(s^*_m)$. Because σ^t_m is the marginal of ϕ^t on S_m, we also have $\phi^t(s'') \leq \sigma^t_m(s''_m)$. We also have $\sigma^t_m(s^*_m) = \frac{\phi^t(s^*)}{\prod_{n \neq m} \sigma^t_n(s^*_n)} \leq \left(\prod_{n \neq m} |S_n|\right) \phi^t(s^*) \leq |\bar{S}|\phi^t(s^*)$, where the first inequality follows from $s^*_n \in \arg\max_{s_n \in S_n} \sigma^t_n(s_n)$ for all $n \in \mathcal{N}$. Further, as derived above, $\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'')$ and $\bar{\sigma}^t(s^*) \geq (1 - N|\bar{S}|\varepsilon_t)\phi^t(s^*)$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'') \leq |\bar{S}|\sigma^t_m(s''_m) \leq |\bar{S}|\varepsilon_t^{|\bar{S}|+2}\sigma^t_m(s^*_m) \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}\phi^t(s^*) \leq \frac{|\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}}{1 - N|\bar{S}|\varepsilon_t}\bar{\sigma}^t(s^*).$$

From the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s^*)$ for all $t \geq T$.

Case 2: $s' \in \mathcal{S}^1$ and $s'' \in \mathcal{S}^1$. Note that for some players l and m, we can write $s' = s^*/s'_l$ and $s'' = s^*/s''_m$. As previously observed, we have $s^*_n \in BR_n(\sigma^t)$ for all $n \in \mathcal{N}$, and so $s' \succ s''$ implies $\alpha_l L_l(\sigma^t/s'_l) < \alpha_m L_m(\sigma^t/s''_m)$. Because σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium of Γ, we have that $\sigma^t_m(s''_m) \leq \varepsilon_t^{|\bar{S}|+2}\sigma^t_l(s'_l)$. Because σ^t_m is the marginal of ϕ^t on S_m, we also have $\phi^t(s'') \leq \sigma^t_m(s''_m)$. We also have $\sigma^t_l(s'_l) = \frac{\phi^t(s')}{\prod_{n \neq l} \sigma^t_n(s^*_n)} \leq \left(\prod_{n \neq l} |S_n|\right) \phi^t(s') \leq |\bar{S}|\phi^t(s')$, where the first inequality follows from $s^*_n \in \arg\max_{s_n \in S_n} \sigma^t_n(s_n)$ for all $n \in \mathcal{N}$. Further, as derived above, $\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'')$ and $\bar{\sigma}^t(s') \geq (1 - |\bar{S}|\varepsilon_t^2)\phi^t(s')$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'') \leq |\bar{S}|\sigma^t_m(s''_m) \leq |\bar{S}|\varepsilon_t^{|\bar{S}|+2}\sigma^t_l(s'_l) \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}\phi^t(s') \leq \frac{|\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}}{1 - |\bar{S}|\varepsilon_t^2}\bar{\sigma}^t(s').$$

Note that $\frac{|\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}}{1-|\bar{S}|\varepsilon_t^2} \leq \frac{|\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}}{1-N|\bar{S}|\varepsilon_t}$. Thus, from the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s')$ for all $t \geq T$.

Case 3: $s' \in \mathcal{S}^1$ and $s'' \in \mathcal{S}^2$. Its construction implies $\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s')$. Further, as derived above, $\bar{\sigma}^t(s') \geq (1 - |\bar{S}|\varepsilon_t^2)\phi^t(s')$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s') \leq \frac{\varepsilon_t^2}{1 - |\bar{S}|\varepsilon_t^2}\bar{\sigma}^t(s').$$

Note that $\frac{\varepsilon_t^2}{1-|\bar{S}|\varepsilon_t^2} \leq \frac{\varepsilon_t^2}{1-N|\bar{S}|\varepsilon_t}$. Thus, from the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s')$ for all $t \geq T$.

Case 4: $s' \in \mathcal{S}^2$ and $s'' \in \mathcal{S}^1$. Let $\hat{s} \in \mathcal{S}^1$ be as specified above for the step of the loop involving s'. By its construction, $\bar{\sigma}^t(s') \geq \varepsilon_t^{|\bar{S}|}\phi^t(\hat{s})$. We also have $\hat{s} \not\succ s' \succ s''$. Note that for some players l and m, we can write $\hat{s} = s^*/\hat{s}_l$ and $s'' = s^*/s''_m$. As previously observed, we have $s^*_n \in BR_n(\sigma^t)$ for all $n \in \mathcal{N}$, and so $\hat{s} \succ s''$ implies $\alpha_l L_l(\sigma^t/\hat{s}_l) < \alpha_m L_m(\sigma^t/s''_m)$. Because σ^t is an $(\alpha, \varepsilon_t^{|\bar{S}|+2})$-extended proper equilibrium of Γ, we have that $\sigma^t_m(s''_m) \leq \varepsilon_t^{|\bar{S}|+2}\sigma^t_l(\hat{s}_l)$. Because σ^t_m is the marginal of ϕ^t on S_m, we also have $\phi^t(s'') \leq \sigma^t_m(s''_m)$. We also have $\sigma^t_l(\hat{s}_l) = \frac{\phi^t(\hat{s})}{\prod_{n \neq l} \sigma^t_n(s^*_n)} \leq \left(\prod_{n \neq l} |S_n|\right) \phi^t(\hat{s}) \leq |\bar{S}|\phi^t(\hat{s})$, where the first inequality follows from $s^*_n \in \arg\max_{s_n \in S_n} \sigma^t_n(s_n)$ for all $n \in \mathcal{N}$. Further, as derived above, $\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'')$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq |\bar{S}|\phi^t(s'') \leq |\bar{S}|\sigma^t_m(s''_m) \leq |\bar{S}|\varepsilon_t^{|\bar{S}|+2}\sigma^t_l(\hat{s}_l) \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}\phi^t(\hat{s}) \leq |\bar{S}|^2\varepsilon_t^2\bar{\sigma}^t(s').$$

From the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s')$ for all $t \geq T$.

Case 5: $s' \in \mathcal{S}^2$ and $s'' \in \mathcal{S}^2$. Let $(\hat{s}', \ell')$ and $(\hat{s}'', \ell'')$ be as specified for the steps of the loop corresponding to s' and s'', respectively. Because $s' \succ s''$, we have $\phi^t(\hat{s}') \geq \phi^t(\hat{s}'')$ and $\ell'' \geq \ell' + 1$. Then

$$\bar{\sigma}^t(s'') = \varepsilon_t^{2+\ell''}\phi^t(\hat{s}'') \geq \varepsilon_t^{3+\ell'}\phi^t(\hat{s}'') = \varepsilon_t\varepsilon_t^{2+\ell'}\phi^t(\hat{s}') = \varepsilon_t\bar{\sigma}^t(s').$$

Case 6: $s' = s^*$ and $s'' \in \mathcal{S}^2$. By its construction, $\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s^*)$. As derived above, $\bar{\sigma}^t(s^*) \geq (1 - N|\bar{S}|\varepsilon_t)\phi^t(s^*)$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s^*) \leq \frac{\varepsilon_t^2}{1 - N|\bar{S}|\varepsilon_t}\bar{\sigma}^t(s^*)$$

Thus, from the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s')$ for all $t \geq T$.

Case 7: $s' \in \mathcal{S}^3$ and $s'' \in \mathcal{S}^1$. We can argue as in case 1 that $\bar{\sigma}^t(s'') \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}\phi^t(s^*)$. Because $s' \in \mathcal{S}^3$, $\bar{\sigma}^t(s') = \phi^t(s') \geq \varepsilon_t\phi^t(s^*)$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+2}\phi^t(s^*) \leq |\bar{S}|^2\varepsilon_t^{|\bar{S}|+1}\bar{\sigma}^t(s').$$

Note that $|\bar{S}|^2\varepsilon_t^{|\bar{S}|+1} \leq |\bar{S}|^2\varepsilon_t^2$. Thus, from the definition of T, we indeed obtain $\bar{\sigma}^t(s'') \leq \varepsilon_t\bar{\sigma}^t(s')$ for all $t \geq T$.

Case 8: $s' \in \mathcal{S}^3$ and $s'' \in \mathcal{S}^2$. By its construction, $\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s^*)$. Because $s' \in \mathcal{S}^3$, $\bar{\sigma}^t(s') = \phi^t(s') \geq \varepsilon_t\phi^t(s^*)$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s^*) \leq \varepsilon_t\bar{\sigma}^t(s').$$

Claim: σ is the projection of a symmetrically proper strategy of $\bar{\Gamma}'$.

Proof of Claim: Because $\prod_{n\in\mathcal{N}}\Delta_n$ is compact, $(\bar{\sigma}^t)_{t=T}^{\infty}$ has a convergent subsequence, the limit of which is therefore a symmetrically proper strategy of $\bar{\Gamma}'$.[35] Furthermore, the marginals of each $\bar{\sigma}^t$ are given by the profile σ^t. Therefore, the projection of the aforementioned limit equals $\lim_{t\to\infty}\sigma^t = \sigma$. □

B.3 Proofs Corresponding to Section 5

Proof of Proposition 5. *Sufficiency.* Suppose (ρ, σ) is a lexicographic Nash equilibrium, where $\rho = (p^1, \ldots, p^K)$ is an LPS that satisfies strong independence. For a player $n \in \mathcal{N}$ and a strategy $s'_n \in S_n$, if $p^1_n(s'_n) > 0$ then by condition *(i)* of Definition 6, for all $s''_n \in S_n$,

$$\sum_{s\in\bar{S}} p^1(s)\pi_n(s/s'_n) \geq \sum_{s\in\bar{S}} p^1(s)\pi_n(s/s''_n).$$

By condition *(ii)* of Definition 6, $\sigma = (p^1_1, \ldots, p^1_N)$. Therefore, we can rephrase the above in the following terms: if $\sigma_n(s'_n) > 0$ then for all $s''_n \in S_n$, $\pi_n(\sigma/s'_n) \geq \pi_n(\sigma/s''_n)$. Thus, σ is a Nash equilibrium.

Necessity. Suppose $\sigma = (\sigma_1, \ldots, \sigma_N)$ is a Nash equilibrium. Let $K = 1$ and define $p^1 = \prod_{n\in\mathcal{N}}\sigma_n$ as the distribution over $\bar{S}$ induced by σ. This produces an LPS ρ such that (ρ, σ) is a lexicographic Nash equilibrium. Furthermore, p^1 is by definition a product distribution. Because $K = 1$, ρ therefore satisfies strong independence. □

[35] Following footnote 34, this limit is in fact a *symmetrically extended proper strategy* of $\bar{\Gamma}'$, where the italicized term is defined in a way analogous to Definition 4.

Proof of Proposition 6. This is a restatement of Proposition 7 from BBD. $\square$

Proof of Proposition 7. This is a restatement of Proposition 8 from BBD. $\square$

Proof of Proposition 8. *Sufficiency.* Suppose (ρ, σ) is a lexicographic Nash equilibrium, where $\rho = (p^1, \ldots, p^K)$ is an LPS that satisfies strong independence, has full support, and respects within-and-across-player preferences. As BBD point out, ρ satisfies strong independence if and only if there exists a sequence $r(t) \in (0,1)^{K-1}$ with $r(t) \to 0$ such that $r(t)\Box\rho$ is a product distribution for all t. For each $n \in \mathcal{N}$, define $\sigma_n(t)$ as the marginal of $r(t)\Box\rho$ on S_n, and let $\sigma(t) = (\sigma_1(t), \ldots, \sigma_N(t))$. Note that $\lim_{t\to\infty} \sigma(t) = (p_1^1, \ldots, p_N^1)$, which by condition *(ii)* of Definition 6 is equal to σ. Also define

$$
\begin{aligned}
p_0 &= \min_{k\in\{1,\ldots,K\}} \min_{\bar{s}\in \mathrm{supp}(p^k)} p^k(\bar{s}) \\
\varepsilon(t) &= \max_{k\in\{1,\ldots,K-1\}} \left\{ \frac{r^k(t)}{[1-r^k(t)]p_0} \right\}
\end{aligned}
$$

Because $r(t) \to 0$, we also have $\varepsilon(t) \to 0$. Letting δ be as in the statement of Lemma 12, let T be such that for all $t \geq T$, $r(t) \in (0,\delta)^{K-1}$. In addition, let $\alpha \in \mathbb{R}^N_{++}$ be as in the definition of "respects within-and-across-player preferences." We claim that for all $t \geq T$, $\sigma(t)$ is an $(\alpha, \varepsilon(t))$-extended proper equilibrium.

Fix some $t \geq T$. From the fact that ρ has full support, we obtain for all $n \in \mathcal{N}$ that $\sigma_n(t) \in \Delta_n^0$. Furthermore, suppose $l, m \in \mathcal{N}$, $s_l' \in S_l$, and $s_m'' \in S_m$ are such that $\alpha_l L_l(\sigma(t)/s_l') > \alpha_m L_m(\sigma(t)/s_m'')$. Equivalently,

$$
\alpha_l \max_{\hat{s}_l\in S_l} [\pi_l(\sigma(t)/\hat{s}_l) - \pi_l(\sigma(t)/s_l')] > \alpha_m \max_{\hat{s}_m\in S_m} [\pi_m(\sigma(t)/\hat{s}_m) - \pi_m(\sigma(t)/s_m'')].
$$

Because $r(t)\Box\rho$ is a product distribution with marginals $\sigma_1(t), \ldots, \sigma_N(t)$, the above inequality is equivalent to

$$
\alpha_l \max_{\hat{s}_l\in S_l} \sum_{s\in\bar{S}} (r\Box\rho)(s)[\pi_l(s/\hat{s}_l) - \pi_l(s/s_l')] > \alpha_m \max_{\hat{s}_m\in S_m} \sum_{s\in\bar{S}} (r\Box\rho)(s)[\pi_m(s/\hat{s}_m) - \pi_m(s/s_m'')].
$$

By Lemma 12, we obtain that for all $s_l^* \in BR_l(\rho)$ and all $s_m^* \in BR_m(\rho)$,

$$
\left[\alpha_l \left(\pi_l(s_l^*|p^k) - \pi_l(s_l'|p^k)\right)\right]_{k=1}^K >_L \left[\alpha_m \left(\pi_m(s_m^*|p^k) - \pi_m(s_m''|p^k)\right)\right]_{k=1}^K,
$$

which, because ρ respects within-and-across-player preferences, implies that $s_l' <_\rho s_m''$. Define $k'' = \min\{k : p_m^k(s_m'') > 0\}$. Because ρ has full support, such a k'' must exist. We then have $\sigma_m(t)(s_m'') \geq r^1(t)\cdots r^{k''-1}(t)[1 - r^{k''}(t)]p_0$, by Lemma 11. Because $s_l' <_\rho s_m''$, we also have $\sigma_l(t)(s_l') \leq r^1(t)\cdots r^{k''}(t)$ by Lemma 11. Combining this with the definition of $\varepsilon(t)$ we have

$$
\begin{aligned}
\varepsilon(t)\sigma_m(t)(s_m'') &\geq \varepsilon(t) r^1(t)\cdots r^{k''-1}(t)[1-r^{k''}(t)]p_0 \\
&\geq \frac{r^{k''}(t)}{[1-r^{k''}(t)]p_0} r^1(t)\cdots r^{k''-1}(t)[1-r^{k''}(t)]p_0 \\
&= r^1(t)\cdots r^{k''}(t) \\
&\geq \sigma_l(t)(s_l').
\end{aligned}
$$

In review, we have shown for all $t \geq T$ that $\sigma(t)$ is a totally mixed strategy profile and $\alpha_l L_l(\sigma(t)/s_l') > \alpha_m L_m(\sigma(t)/s_m'')$ implies $\sigma_l(t)(s_l') \leq \varepsilon(t)\sigma_m(t)(s_m'')$. Consequently, $\sigma(t)$ is an

$(\alpha, \varepsilon(t))$-extended proper equilibrium. Because $\varepsilon(t) \to 0$ and $\sigma(t) \to \sigma$, σ is an extended proper equilibrium.

Necessity. Suppose $\sigma = (\sigma_1, \ldots, \sigma_N)$ is an extended proper equilibrium. Then there exists a scaling vector $\alpha \in \mathbb{R}^N_{++}$, a sequence of positive numbers $\varepsilon(t)$ converging to zero, and a sequence of totally mixed strategy profiles $\sigma(t)$ converging to σ where each $\sigma(t)$ is an $(\alpha, \varepsilon(t))$-extended proper equilibrium. Let $\bar{\sigma}(t) = \prod_{n \in \mathcal{N}} \sigma_n(t)$ be the distribution over $\bar{S}$ induced by σ. By Proposition 2 of BBD, there is an LPS $\rho = (p^1, \ldots, p^K)$ on $\bar{S}$ such that a subsequence $\bar{\sigma}(\tau)$ of $\bar{\sigma}(t)$ can be written as $\bar{\sigma}(\tau) = r(\tau)\Box\rho$ for a sequence $r(\tau) \in (0,1)^{K-1}$ with $r(\tau) \to 0$. We claim that ρ satisfies strong independence, has full support, respects within-and-across-player comparisons, and is such that (ρ, σ) meets condition *(ii)* of Definition 6.

First, condition *(ii)* of Definition 6 is satisfied, since $p^1 = \lim_{\tau\to\infty} r(\tau)\Box\rho = \lim_{\tau\to\infty} \bar{\sigma}(\tau) = \lim_{\tau\to\infty} \prod_{n\in\mathcal{N}} \sigma_n(\tau) = \prod_{n\in\mathcal{N}} \sigma_n$. Second, since each $\sigma(\tau)$ is totally mixed, each $\bar{\sigma}(\tau)$ has full support, and therefore ρ must have full support. Third, the definition of ρ implies that, for all τ, $r(\tau)\Box\rho$ is a product distribution; moreover, $r(\tau) \to 0$. BBD point out that these facts imply that ρ satisfies strong independence. Finally, we will show that ρ respects within-and-across-player preferences (with the same choice of α). Suppose that $l, m \in \mathcal{N}$, $s'_l \in S_l$, $s''_m \in S_m$, $s^*_l \in BR_l(\rho)$, and $s^*_m \in BR_m(\rho)$ are such that

$$\left[\alpha_l \left(\pi_l(s^*_l|p^k) - \pi_l(s'_l|p^k)\right)\right]_{k=1}^K >_L \left[\alpha_m \left(\pi_m(s^*_m|p^k) - \pi_m(s''_m|p^k)\right)\right]_{k=1}^K. \tag{6}$$

Letting δ be as in the statement of Lemma 12, let T be such that for all $\tau \geq T$, $r(\tau) \in (0,\delta)^{K-1}$. By Lemma 12, for all $\tau \geq T$,

$$\alpha_l \max_{\hat{s}_l \in S_l} \sum_{s\in\bar{S}} (r(\tau)\Box\rho)(s)[\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] > \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s\in\bar{S}} (r(\tau)\Box\rho)(s)[\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)].$$

Because $r(\tau)\Box\rho$ is a product distribution with marginals $\sigma_1(\tau), \ldots, \sigma_N(\tau)$, the above inequality is equivalent to

$$\alpha_l \max_{\hat{s}_l \in S_l}[\pi_l(\sigma(\tau)/\hat{s}_l) - \pi_l(\sigma(\tau)/s'_l)] > \alpha_m \max_{\hat{s}_m \in S_m}[\pi_m(\sigma(\tau)/\hat{s}_m) - \pi_m(\sigma(\tau)/s''_m)],$$

or, equivalently,

$$\alpha_l L_l(\sigma(\tau)/s'_l) > \alpha_m L_m(\sigma(\tau)/s''_m). \tag{7}$$

Define $k' = \min\{k : p^k_l(s'_l) > 0\}$. Because ρ has full support, k' must exist. To accommodate the case of $k' = K$, all that follows remains valid if we define $r^K(\tau) = 0$ for all τ. Because $r(\tau) \to 0$ and $\varepsilon(\tau) \to 0$, there exists some $\tau^* \geq T$ for which $\varepsilon(\tau^*) < [1 - r^{k'}(\tau^*)]p^{k'}_l(s'_l)$. Using Lemma 11 and then this inequality,

$$\begin{aligned}\sigma_l(\tau^*)(s'_l) &\geq r^1(\tau^*)\cdots r^{k'-1}(\tau^*)[1 - r^{k'}(\tau^*)]p^{k'}_l(s'_l) \\ &> \varepsilon(\tau^*)r^1(\tau^*)\cdots r^{k'-1}(\tau^*).\end{aligned}$$

Suppose by way of contradiction that $s''_m \leq_\rho s'_l$, so that $\min\{k : p^k_m(s''_m) > 0\} \geq k'$, and thus $\sigma_m(\tau^*)(s''_m) \leq r^1(\tau^*)\cdots r^{k'-1}(\tau^*)$ by Lemma 11. Then we would have $\sigma_l(\tau^*)(s'_l) > \varepsilon(\tau^*)\sigma_m(s''_m)$. That, taken together with (7), would contradict the fact that $\sigma(\tau^*)$ is an $(\alpha, \varepsilon(\tau^*))$-extended proper equilibrium. It must therefore be the case that $s'_l <_\rho s''_m$. In review, we have shown that equation (6) implies that $s'_l <_\rho s''_m$, which establishes that ρ respects within-and-across-player preferences. □

B.4 Proofs Corresponding to Section 6

Proof of Proposition 9. Assume $M > v_1$, and assume $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$, as in footnote 25. Let b^* be an extended proper equilibrium.[36] By Proposition 8, there exists an LPS ρ that satisfies strong independence, has full support, and respects within-and-across-player preferences for which (ρ, b^*) is a lexicographic Nash equilibrium. Suppose, by way of contradiction, that b^* is not locally envy-free*. Let i' be the largest index for which the locally envy-free* inequality is violated:

$$\kappa_{i'} \left[v_{g(i')} - b^*_{g(i'+1)} \right] < \kappa_{i'-1} \left[v_{g(i')} - \left(b^*_{g(i')} + \frac{1}{m} \right) \right]. \tag{8}$$

The proof consists of two parts. First, we demonstrate that bidders who bid lower than $g(i')$ are sorted by their values for clicks. We then argue that, against ρ, bidder $g(i')$ prefers the deviation $b^*_{g(i')} + \frac{1}{m}$ over its equilibrium bid $b^*_{g(i')}$.

Part 1: Let $i \in \{i'+1, \dots, N\}$. Because i' is the largest index for which the locally envy-free* inequality is violated,

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

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

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

Manipulating these inequalities yields

$$(\kappa_{i-1} - \kappa_i) \left(v_{g(i-1)} - v_{g(i)} \right) + \frac{\kappa_{i-1}}{m} \geq 0$$

We have labeled positions so that $\kappa_{i-1} > \kappa_i$. Thus, if $v_{g(i-1)} < v_{g(i)}$, we would obtain a contradiction: the left-hand side would be bounded above by $-\Delta_\kappa \Delta_v + \frac{\kappa_1}{m}$, which, because $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$, is negative. We conclude that $v_{g(i-1)} > v_{g(i)}$.

Part 2: Suppose bidder $g(i')$ contemplates raising its bid from $b^*_{g(i')}$ to $b^*_{g(i')} + \frac{1}{m}$. This does not make a difference against the equilibrium profile b^* for the following reason. It could make a difference only if $b^*_{g(i'-1)} = b^*_{g(i')} + \frac{1}{m}$. But we cannot have $b^*_{g(i'-1)} = b^*_{g(i')} + \frac{1}{m}$, for otherwise equation (8) would imply that $b^*_{g(i')} + \frac{1}{m}$ would be a profitable deviation from the equilibrium for bidder $g(i')$.

Rather, this deviation makes a difference only if at least one of bidder $g(i')$'s opponents trembles within the set $\{b^*_{g(i')}, b^*_{g(i')} + \frac{1}{m}\}$. Thus, the relevant profiles can be partitioned in the following way:

- *Case 1(a):* $b_{g(i'')} = b^*_{g(i')} + \frac{1}{m}$ for some $i'' < i'$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i''\}$
- *Case 1(b):* $b_{g(i'')} = b^*_{g(i')}$ for some $i'' < i'$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i''\}$
- *Case 2(a):* $b_{g(i'+1)} = b^*_{g(i')} + \frac{1}{m}$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i'+1\}$
- *Case 2(b):* $b_{g(i'+1)} = b^*_{g(i')}$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i'+1\}$
- *Case 3(a):* $b_{g(i'')} = b^*_{g(i')} + \frac{1}{m}$ for some $i'' > i'+1$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i''\}$
- *Case 3(b):* $b_{g(i'')} = b^*_{g(i')}$ for some $i'' > i'+1$ and $b_{g(i)} = b^*_{g(i)}$ for all $i \notin \{i', i''\}$

[36] Throughout the proof, we use Lemma 17, which says that for all $i \in \{1, \dots, N-1\}$, $b^{(i)} > b^{(i+1)}$. It follows that for all $i \in \{1, \dots, N\}$, it is well defined to let $g(i)$ denote the identity of the ith highest bidder.

- *Case 4:* $b_{g(i'')} \in \{b^*_{g(i')}, b^*_{g(i')} + \frac{1}{m}\}$ for some $i'' \neq i'$ and $b_{g(i)} \neq b^*_{g(i)}$ for some $i \notin \{i', i''\}$

We then observe the following points. Together, they show that for each of these relevant profiles: either (i) bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$ than playing $b^*_{g(i')}$ against that profile, or (ii) that profile is infinitely less likely under ρ than another relevant profile against which bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$.

- *Case 4:* Each profile in case 4 corresponds to a profile in one of the other cases in which strictly fewer bidders are deviating (i.e., the profile in which $g(i'')$ is the only deviator). Because ρ satisfies strong independence, it follows that the former must be infinitely less likely under ρ than the latter.[37]

- *Cases 1(a) and 1(b):* Against each profile in case 1(a), bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$. Indeed, without deviating, bidder $g(i')$ receives its equilibrium payoff $\kappa_{i'}\left[v_{g(i')} - b^*_{g(i'+1)}\right]$. Deviating to $b^*_{g(i')} + \frac{1}{m}$, bidder $g(i')$ receives a uniform lottery between that and $\kappa_{i'-1}\left[v_{g(i')} - \left(b^*_{g(i')} + \frac{1}{m}\right)\right]$, which is larger by equation (8). A similar argument establishes the same conclusion for each profile in case 1(b).

  A special case is the profile in case 1(a) for which $i'' = i' - 1$. Clearly, $b^*_{g(i')} + \frac{1}{m}$ is a best response to the equilibrium b^* for bidder $g(i' - 1)$. Therefore, this profile has the following feature: only a single bidder is deviating from b^* and that deviation is to a best response to b^*.

- *Cases 3(a) and 3(b):* For all $i'' > i'+1$, $b^*_{g(i')} + \frac{1}{m}$ is an inferior response to the equilibrium b^* for bidder $g(i'')$. Indeed, bidder $g(i'')$'s equilibrium payoff is $\kappa_{i''}\left[v_{g(i'')} - b^*_{g(i''+1)}\right]$. By deviating from the equilibrium to $b^*_{g(i')} + \frac{1}{m}$, bidder $g(i'')$ receives the payoff $\kappa_{i'}\left[v_{g(i'')} - b^*_{g(i')}\right]$. Because i' was defined as the largest index for which the locally envy-free* inequality is violated, we have

$$\kappa_{i''}\left[v_{g(i'')} - b^*_{g(i''+1)}\right] \geq \kappa_{i''-1}\left[v_{g(i'')} - \left(b^*_{g(i'')} + \frac{1}{m}\right)\right] \tag{9}$$

$$\forall i \in \{i'+1, \ldots, i''-1\} : \kappa_i\left[v_{g(i)} - b^*_{g(i+1)}\right] \geq \kappa_{i-1}\left[v_{g(i)} - \left(b^*_{g(i)} + \frac{1}{m}\right)\right] \tag{10}$$

  From part 1 of this proof, we have that $i \in \{i'+1, \ldots, i''-1\}$, $v_{g(i)} \geq v_{g(i'')} + \Delta_v$. Combining that with (10), we obtain

$$\begin{aligned}\forall i \in \{i'+1, \ldots, i''-1\} : \kappa_i\left[v_{g(i'')} - b^*_{g(i+1)}\right] &\geq \kappa_{i-1}\left[v_{g(i'')} - \left(b^*_{g(i)} + \frac{1}{m}\right)\right] + \Delta_\kappa \Delta_v \\ &\geq \kappa_{i-1}\left[v_{g(i'')} - b^*_{g(i)}\right] + \Delta_\kappa \Delta_v - \frac{\kappa_1}{m} \end{aligned} \tag{11}$$

  Summing (9) and (11), then canceling like terms, we obtain

$$\kappa_{i''}\left[v_{g(i'')} - b^*_{g(i''+1)}\right] \geq \kappa_{i'}\left[v_{g(i'')} - b^*_{g(i'+1)}\right] + (i''-i'-1)\left(\Delta_\kappa \Delta_v - \frac{2\kappa_1}{m}\right) + (i''-i'-2)\frac{\kappa_1}{m}$$

[37] This can be shown formally using Lemmas 14 and 16. Let b' denote a profile in case 4. Let b'' denote any corresponding profile in one of the other cases in which strictly fewer bidders are deviating. Let $g(i'')$ denote the deviating bidder in b''. Let $g(i''')$ denote some bidder who deviates in b' but not in b''. Then $b''_{g(i''')} >_\rho b'_{g(i''')}$ by Lemma 14*(ii)*. Because $b''_{g(i'')} = b'_{g(i'')}$, we also have $b''_{g(i'')} =_\rho b'_{g(i'')}$. And we have $b''_{g(i)} \geq_\rho b'_{g(i)}$ for all $i \notin \{i'', i'''\}$ by Lemma 14*(i)*. Combining all this and applying Lemma 16, we have $b'' >_\rho b'$.

Using $b^*_{g(i'+1)} < b^*_{g(i')}$, $i'' \geq i'+2$, and $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$, we obtain

$$\kappa_{i''}\left[v_{g(i'')} - b^*_{g(i''+1)}\right] > \kappa_{i'}\left[v_{g(i'')} - b^*_{g(i')}\right],$$

as desired. Because ρ satisfies strong independence and respects within-and-across-player preferences, it follows that each profile in case 3(a) is infinitely less likely under ρ than the special case mentioned in the previous bullet (i.e., the profile in case 1(a) for which $i'' = i'-1$).[38] A similar argument establishes the same conclusion for each profile in case 3(b).

- *Cases 2(a) and 2(b):* For bidder $g(i'+1)$, $b^*_{g(i')} + \frac{1}{m}$ might be either an inferior response to the equilibrium b^* or a best response to b^*. In the case of an inferior response, the profile in case 2(a) can be handled in the same way as the profiles in case 3(a), which were discussed in the previous bullet. It remains to consider the case where $b^*_{g(i')} + \frac{1}{m}$ is a best response to b^* for the deviating bidder $b_{g(i'+1)}$.

  Bidder $g(i'+1)$'s equilibrium payoff is $\kappa_{i'+1}\left[v_{g(i'+1)} - b^*_{g(i'+2)}\right]$. By deviating from the equilibrium to $b^*_{g(i')} + \frac{1}{m}$, bidder $g(i'+1)$ receives the payoff $\kappa_{i'}\left[v_{g(i'+1)} - b^*_{g(i')}\right]$. We therefore proceed under the assumption that

$$\kappa_{i'+1}\left[v_{g(i'+1)} - b^*_{g(i'+2)}\right] = \kappa_{i'}\left[v_{g(i'+1)} - b^*_{g(i')}\right]. \tag{12}$$

  Against the profile in case 2(a), bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$. Indeed, without deviating, bidder $g(i')$ receives the payoff $\kappa_{i'+1}\left[v_{g(i')} - b^*_{g(i'+2)}\right]$. Deviating to $b^*_{g(i')} + \frac{1}{m}$, bidder $g(i')$ receives a uniform lottery between that and $\kappa_{i'}\left[v_{g(i')} - \left(b^*_{g(i')} + \frac{1}{m}\right)\right]$. To establish the conclusion, we argue:

$$\begin{aligned}
\kappa_{i'}\left[v_{g(i')} - \left(b^*_{g(i')} + \frac{1}{m}\right)\right] &= \kappa_{i'}\left[v_{g(i'+1)} - b^*_{g(i')}\right] - \frac{\kappa_{i'}}{m} + \kappa_{i'}(v_{g(i')} - v_{g(i'+1)}) \\
&= \kappa_{i'+1}\left[v_{g(i'+1)} - b^*_{g(i'+2)}\right] - \frac{\kappa_{i'}}{m} + \kappa_{i'}(v_{g(i')} - v_{g(i'+1)}) \\
&= \kappa_{i'+1}\left[v_{g(i')} - b^*_{g(i'+2)}\right] - \frac{\kappa_{i'}}{m} + (\kappa_{i'} - \kappa_{i'+1})(v_{g(i')} - v_{g(i'+1)}) \\
&\geq \kappa_{i'+1}\left[v_{g(i')} - b^*_{g(i'+2)}\right] - \frac{\kappa_1}{m} + \Delta_\kappa \Delta_v \\
&> \kappa_{i'+1}\left[v_{g(i')} - b^*_{g(i'+2)}\right],
\end{aligned}$$

  as desired. In the above, the first and third steps are algebra. The second step uses (12). The fourth step uses $\kappa_{i'} \leq \kappa_1$, $\kappa_{i'} - \kappa_{i'+1} \geq \Delta_\kappa$, and $v_{g(i')} - v_{g(i'+1)} \geq \Delta_v$. (Recall that step one of this proof established $v_{g(i')} > v_{g(i'+1)}$.) The fifth step uses $m > \frac{4\kappa_1}{\Delta_\kappa \Delta_v}$. A similar argument establishes the same conclusion for the profile in case 2(b).

In summary, we have shown that for each of the relevant profiles: either (*i*) bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$ than playing $b^*_{g(i')}$ against that profile, or (*ii*) that profile is infinitely less likely under ρ than another relevant profile against which bidder $g(i')$ is strictly better off deviating to $b^*_{g(i')} + \frac{1}{m}$. It follows that bidder $g(i')$ lexicographically prefers $b^*_{g(i')} + \frac{1}{m}$ over $b^*_{g(i')}$ against ρ, which contradicts (ρ, b^*) being a lexicographic Nash equilibrium. $\square$

[38] This can be shown formally using Lemmas 13, 14, and 16. Let b'' denote a profile in case 3(a). Let b' denote the aforementioned profile in case 1(a). Let $g(i'')$ denote the deviating bidder in b''. By definition, $g(i'-1)$ is the deviating bidder in b'. Then $b'_{g(i'-1)} >_\rho b''_{g(i'')}$ by Lemma 13. And we have $b'_{\mathcal{N}\setminus\{g(i'-1)\}} \geq_\rho b''_{\mathcal{N}\setminus\{g(i'')\}}$ by Lemma 14(*i*). Combining all this and applying Lemma 16, we have $b' >_\rho b''$.

C Connections to Test-Set Equilibrium

Test-set equilibrium (Milgrom and Mollner, 2018), like extended proper equilibrium, can be interpreted as imposing across-player restrictions on the likelihood of trembles. But in general, neither concept implies the other. After reviewing the definition of test-set equilibrium, this appendix further discusses its connection to extended proper equilibrium.

C.1 Definition of Test-Set Equilibrium

We begin by reviewing the formal definition of test-set equilibrium. Given a mixed strategy profile σ, let

$$T(\sigma) = \bigcup_{n=1}^{N} \{\sigma/s_n \,|\, s_n \in BR_n(\sigma)\}\,.$$

We refer to $T(\sigma)$ as the *test set* associated with σ. In words, $T(\sigma)$ consists of the strategy profiles in which a single player deviates from σ and to a strategy that is a best response to 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 opponent strategy profiles consistent with some $\sigma' \in T(\sigma)$:

Definition 12. A mixed strategy profile σ satisfies the *test-set condition* if and only if, for all $n \in \mathcal{N}$, there is no $\hat{\sigma}_n \in \Delta_n$ 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 $\sigma' \in T(\sigma)$, $\pi_n(\sigma'/\hat{\sigma}_n) > \pi_n(\sigma'/\sigma_n)$.

Definition 13. 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.

For finite games, one can use the separating hyperplane theorem to recast the undominatedness requirements of test-set equilibrium in terms of trembles. Test-set equilibrium implicitly requires each player to use a strategy that remains optimal even though its opponents may tremble—for some mode of trembling with the structure that every strategy profile outside the test set is overwhelmingly less likely than every profile in the test set.

Given the definition of the test set, test-set equilibrium effectively requires any unilateral deviation to a strategy that is a best response to equilibrium to be treated as overwhelmingly more likely than any unilateral deviation to a strategy that is an inferior response to the equilibrium, whether by the same or by another player. As observed in the text, extended proper equilibrium implies the same. But there are also differences, and in general, neither solution concept implies the other in games with three or more players.

- Extended proper equilibrium may prohibit some trembles that are implicitly allowed by test-set equilibrium, and it therefore can sometimes be more demanding. For example, given two strategy profiles in the test set, extended proper equilibrium may require one to be overwhelming less likely than another. This possibility was illustrated by the example of Figure 2(b). The same can also be true for strategy profiles outside the test set, which is illustrated by the example of Figure 5, below.

- Conversely, extended proper equilibrium may allow some trembles that are implicitly prohibited by test-set equilibrium, and it therefore can sometimes be less demanding. For example, extended proper equilibrium is consistent with a deviation involving two non-equilibrium best responses (and therefore outside the test set) being more likely than another deviation involving only a single non-equilibrium best response (and therefore in the test set). This possibility is central to the example of Figure 3 in which no test-set equilibrium exists.

C.2 Examples

In this appendix, we formally apply the definition of test-set equilibrium to the example games discussed in the main text.

Figure 1. First, return to the game in Figure 1. As mentioned, $(Up, Left, East)$ is a proper equilibrium of this game, but it is not extended proper. It is also not a test-set equilibrium. Indeed, *West* weakly dominates *East* in the test set:

$$T(Up, Left, East) = \{(Up, Left, East), (Up, Right, East), (Up, Left, West)\}.$$

In particular, $\pi_{geo}(\sigma'/West) \geq \pi_{geo}(\sigma'/East)$ for all $\sigma' \in T(Up, Left, East)$, and the inequality is strict for $\sigma' = (Up, Right, East)$. An analogous argument also rules out equilibria in which Geo mixes between *East* and *West*, leaving $(Up, Left, West)$ as the lone remaining candidate. And indeed, it is easily checked that $(Up, Left, West)$ is in fact a test-set equilibrium.

Thus, test-set equilibrium and extended proper equilibrium coincide in this game. In words, test-set equilibrium implicitly requires Geo to believe that the test set element $(Up, Right, \sigma_{geo})$ is more likely than $(Down, Left, \sigma_{geo})$, which is outside the test set. As observed in the text, extended proper equilibrium requires the same thing.

Figure 2(b). Next, return to the game in Figure 2(b). As mentioned, (S_1, S_2, S_3) is not an extended proper equilibrium of this game. But it is a proper equilibrium, and it is also a test-set equilibrium. In this case, players 2 and 3 are each using their dominant strategies S_2 and S_3, so we need only check that player 1's strategy S_1 is neither weakly dominated in the test set nor weakly dominated in the game. The test set is

$$T(S_1, S_2, S_3) = \{(S_1, S_2, S_3), (C_1, S_2, S_3), (S_1, C_2, S_3), (S_1, S_2, C_3)\}.$$

Letting σ' denote the test set element (S_1, C_2, S_3), we have $\pi_1(\sigma'/S_1) > \pi_1(\sigma'/C_1)$, so that S_1 is indeed not dominated in this way. In words, test-set equilibrium implicitly allows player 1 to believe that opponents' trembling will be such that the test set element (S_1, C_2, S_3) is more likely than (S_1, S_2, C_3). But as observed in the text, extended proper equilibrium requires the reverse.

Figure 3. Next, return to the game in Figure 3. As mentioned, $(Up, Left, West)$ is the unique Nash equilibrium of this game, and it is therefore also a perfect, proper, and extended proper equilibrium. But it is not a test-set equilibrium. Indeed, *East* weakly dominates *West* in the test set:

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

In particular, $\pi_{geo}(\sigma'/East) \geq \pi_{geo}(\sigma'/West)$ for all $\sigma' \in T(Up, Left, West)$, and the inequality is strict for $\sigma' = (Up, Center, West)$. In words, test-set equilibrium implicitly requires Geo to believe that the test set element $(Up, Center, West)$ is more likely than $(Down, Right, West)$, which is outside the test set. But as observed in the text, extended proper equilibrium allows the reverse.

Figure 5. Finally, consider a new example: the game in Figure 5. For Row, *Down* is strictly dominated by *Up*. For Column, *Right* is strictly dominated by *Center*, which in turn is strictly dominated by *Left*. In consequence, $(Up, Left)$ is the unique Nash equilibrium of this game—it is also a perfect, proper, extended proper, and test-set equilibrium. But what we wish to illustrate with this game are two ways in which proper equilibrium (and therefore also extended proper equilibrium) can sometimes prohibit trembles that test-set equilibrium implicitly allows:

- First, proper equilibrium requires a deviation by Column to *Right* to be less likely than a deviation to *Center*. Formally, if σ^t is an ε_t-proper equilibrium then

  $$\sigma^t_{col}(Right) \leq \varepsilon_t \sigma^t_{col}(Center).$$

  In contrast, test-set equilibrium implicitly allows the reverse: both *Center* and *Right* are inferior responses to the equilibrium, so each results in a strategy profile outside the test set, and the concept does not rank profiles that are both outside the test set.

- Second, proper equilibrium requires the joint deviation to $(Down, Right)$ to be less likely than either unilateral deviation. Essentially, this is because proper equilibrium requires the trembles of different players to be statistically independent, so that σ^t induces a product distribution over the set of strategy profiles. Formally, we have $\sigma^t_{col}(Right) \leq \varepsilon_t \sigma^t_{col}(Left)$ and $\sigma^t_{row}(Down) \leq \varepsilon_t \sigma^t_{row}(Up)$, which together imply

  $$\sigma^t_{row}(Down)\sigma^t_{col}(Right) \leq \varepsilon_t \max\{\sigma^t_{row}(Down)\sigma^t_{col}(Left), \sigma^t_{row}(Up)\sigma^t_{col}(Right)\}.$$

  In contrast, test-set equilibrium implicitly allows the reverse: both a unilateral deviation to an inferior response and the joint deviation result in strategy profiles outside the test set, and the concept does not rank profiles that are both outside the test set. Thus, test-set equilibrium does not embed all the restrictions that are implied if players tremble independently of each other. In other words, the trembles associated with a test-set equilibrium may be certain *correlated* strategy profiles.

In these ways, proper equilibrium (and therefore also extended proper equilibrium) can sometimes prohibit trembles that test-set equilibrium implicitly allows. What is more, these additional restrictions could translate into sharper predictions for equilibrium play if there were a third player whose best strategy depended on the relative likelihoods of the trembles considered above.[39]

Figure 5: A two player game†

	Left	*Center*	*Right*
Up	$1, 2$	$1, 1$	$1, 0$
Down	$0, 2$	$0, 1$	$0, 0$

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

[39] In fact, the GSP auction example discussed in Appendix C.3 is such a game. As we discuss, the predictions of test-set equilibrium are less sharp—precisely because it does not embed all the restrictions that are implied if players tremble independently of each other.

C.3 Another GSP Auction Example

In the GSP auction game, extended proper equilibrium and test-set equilibrium similarly imply locally envy-free equilibrium, in the sense described in Section 6. But despite this similarity, extended proper equilibrium can deliver conclusions that are stronger than those that test-set equilibrium can deliver.

To illustrate, we consider another specific example. There are three bidders, with per-click values $(v_1, v_2, v_3) = (6, 2, 1)$. There are two ad positions, with click rates $(\kappa_1, \kappa_2) = (2, 1)$. For the discretized bid set, it suffices to let $m = 3$ and $M = 4$.

The profile $b^* = (3, 2, 1)$ is a test-set equilibrium.[40] But this bid profile is not an extended proper equilibrium. In fact, it is not even trembling-hand perfect. The subsequent analysis will show the problem to be that bidder 1 is bidding too low if the trembles of its different opponents must be statistically independent.

Building on Proposition 6, suppose to the contrary that ρ is an LPS that satisfies strong independence and has full support for which (ρ, b^*) is a lexicographic Nash equilibrium. To see that this produces a contradiction, suppose bidder 1 contemplates raising its bid from $b_1^* = 3$ to $\frac{10}{3}$. This change makes a difference only if at least one of its opponents trembles within the set $\{3, \frac{10}{3}\}$. Thus, the relevant profiles can be partitioned into the following way:

- *Case 1(a):* $b_2 = \frac{10}{3}$ and $b_3 \in \{0, \frac{1}{3}\}$
- *Case 1(b):* $b_2 \in \{0, \frac{1}{3}\}$ and $b_3 = \frac{10}{3}$
- *Case 2(a):* $b_2 = \frac{10}{3}$ and $b_3 = 1$
- *Case 2(b):* $b_2 = 2$ and $b_3 = \frac{10}{3}$
- *Case 3:* all remaining profiles in which $b_2 \in \{3, \frac{10}{3}\}$ and/or $b_3 \in \{3, \frac{10}{3}\}$

In general, bidder 1 prefers to make a lower bid instead of a higher bid only when one of its opponents bids relatively high and the other opponent bids relatively low, so that position 2 can be obtained for a relatively low price while position 1 can be obtained only for a relatively high price. Computing, we find that cases 1(a) and 1(b) are the only relevant profiles against which bidder 1 is strictly worse off deviating to $\frac{10}{3}$. For example, consider $b_2 = \frac{10}{3}$ and $b_3 = 0$:

$$\pi_1(\tfrac{10}{3}, \tfrac{10}{3}, 0) = \tfrac{17}{3} < 6 = \pi_1(3, \tfrac{10}{3}, 0).$$

However, note that each profile in case 1(a) involves both bidders 2 and 3 deviating from their equilibrium bids. In contrast, the profile in case 2(a) involves the same deviation by bidder 2 but no deviation by bidder 3, which, by strong independence, must be infinitely more likely under ρ. Moreover, against the profile in case 2(a), bidder 1 is strictly better off deviating to $\frac{10}{3}$:

$$\pi_1(\tfrac{10}{3}, \tfrac{10}{3}, 1) = \tfrac{31}{6} > 5 = \pi_1(3, \tfrac{10}{3}, 1).$$

Similarly, each profile in case 1(b) involves both bidders 2 and 3 deviating from their equilibrium bids. In contrast, the profile in case 2(b) involves the same deviation by bidder 3 but no deviation

[40] No bidder is using a weakly dominated bid. Indeed: for bidder 1, the dominated bids are those greater than 6 and those less than 3; for bidder 2, the dominated bids are those greater than 2 and those less than 1; and for bidder 3, the dominated bids are those greater than 1 and those less than $\frac{1}{2}$.

Furthermore, no bidder is using a bid that is weakly dominated in the test set. Bidder 1 best responds to its opponents' equilibrium bids $(2, 1)$ with any bid greater than 2. Bidder 2 best responds to its opponents' equilibrium bids $(3, 1)$ with any bid greater than 1 and less than 3. Bidder 3 best responds to its opponents' equilibrium bids $(3, 2)$ with any bid less than 2. From these observations, the test set can be constructed and the test-set condition checked.

by bidder 2, which, by strong independence, must be infinitely more likely under ρ. And against the profile in case 2(b), bidder 1 is strictly better off deviating to $\frac{10}{3}$:

$$\pi_1(\tfrac{10}{3}, 2, \tfrac{10}{3}) = \tfrac{14}{3} > 4 = \pi_1(3, 2, \tfrac{10}{3}).$$

We have therefore shown that for each of the relevant profiles, either (i) bidder 1 is strictly better off deviating to $\frac{10}{3}$, or (ii) that profile is infinitely less likely under ρ than another relevant profile against which bidder 1 is strictly better off deviating to $\frac{10}{3}$. It follows that $\frac{10}{3}$ is a profitable deviation for bidder 1 against ρ, which contradicts (ρ, b^*) being a lexicographic Nash equilibrium.

The key issue is that test-set equilibrium does not require the trembles of different players to be statistically independent. The profiles in cases 1(a) and 2(a) are all outside the test set, as all involve bidder 2 deviating to an inferior response. Thus, test-set equilibrium implicitly allows bidder 1 to believe that either case is the more likely one. And bidder 1 can justify $b_1^* = 3$ with a belief that case 1(a) is more likely. In contrast, perfect equilibrium (and hence extended proper equilibrium) does require players to tremble independently of each other, and thus it requires case 2(a) to be more likely, so that bidder 1 cannot justify $b_1^* = 3$. Indeed, case 2(a) involves only bidder 2 deviating from equilibrium play, while case 1(a) involves bidder 2 making the same deviation as well as an additional deviation by bidder 3.

C.4 Weak Test-Set Equilibrium

As some of the preceding examples have illustrated (e.g., the game in Figure 3), extended proper equilibrium does not always imply test-set equilibrium. Nevertheless, extended proper equilibrium does always imply a weakened version of test-set equilibrium, which we define below. We provide this new definition as it may facilitate comparisons between extended proper equilibrium and test-set equilibrium.

C.4.1 Formalities

For any strategy profile σ, let

$$\hat{T}(\sigma) = \bigcup_{J \subseteq \mathcal{N}} \{\sigma/s_J \,|\, (\forall n \in J)\, s_n \in BR_n(\sigma)\}$$

In words, $\hat{T}(\sigma)$ consists of all strategy profiles that can be constructed from σ by replacing the strategies of any number of players with best responses to σ. We will refer to $\hat{T}(\sigma)$ as the *generalized test set* associated with σ. Note that $T(\sigma) \subseteq \hat{T}(\sigma)$. We also define a partial order on $\hat{T}(\sigma)$ in the following way.

Definition 14. For $\sigma', \sigma'' \in \hat{T}(\sigma)$, $\sigma'' \rhd \sigma'$ if and only if

(i) for all $n \in \mathcal{N}$, $\sigma''_n \neq \sigma_n \implies \sigma'_n = \sigma''_n$; and

(ii) for some $n \in \mathcal{N}$, $\sigma''_n = \sigma_n$ and $\sigma'_n \notin \text{supp}(\sigma_n)$.

In words, $\sigma'' \rhd \sigma'$ if σ' entails all the deviations from σ that are entailed by σ'', as well as at least one additional deviation to a strategy that is played with zero probability under σ. In the case where σ is a pure strategy profile: (i) σ itself is an element of $T(\sigma)$, and it is the unique maximum element under $\rhd$; (ii) next in the order are the remaining elements of the test set, $T(\sigma) \setminus \{\sigma\}$, which are single deviations from σ; (iii) finally, we have the remaining elements of the generalized test set, $\hat{T}(\sigma) \setminus T(\sigma)$, which are joint deviations from σ. With this notation in place, we can state the following definitions:

Definition 15. A mixed-strategy profile σ satisfies the *weak test-set condition* if and only if, for all $n \in \mathcal{N}$, there is no $\hat{\sigma}_n \in \Delta_n$ such that both

(*i*) for all $\sigma' \in \hat{T}(\sigma)$ with $\pi_n(\sigma'/\hat{\sigma}_n) < \pi_n(\sigma'/\sigma_n)$, there exists a $\sigma'' \in \hat{T}(\sigma)$ such that $\sigma'' \rhd \sigma'$ and $\pi_n(\sigma''/\hat{\sigma}_n) > \pi_n(\sigma''/\sigma_n)$; and

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

Definition 16. A mixed strategy profile σ is a *weak test-set equilibrium* if and only if it is a Nash equilibrium in undominated strategies that satisfies the weak test-set condition.

The following result implies that, as the nomenclature suggests, every test-set equilibrium is a weak test-set equilibrium.

Proposition 18. *Any Nash equilibrium that satisfies the test-set condition also satisfies the weak test-set condition.*

Proof of Proposition 18. We prove the contrapositive. Suppose σ is a Nash equilibrium that does not satisfy the weak test-set condition. Thus, there exists some $n \in \mathcal{N}$ and some $\hat{\sigma}_n \in \Delta_n$ that satisfies both conditions (*i*) and (*ii*) of Definition 15. We will show that the same n and the same $\hat{\sigma}_n$ satisfy the conditions of Definition 12. Because condition (*ii*) of Definition 12 is identical to condition (*ii*) of Definition 15, we need only check condition (*i*). Suppose, by way of contradiction, that condition (*i*) of Definition 12 does not hold. That is, there exists a $\sigma' \in T(\sigma)$ such that $\pi_n(\sigma'/\hat{\sigma}_n) < \pi_n(\sigma'/\sigma_n)$. But for $\sigma' \in T(\sigma)$, we can have $\sigma'' \rhd \sigma'$ only if $\sigma'' = \sigma$.[41] And $\pi_n(\sigma/\hat{\sigma}_n) > \pi_n(\sigma/\sigma_n)$ contradicts the assumption that σ is a Nash equilibrium. □

C.4.2 Example: The Game in Figure 3

To illustrate the definitions that we have just given, consider again the game in Figure 3. As mentioned, (*Up*, *Left*, *West*) is the unique Nash equilibrium of this game. We previously pointed out that the test set associated with this strategy profile is

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

Let us now construct the the generalized test set associated with this strategy profile. Inspecting the payoff matrix, every strategy is a best response to (*Up*, *Left*, *West*). Therefore, $\hat{T}(\mathit{Up}, \mathit{Left}, \mathit{West})$ is in fact the entire set of pure strategy profiles. And in this case, the partial order $\rhd$ on $\hat{T}(\mathit{Up}, \mathit{Left}, \mathit{West})$ can be represented by the directed graph depicted in Figure 6, where $\sigma'' \rhd \sigma'$ iff there exists a directed path from σ'' to σ'.

Recall that (*Up*, *Left*, *West*) did not satisfy the test-set condition because *East* met the requirements set out for $\hat{\sigma}_n$ set out in Definition 12. Let us now check whether it also meets the requirements for $\hat{\sigma}_n$ set out in Definition 15. It does not. In particular, consider $\sigma' = (\mathit{Down}, \mathit{Right}, \mathit{West})$. This profile is not an element of $T(\sigma)$, and so it was not relevant to Definition 12. But it is an element of $\hat{T}(\sigma)$, and so it must now be considered when checking condition (*i*) of Definition 15. With this choice of σ', $\pi_{geo}(\sigma'/\mathit{East}) < \pi_{geo}(\sigma'/\mathit{West})$. Moreover, there is no $\sigma'' \rhd \sigma'$ for which $\pi_{geo}(\sigma''/\mathit{East}) > \pi_{geo}(\sigma''/\mathit{West})$.[42] Thus, *East* does not meet the requirements. Furthermore, as this argument suggests, it can be checked that (*Up*, *Left*, *West*) does indeed satisfy the weak test-set condition and is a weak test-set equilibrium.

[41] And this is possible only if σ is itself an element of $\hat{T}(\sigma)$.

[42] In particular, we have $\pi_{geo}(\sigma''/\mathit{East}) > \pi_{geo}(\sigma''/\mathit{West})$ iff $(\sigma''_{row}, \sigma''_{col}) = (\mathit{Up}, \mathit{Center})$. But it is apparent from Figure 6 that no such σ'' also satisfies $\sigma'' \rhd (\mathit{Down}, \mathit{Right}, \mathit{West})$.

Figure 6: The partial order $\triangleright$ on $\hat{T}(\mathit{Up}, \mathit{Left}, \mathit{West})$ for the game in Figure 3

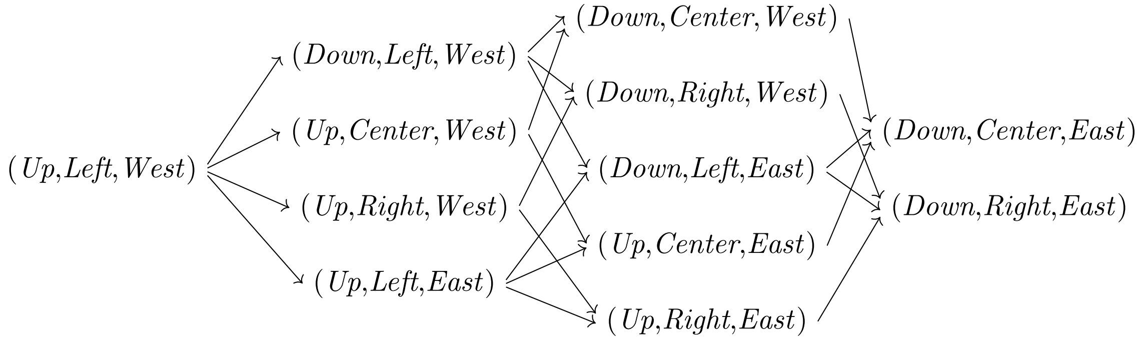

C.4.3 Main Result

We had previously said that $(\mathit{Up}, \mathit{Left}, \mathit{West})$ is an extended proper equilibrium of the game in Figure 3. And we have just shown that it is also a weak test-set equilibrium. The main result of this appendix states that this relationship is in fact true in general:

Proposition 19. *Every extended proper equilibrium is a weak test-set equilibrium.*

Roughly speaking, we prove the proposition by arguing that extended proper equilibrium requires (i) each strategy profile outside $\hat{T}(\sigma)$ to be infinitely less likely than every strategy profile inside $T(\sigma)$,[43] and (ii) if $\sigma', \sigma'' \in \hat{T}(\sigma)$ are such that $\sigma'' \triangleright \sigma'$, then σ' is infinitely less likely than σ''.

In addition to facilitating comparisons between extended proper equilibrium and test-set equilibrium, Proposition 19 might also be useful in its own right. For example, it could have been used to simplify the proof of Proposition 9: whereas our direct proof showed that every pure extended proper equilibrium of the GSP auction game is locally envy-free*, given Proposition 19, it would have sufficed to show that this property is possessed by the pure weak test-set equilibria.

Proof of Proposition 19. Suppose σ is an extended proper equilibrium. By Proposition 8, there exists some LPS ρ that satisfies strong independence, has full support, respects within-and-across-player preferences, and for which (ρ, σ) is a lexicographic Nash equilibrium. It is immediate that σ is a Nash equilibrium in undominated strategies, so it only remains to verify that the weak test-set condition is satisfied.

Suppose, by way of contradiction, that the weak test-set condition does not hold. Thus, there exists some $n \in \mathcal{N}$ and some $\hat{\sigma}_n \in \Delta_n$ that satisfies both conditions *(i)* and *(ii)* of Definition 15. Condition *(ii)* says there exists a $\sigma' \in T(\sigma)$ such that $\pi_n(\sigma'/\hat{\sigma}_n) > \pi_n(\sigma'/\sigma_n)$. From the definition of $T(\sigma)$, there must exist the following: some player $m_0 \neq n$, some $s'_{m_0} \in BR_{m_0}(\sigma)$, and for all $m \neq m_0$, some $s'_m \in S_m$ with $\sigma_m(s'_m) > 0$, such that $\pi_n(s'/\hat{\sigma}_n) > \pi_n(s'/\sigma_n)$. On the other hand, suppose a pure strategy profile s'' is such that $\pi_n(s''/\hat{\sigma}_n) < \pi_n(s''/\sigma_n)$. There are two cases, depending on whether $s''/\sigma_n \in \hat{T}(\sigma)$:

- *Case 1:* $s''/\sigma_n \notin \hat{T}(\sigma)$. By the definition of $\hat{T}(\sigma)$, there exists some player $m_1 \neq n$ such that $s''_{m_1} \notin BR_{m_1}(\sigma)$. On the other hand, we had $s'_{m_0} \in BR_{m_0}(\sigma)$. Because ρ respects within-and-across-player preferences, it follows from Lemma 13 that $s'_{m_0} >_\rho s''_{m_1}$. Because $\sigma_m(s'_m) > 0$ for all $m \neq m_0$, it follows from Lemma 14*(i)* that $s'_{\mathcal{N}\setminus\{m_0\}} \geq_\rho s''_{\mathcal{N}\setminus\{m_1\}}$. Applying Lemma 16, we conclude $s' >_\rho s''$.

[43] On the other hand, extended proper equilibrium does not require the stronger condition that each strategy profile outside $T(\sigma)$ be infinitely less likely than every strategy profile inside $T(\sigma)$. This is why extended proper equilibrium does not imply the test-set condition itself.

- *Case 2:* $s''/\sigma_n \in \hat{T}(\sigma)$. Condition (*i*) of Definition 15 says there exists a $\sigma''' \in \hat{T}(\sigma)$ such that $\sigma''' \triangleright s''/\sigma_n$ and $\pi_n(\sigma'''/\hat{\sigma}_n) > \pi_n(\sigma'''/\sigma_n)$. By the definition of the partial order $\triangleright$, there exist some $J \subseteq \mathcal{N} \setminus \{n\}$ and some $m_2 \in \mathcal{N} \setminus (J \cup \{n\})$ with $s''_{m_2} \notin \text{supp}(\sigma_{m_2})$ such $\sigma''' = (s''_J, \sigma_{\mathcal{N}\setminus J})$. Furthermore, for all $m \in \mathcal{N} \setminus J$, there must exist some $s'''_m \in S_m$ with $\sigma_m(s'''_m) > 0$ such that $\pi_n(s''_J, s'''_{\mathcal{N}\setminus J}/\hat{\sigma}_n) > \pi_n(s''_J, s'''_{\mathcal{N}\setminus J}/\sigma_n)$. Because $\sigma_{m_2}(s'''_{m_2}) > 0$ but $\sigma_{m_2}(s''_{m_2}) = 0$, it follows from Lemma 14(*ii*) that $s'''_{m_2} >_\rho s''_{m_2}$. Because $\sigma_m(s'''_m) > 0$ for all $m \in \mathcal{N} \setminus (J \cup \{m_2\})$, it follows from Lemma 14(*i*) that $s'''_{\mathcal{N}\setminus(J\cup\{m_2\})} \geq_\rho s''_{\mathcal{N}\setminus(J\cup\{m_2\})}$. Applying Lemma 16, we have $s'''_{\mathcal{N}\setminus J} >_\rho s''_{\mathcal{N}\setminus J}$, and $(s''_J, s'''_{\mathcal{N}\setminus J}) >_\rho s''$.

In summary, for every pure strategy profile s'' for which $\pi_n(s''/\hat{\sigma}_n) < \pi_n(s''/\sigma_n)$, we have shown that there exists another profile that is infinitely more likely under the LPS ρ and under which the inequality is reversed (i.e., s' in case 1 and $(s''_J, s'''_{\mathcal{N}\setminus J})$ in case 2). We conclude that player n lexicographically prefers $\hat{\sigma}_n$ over σ_n against ρ. This contradicts (ρ, σ) being a lexicographic Nash equilibrium. □

D An Illustration of the Proof of Theorem 4

Let Γ be the game in Figure 7. For Column, *Left* always pays 2, *Center* always pays 1, and *Right* always pays 0. For Row, *Up* always pays 1; *Down* also pays 1 if Column plays its dominant strategy of *Left* but pays 0 otherwise. In consequence, $\sigma = (Up, Left)$ is clearly an extended proper equilibrium. We use this game to illustrate our strategy for proving the sufficiency of extended proper equilibrium in Theorem 4. To do so, we demonstrate that $\bar{\sigma} = (Up, Left)$—which clearly has σ as its projection—is a symmetrically proper strategy of $\bar{\Gamma}$.

Figure 7: A two player game†

	Left	*Center*	*Right*
Up	1, 2	1, 1	1, 0
Down	1, 2	0, 1	0, 0

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

The following strategy profile is an (α, δ)-extended proper equilibrium with $\alpha = (1, 1)$ and all sufficiently small $\delta > 0$:

$$\sigma^\delta_{row} = (1 - \delta^2, \delta^2) \quad \sigma^\delta_{col} = (1 - \delta^3 - \delta^4, \delta^3, \delta^4).$$

In particular, *Down* has vanishingly small loss and is played with probability δ^2; *Center* has a loss of 1 and is played with probability δ^3; and *Right* has a loss of 2 and is played with probability δ^4. Let ϕ^δ be the distribution over $\bar{S}$ induced by $\sigma^\delta = (\sigma^\delta_{row}, \sigma^\delta_{col})$:

$$\phi^\delta(Up, Left) = 1 - \delta^2 - \delta^3 - \delta^4 + \delta^5 + \delta^6 \qquad \phi^\delta(Down, Left) = \delta^2 - \delta^5 - \delta^6$$
$$\phi^\delta(Up, Center) = \delta^3 - \delta^5 \qquad \phi^\delta(Down, Center) = \delta^5$$
$$\phi^\delta(Up, Right) = \delta^4 - \delta^6 \qquad \phi^\delta(Down, Right) = \delta^6$$

Mechanically, ϕ^δ is a distribution over $\bar{S}$ that has σ^δ as its projection (although not the unique such distribution). As remarked in the text, $(\phi^\delta, \ldots, \phi^\delta)$ might not itself be an ε-proper equilibrium of

$\bar{\Gamma}$. That is indeed the case here. To see that, let $\bar{s}' = (\textit{Down}, \textit{Center})$ and $\bar{s}'' = (\textit{Up}, \textit{Right})$. Then $\bar{\pi}_1(\bar{s}', \phi^\delta) > \bar{\pi}_1(\bar{s}'', \phi^\delta)$ for all sufficiently small $\delta > 0$:

$$\bar{\pi}_1(\bar{s}', \phi^\delta) = \frac{1}{2}\left[\pi_{row}\left(\textit{Down}, \mathrm{proj}_{col}(\phi^\delta)\right) + \pi_{col}\left(\mathrm{proj}_{row}(\phi^\delta), \textit{Center}\right)\right] = \frac{1}{2}\left[(1 - \delta^3 - \delta^4) + 1\right]$$
$$\bar{\pi}_1(\bar{s}'', \phi^\delta) = \frac{1}{2}\left[\pi_{row}\left(\textit{Up}, \mathrm{proj}_{col}(\phi^\delta)\right) + \pi_{col}\left(\mathrm{proj}_{row}(\phi^\delta), \textit{Right}\right)\right] = \frac{1}{2}\left[1 + 0\right]$$

And whereas ε-proper equilibrium would therefore require $\phi^\delta(\textit{Up}, \textit{Right}) \leq \varepsilon \cdot \phi^\delta(\textit{Down}, \textit{Center})$, we instead have $\phi^\delta(\textit{Up}, \textit{Right}) > \phi^\delta(\textit{Down}, \textit{Center})$ when $\delta > 0$ is small.

More generally, given a distribution over $\bar{S}$ whose projection is σ^δ, a necessary condition for it to constitute the desired ε-proper equilibrium of $\bar{\Gamma}$ is that it satisfy the following: $(\textit{Up}, \textit{Left})$ is the most likely profile, followed by $(\textit{Down}, \textit{Left})$, $(\textit{Up}, \textit{Center})$, $(\textit{Down}, \textit{Center})$, $(\textit{Up}, \textit{Right})$, and $(\textit{Down}, \textit{Right})$. As shown above, ϕ^δ itself does not satisfy this necessary condition when δ is small. However, ϕ^δ is not the only distribution over $\bar{S}$ whose projection is σ^δ. The rest of the discussion demonstrates how we construct another such distribution with the necessary properties. To illustrate our approach, we partition the elements of $\bar{S}$ in the following way:

- Let $s^* = (\textit{Up}, \textit{Left})$. This is the profile that is played in equilibrium.
- Let $\mathcal{S}^1 = \{(\textit{Up}, \textit{Center}), (\textit{Up}, \textit{Right}), (\textit{Down}, \textit{Left})\}$. These are all single deviations from s^*.
- Let $\mathcal{S}^2 = \{(\textit{Down}, \textit{Center}), (\textit{Down}, \textit{Right})\}$. These are all joint deviations from s^*.

The key is to adjust the probabilities assigned to the elements of $\mathcal{S}^2$ in an appropriate way, then also adjusting other probabilities so as to keep the projection unchanged. Let $\bar{\sigma}^{\delta,\varepsilon}$ denote the new distribution over $\bar{S}$ that we will obtain. Applying the proof's general construction to this example, $\bar{\sigma}^{\delta,\varepsilon}$ attaches the following probabilities to the elements of $\mathcal{S}^2$:

$$\bar{\sigma}^{\delta,\varepsilon}(\textit{Down}, \textit{Center}) = \varepsilon^2(\delta^3 - \delta^5) \qquad\qquad \bar{\sigma}^{\delta,\varepsilon}(\textit{Down}, \textit{Right}) = \varepsilon^3(\delta^4 - \delta^6)$$

Relative to ϕ^δ, notice that we have added the weight $\varepsilon^2(\delta^3 - \delta^5) - \delta^5$ to $(\textit{Down}, \textit{Center})$. So as to ensure that the projection will remain unchanged, we therefore subtract the same amount of weight from each of the corresponding single deviations in $\mathcal{S}^1$—in this case, $(\textit{Up}, \textit{Center})$ and $(\textit{Down}, \textit{Left})$—and we also add the same amount of weight to s^*. Similarly, we have added the weight $\varepsilon^3(\delta^4 - \delta^6) - \delta^6$ to $(\textit{Down}, \textit{Right})$. We therefore subtract the same amount from both $(\textit{Up}, \textit{Right})$ and $(\textit{Down}, \textit{Left})$, and we also add the same amount to s^*. In summary, we obtain

$$\begin{aligned}
\bar{\sigma}^{\delta,\varepsilon}(\textit{Up}, \textit{Left}) &= 1 - \delta^2 - \delta^3 - \delta^4 + \varepsilon^2(\delta^3 - \delta^5) + \varepsilon^3(\delta^4 - \delta^6) \\
\bar{\sigma}^{\delta,\varepsilon}(\textit{Down}, \textit{Left}) &= \delta^2 - \varepsilon^2(\delta^3 - \delta^5) - \varepsilon^3(\delta^4 - \delta^6) \\
\bar{\sigma}^{\delta,\varepsilon}(\textit{Up}, \textit{Center}) &= \delta^3 - \varepsilon^2(\delta^3 - \delta^5) \\
\bar{\sigma}^{\delta,\varepsilon}(\textit{Up}, \textit{Right}) &= \delta^4 - \varepsilon^3(\delta^4 - \delta^6)
\end{aligned}$$

It is easily verified that σ^δ is the projection of $\bar{\sigma}^{\delta,\varepsilon}$—just as it was the projection of ϕ^δ. Furthermore, if we choose δ to be a sufficiently high power of ε, then this constitutes the desired ε-proper equilibrium of $\bar{\Gamma}$. Indeed, plugging in $\delta = \varepsilon^4$, we obtain the following distribution over $\bar{S}$:[44]

$$\begin{aligned}
\bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Left}) &= 1 - \varepsilon^8 - \varepsilon^{12} + \varepsilon^{14} - \varepsilon^{16} + \varepsilon^{19} - \varepsilon^{22} - \varepsilon^{27} \\
\bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Left}) &= \varepsilon^8 - \varepsilon^{14} - \varepsilon^{19} + \varepsilon^{22} + \varepsilon^{27} \\
\bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Center}) &= \varepsilon^{12} - \varepsilon^{14} + \varepsilon^{22} \\
\bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Center}) &= \varepsilon^{14} - \varepsilon^{22} \\
\bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Right}) &= \varepsilon^{16} - \varepsilon^{19} + \varepsilon^{27} \\
\bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Right}) &= \varepsilon^{19} - \varepsilon^{27}
\end{aligned}$$

We see that when $\varepsilon > 0$ is sufficiently small, $\bar{\sigma}^\varepsilon$ satisfies the aforementioned necessary condition—and what is more, we have the following:

$$\begin{aligned}
&\bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Right}) \leq \varepsilon \cdot \bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Right}) \leq \varepsilon \cdot \bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Center}) \\
&\quad \leq \varepsilon \cdot \bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Center}) \leq \varepsilon \cdot \bar{\sigma}^\varepsilon(\mathit{Down}, \mathit{Left}) \leq \varepsilon \cdot \bar{\sigma}^\varepsilon(\mathit{Up}, \mathit{Left}),
\end{aligned}$$

from which it follows that $(\bar{\sigma}^\varepsilon, \ldots, \bar{\sigma}^\varepsilon)$ is an ε-proper equilibrium of $\bar{\Gamma}$. We conclude that $\bar{\sigma} = \lim_{\varepsilon\to 0} \bar{\sigma}^\varepsilon = (\mathit{Up}, \mathit{Left})$ is a symmetrically proper strategy of $\bar{\Gamma}$. Moreover, $\sigma = (\mathit{Up}, \mathit{Left})$, which is the extended proper equilibrium of Γ that we began with, is trivially the projection of $\bar{\sigma}$, as required.

[44] The general proof uses $\delta = \varepsilon^{|\bar{S}|+2}$, which always suffices. But for this example, even $\delta = \varepsilon^4$ suffices.