

Extended Proper Equilibrium*

Paul Milgrom Joshua Mollner

October 15, 2019

Abstract

We introduce *extended proper equilibrium*, which refines proper equilibrium (?) 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 to make a less costly one. Proper equilibrium, however, still permits the possibility that one player has an arbitrarily greater propensity to tremble than another, so 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, the strategy *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

*For helpful comments, we thank Gabriel Carroll, Piotr Dworzak, Fuhito Kojima, 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.

as likely that Row will deviate to play *Down* as that Column will deviate to play *Right*, 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[†]

[<i>West</i>]	<i>Left</i>	<i>Right</i>		[<i>East</i>]	<i>Left</i>	<i>Right</i>
<i>Up</i>	1, 1, 0	1, 1, 1		<i>Up</i>	1, 1, 0	1, 1, 0
<i>Down</i>	0, 1, 0	1, 0, 0		<i>Down</i>	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 this condition is required to hold only for some scaling, the set of extended proper equilibria is invariant to 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 first deviation has zero cost and the second has positive cost.

The idea that players should not differ arbitrarily in their propensities to err 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. Since 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 $(\bar{\sigma}^\varepsilon, \dots, \bar{\sigma}^\varepsilon)$ 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

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

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

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

²A 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

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 players being arbitrarily 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 two 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 ?), which can be interpreted as capturing beliefs about how the game will be played. This characterization functions as a useful step 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 most closely related to *proper equilibrium* (?): our definition not only refines proper equilibrium but also connects to it through the meta-game model that we describe. Despite its various critiques, proper equilibrium remains a prominent refinement and serves as a natural starting point.

A handful of other refinements of proper equilibrium have also been proposed elsewhere in the literature. *Finely settled equilibrium* (?) 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* (?) 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* (?) refines strictly perfect equilibrium (?) by adding a certain continuity assumption; it also refines proper equilibrium. *Truly proper*

induced strategy profile for the original game is the profile of the component-wise marginal distributions of $\bar{\sigma}^\varepsilon$ (cf. Definition 5).

equilibrium and *more-than-proper equilibrium* (?) also refine proper equilibrium, yet are not as demanding as strictly proper equilibrium. *Strongly proper equilibrium* (?) refines proper equilibrium by adding the requirement that if two strategies are payoff-equivalent for a player, then they must be trembled to with the same probability. *Persistent equilibrium* (?) 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* (?) contains a proper equilibrium.

Another related refinement is *test-set equilibrium* (?), 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 may fail to exist for some finite games, whereas every such game contains at least one extended proper equilibrium. Another difference is that test-set equilibrium does not require players to tremble independently. In contrast, perfect equilibrium—and hence extended proper equilibrium—does require such independence, which can lead to additional restrictions.³

Finally, a number of other solution concepts can be thought of as imposing restrictions on the likelihood of deviations akin to those across-player restrictions 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* (?), *divine equilibrium* (?), and the *compatibility criterion* (?). Closely related to the last of these is *player-compatible equilibrium* (?), 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* (?) 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 $\mathcal{N} = \{1, \dots, N\}$ 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 also use \bar{S} to denote $\prod_{n \in \mathcal{N}} S_n$, the set of pure strategy profiles.

³Appendix C contains a formal definition of test-set equilibrium and a further discussion of its relationship with extended proper equilibrium.

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 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 $\sigma = (\sigma_1, \dots, \sigma_N)$.

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 :

$$L_n(\sigma) := \max_{\hat{s}_n \in S_n} \pi_n(\sigma/\hat{s}_n) - \pi_n(\sigma).$$

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 $\sigma \in \prod_{n \in \mathcal{N}} \Delta_n^0$ 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 (?). 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.⁴

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 $(\varepsilon_t)_{t=1}^\infty$ and $(\sigma^t)_{t=1}^\infty$ such that:

- (i) each $\varepsilon_t > 0$ and $\lim_{t \rightarrow \infty} \varepsilon_t = 0$,
- (ii) each σ^t is an (α, ε_t) -extended proper equilibrium, and
- (iii) $\lim_{t \rightarrow \infty} \sigma^t = \sigma$.

The scaling vector α can play two roles. In traditional non-cooperative game theory, payoffs are expressed in terms of von Neumann-Morgenstern utilities, which are unique only up to affine transformations. With that interpretation, game-theoretic predictions should

⁴For purposes of comparison, we reiterate the following definitions from ?. An ε -proper equilibrium is a profile of totally mixed strategies $\sigma \in \prod_{n \in \mathcal{N}} \Delta_n^0$ 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 $\sigma \in \prod_{n \in \mathcal{N}} \Delta_n^0$ 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 $\varepsilon \rightarrow 0$, in the same sense as Definition 2.

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.⁵

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.

3.3 Three Basic Facts About Extended Proper Equilibrium

The first three theorems 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., ?). 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., ?). \square

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 ?.

⁵For those settings, an alternative approach would be to treat all players as equally likely to tremble. This would correspond to requiring $\alpha = 1$ in the definition of extended proper equilibrium. This would also result in definition that is more restrictive than the one given here, although not so restrictive that an equilibrium satisfying this stronger condition might fail to exist in finite games. In fact, the proof of Theorem 3 proceeds by establishing existence of this stronger concept.

⁶To see this, suppose that σ is an extended proper equilibrium, so that it is the limit of a sequence $(\sigma^t)_{t=1}^\infty$ 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 $\sigma_l^t(s'_l) \leq \varepsilon_t \sigma_m^t(s''_m)$ —for sufficiently large values of t .

Corollary 2. *An extended proper equilibrium of a normal form game induces a quasi-perfect equilibrium (?)—and hence a sequential equilibrium (?)—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 3. *Every finite normal form game has at least one extended proper equilibrium.*

The proof of Theorem 3 is deferred to Appendix B, as are the proofs of most subsequent results. The proof is similar to the one used by ? 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 4. *In two-player games, the sets of proper equilibria and extended proper equilibria coincide.*

The fact that proper equilibrium and extended proper equilibrium coincide in two-player games is perhaps reassuring. 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. With two players, each player has only a single opponent, and there are none of these across-opponent comparisons to be made. Thus, while extended proper equilibrium refines proper equilibrium by imposing additional structure on these across-opponent comparisons, the result indicates that extended proper equilibrium does not impose much more.

Theorem 4 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. 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. Lexicographic probability systems, discussed in Section 5, are useful in this construction: for a proper equilibrium σ , we use the machinery to define a sequence of subsets of $\prod_{n \in \mathcal{N}} \Delta_n^0$ converging to σ within which to use a fixed point theorem to find the requisite (α, ε) -extended proper equilibria.

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 standard tremble-based refinements do not incorporate such symmetry

in their definitions. Rather, both perfect and proper equilibrium permit one player to have an arbitrarily greater propensity to tremble than another.

Although there might be many reasons for players to be similar in their degrees of strategic sophistication, 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 $\bar{\Gamma}$ itself 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 .⁷ 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

$$\bar{\pi}_n(\bar{s}_1, \dots, \bar{s}_N) = \frac{1}{N!} \sum_{\tau \in \text{Sym}(\mathcal{N})} \pi_{\tau(n)}\left(\text{proj}_1(\bar{s}_{\tau^{-1}(1)}), \dots, \text{proj}_N(\bar{s}_{\tau^{-1}(N)})\right).⁸$$

In the biological game theory literature, this meta-game construction is used in order to apply the concept of an evolutionarily stable strategy—which was originally defined only for symmetric games—to asymmetric games (e.g., ???). Similarly, we exploit the symmetric structure of $\bar{\Gamma}$ to apply symmetric versions of the standard tremble-based refinements to

⁷Thus, 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.

⁸In this sum, $\text{Sym}(\mathcal{N})$ denotes the set of all permutations of \mathcal{N} . Under a permutation τ , the role of player m in 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 $\text{proj}_m(\bar{s}_{\tau^{-1}(m)})$. By averaging the payoffs to player n in $\bar{\Gamma}$ over all permutations, we obtain the given form for $\bar{\pi}_n$.

that game. Note that the following definitions require symmetry not only of the equilibrium strategies but also of the trembles (as in e.g., ?).

Definition 4. In a finite symmetric game Γ ,

- (i) a strategy σ is a *symmetrically Nash strategy* if (σ, \dots, σ) 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 totally mixed 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 totally mixed 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, analogous results are true for the 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 5. *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 $\bar{\Gamma}'$.*

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 to Γ .⁹

5 Lexicographic Characterization

In this section, we employ the framework of ? to characterize extended proper equilibrium directly in terms of players' beliefs. Toward that end, let $\rho = (p^1, \dots, p^K)$ be a sequence of probability measures on \bar{S} , the set of pure strategy profiles. ? refer to such a sequence as a *lexicographic probability system* (LPS). 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.¹⁰

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.¹¹

$$BR_n(\rho) = \left\{ s'_n \in S_n \mid \forall s''_n \in S_n : \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K \geq_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K \right\}.$$

As additional notation, given any $J \subseteq \mathcal{N}$, let p_J^k be the marginal distribution of p^k on $\prod_{j \in J} S_j$. An LPS ρ gives rise to a partial order on the elements of the set $\bigcup_{J \subseteq \mathcal{N}} \prod_{j \in J} S_j$, which is the set of pure strategy profiles for subsets of \mathcal{N} . If $s_J \in \prod_{j \in J} S_j$ and $s_I \in \prod_{i \in I} S_i$, 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 $\min\{k : p_J^k(s_J) > 0\} < \min\{k : p_I^k(s_I) > 0\}$. We write $s_J \geq_\rho s_I$ to mean it is not the case that $s_I >_\rho s_J$.

Next, we define what ? refer to as a lexicographic Nash equilibrium. Such an object 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 $s_n \in S_n$, $p_n^1(s_n) > 0$ implies $s_n \in BR_n(\rho)$, and
- (ii) for all $s \in \bar{S}$, $p^1(s) = \prod_{n \in \mathcal{N}} \sigma_n(s_n)$.

⁹In the proof of Theorem 5, 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.

¹⁰Our approach and notation differ slightly from ?. One difference is that we directly assume the beliefs that players possess about the strategies of their opponents are the marginals of a “common prior LPS,” whereas ? only add that assumption when they extend 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 impose this common prior assumption from the outset. There are also small notational differences. For instance, while they designate players with superscripts and levels of an LPS with subscripts, we do the reverse.

¹¹Formally, 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 $a^j > b^j$.

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

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

Definition 7. An LPS $\rho = (p^1, \dots, p^K)$ on \bar{S} respects *within-person preferences* if for all $n \in \mathcal{N}$ and all $s'_n, s''_n \in S_n$ with $s'_n \geq_\rho s''_n$, it is the case that

$$\left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K \geq_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K.$$

Definition 8. An LPS $\rho = (p^1, \dots, p^K)$ on \bar{S} respects *within-and-across-person 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 $s'_l \in S_l$, $s''_m \in S_m$ with $s'_l \geq_\rho s''_m$, it is the case that

$$\left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] \right]_{k=1}^K \leq_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s''_m) - \pi_m(s/s_m^*)] \right]_{k=1}^K.$$

Definition 7 requires that under the beliefs, player n is infinitely less likely to use a strategy that is a “worse response” to its beliefs about its opponents 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-person 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 ?. These are generalizations for LPSs of what it means for a probability measure to have full support or to be a product measure.

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

Definition 10. An LPS $\rho = (p^1, \dots, p^K)$ on \bar{S} satisfies *strong independence* if there is an equivalent \mathbb{F} -valued probability measure that is a product measure, where \mathbb{F} is some non-Archimedean ordered field that is a proper extension of \mathbb{R} .¹²

¹²An ordered field is non-Archimedean if it contains an element $\varepsilon > 0$ such that $\varepsilon < \frac{1}{n}$ for all $n \in \mathbb{N}$. Such an element is called an infinitesimal. An LPS $\rho = (p^1, \dots, p^K)$ on \bar{S} and an \mathbb{F} -valued probability measure \hat{p} on \bar{S} , are said to be *equivalent* if there exists a vector of positive infinitesimals $\varepsilon = (\varepsilon_1, \dots, \varepsilon_{K-1}) \in \mathbb{F}^{K-1}$

Using these definitions, Proposition 9 provides a characterization of extended proper equilibrium. For comparison, we also incorporate as Propositions 6–8 below the characterizations of Nash equilibrium, perfect equilibrium, and proper equilibrium, which are due to ?.¹³ These characterizations are useful for two reasons. First, they simplify some proofs (particularly that of Theorem 4), 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 6. *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 7. *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 8. *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-person preferences for which (ρ, σ) is a lexicographic Nash equilibrium.*

Proposition 9. *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-person preferences for which (ρ, σ) is a lexicographic Nash equilibrium.*

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 ? who also proposed *locally envy-free equilibrium* as a refinement of the auction’s many equilibria.¹⁴ That refinement, however, is defined directly in terms of the GSP auction game and does not apply

such that $\hat{p}(s) = \varepsilon \square \rho(s)$ for all $s \in \bar{S}$, where $\varepsilon \square \rho$ denotes the probability measure

$$\begin{aligned} \varepsilon \square \rho = & (1 - \varepsilon_1)p^1 + \varepsilon_1 \left[(1 - \varepsilon_2)p^2 + \varepsilon_2 \left[(1 - \varepsilon_3)p^3 \right. \right. \\ & \left. \left. + \varepsilon_3 \left[\cdots + \varepsilon_{K-2} \left[(1 - \varepsilon_{K-1})p^{K-1} + \varepsilon_{K-1}p^K \right] \cdots \right] \right] \right]. \end{aligned}$$

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

¹³Propositions 7 and 8 are restatements of Propositions 7 and 8 of ?. While Proposition 6 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 in Appendix B. The differences are: (i) their proposition does not require the beliefs that players possess about the strategies of their opponents to be the marginals of a “common prior LPS,” and (ii) their proposition does not require the LPS to satisfy strong independence. The characterization is true regardless of whether these requirements are imposed. As discussed in footnote 10, it is natural, given our approach, for us to make the common prior assumption at the outset. In addition, we also impose strong independence to provide more continuity with the other characterizations.

¹⁴? studies the same model and makes the same equilibrium selection, calling these *symmetric equilibria*. For formal definitions of the GSP auction model and the locally envy-free equilibrium refinement, see the aforementioned papers or see ?, Section 3.2.

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.

Bids in a GSP auction are submitted simultaneously. After breaking any ties uniformly at random, the i th ad position is allocated to the i th highest bidder, who makes a per-click payment that is equal to the next highest bid. (In the case of the lowest bidder, the payment is zero.)

In the original analysis, allowable bids are the nonnegative reals. However, the refinements of proper equilibrium and extended proper equilibrium are usually defined only for finite games. Consequently, the following discussion restricts bids to the finite set $\{0, \frac{1}{m}, \frac{2}{m}, \dots, \frac{mM}{m}\}$ for positive integers m and M .

Furthermore, for most of the discussion below, we focus on the following specific example. There are three bidders, whose values per click are 4, 2, and 1, respectively. There are also three ad positions, which attract four clicks, two clicks, and one click, respectively. For the discretized bid set, let $m = 6$ and $M = 3$, although the arguments can easily be extended to accommodate all even $m \geq 6$ and all $M \geq 3$.

The equilibrium $b^* = (\frac{5}{2}, 1, \frac{1}{2})$ is not locally envy-free.¹⁵ Nevertheless, it is a proper equilibrium. Indeed, in ?, 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 that it is 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 9, suppose to the contrary that ρ is an LPS that satisfies strong independence, has full support, and respects within-and-across-person 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:* $b_1 = \frac{7}{6}$ and $b_3 \neq \frac{1}{2}$
- *Case 3(a):* $b_1 = 1$ and b_3 unrestricted
- *Case 3(b):* $b_1 \notin \{1, \frac{7}{6}\}$ and $b_3 \in \{1, \frac{7}{6}\}$

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

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

¹⁵This is because bidder 2 “envies” bidder 1 in the sense described by ?: 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, then it would obtain the higher payoff $\pi_2(1, \frac{5}{2}, \frac{1}{2}) = 4$.

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 3(a) or 3(b), at least one bidder is playing an inferior response to b^* .¹⁶ 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 respecting within-and-across-person preferences that for any (b_1, b_3) in case 3(a) or 3(b), $(\frac{7}{6}, \frac{1}{2}) >_\rho (b_1, b_3)$. Similarly, any profile in case 2 features bidder 3 deviating from its equilibrium bid of $b_3^* = \frac{1}{2}$. It therefore follows from ρ satisfying strong independence that for any $(\frac{7}{6}, b_3)$ in case 2, $(\frac{7}{6}, \frac{1}{2}) >_\rho (\frac{7}{6}, b_3)$. We conclude from this analysis that $\frac{7}{6}$ is a profitable deviation against ρ for bidder 2, which contradicts (ρ, b^*) being a lexicographic Nash equilibrium.

In fact, the conclusion of our example can be shown to apply more generally: in all discretized GSP auction games with m and M sufficiently large, every pure extended proper equilibrium is locally envy-free.¹⁷ And, again, that is not a conclusion that could be obtained from proper equilibrium alone.

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 of strategic sophistication—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, each player should regard costly mistakes 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 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

¹⁶If bidder 1 deviates from $b_1^* = \frac{5}{2}$ to $b_1 = 1$ while the others continue to play their equilibrium bids of $b_2^* = 1$ and $b_3^* = \frac{1}{2}$, then that bidder'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 $b_3 \in \{1, \frac{7}{6}\}$.

¹⁷In ?, Section 3.2, we 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, as noted in Section 2, extended proper equilibrium always exist in finite games while test-set equilibrium may fail to exist. Moreover, 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).

equilibrium may be a useful refinement of proper equilibrium. The generalized second price auction model (??) is one such application.

A Lemmas

A.1 Fixed Point Theorem

The proof of Theorem 4 requires a more powerful fixed point theorem than that of ?. We use the following corollary of the ? fixed point result, the statement of which is taken from ?.

Lemma 10 (Eilenberg-Montgomery). *Suppose that a compact metric space (X, d) is an absolute retract 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. For the purposes of this appendix, it is useful to recall the definition of the “square operator” introduced in footnote 12. If $r = (r^1, \dots, r^{K-1})$ is a vector, and $\rho = (p^1, \dots, p^K)$ is an LPS, then

$$r \square \rho = (1 - r^1)p^1 + r^1 \left[(1 - r^2)p^2 + r^2 \left[(1 - r^3)p^3 + r^3 \left[\dots + r^{K-2} \left[(1 - r^{K-1})p^{K-1} + r^{K-1}p^K \right] \dots \right] \right] \right].$$

Lemma 11 is an immediate consequence of this definition.

Lemma 11. *Suppose $\rho = (p^1, \dots, p^K)$ is an LPS on \bar{S} and $r \in (0, 1)^{K-1}$. For any $s'_n \in S_n$, define $k' = \min\{k : p_n^k(s'_n) > 0\}$ and $m' = p_n^{k'}(s'_n)$. Letting σ_n be the marginal of $r \square \rho$ on S_n ,*

$$r^1 \dots r^{k'-1} (1 - r^{k'}) m' \leq \sigma_n(s'_n) \leq r^1 \dots r^{k'-1}.$$

Proof of Lemma 11. First observe that $\sigma_n(s'_n) = \sum_{s \in \bar{S}: s_n = s'_n} r \square \rho(s)$, which can be expressed as

$$\begin{aligned} & (1 - r^1) \sum_{s \in \bar{S}: s_n = s'_n} p^1(s) + r^1 \left[(1 - r^2) \sum_{s \in \bar{S}: s_n = s'_n} p^2(s) + r^2 \left[(1 - r^3) \sum_{s \in \bar{S}: s_n = s'_n} p^3(s) \right. \right. \\ & \quad \left. \left. + r^3 \left[\dots + r^{K-2} \left[(1 - r^{K-1}) \sum_{s \in \bar{S}: s_n = s'_n} p^{K-1}(s) + r^{K-1} \sum_{s \in \bar{S}: s_n = s'_n} p^K(s) \right] \dots \right] \right] \right] \\ & = (1 - r^1) p_n^1(s'_n) + r^1 \left[(1 - r^2) p_n^2(s'_n) + r^2 \left[(1 - r^3) p_n^3(s'_n) \right. \right. \\ & \quad \left. \left. + r^3 \left[\dots + r^{K-2} \left[(1 - r^{K-1}) p_n^{K-1}(s'_n) + r^{K-1} p_n^K(s'_n) \right] \dots \right] \right] \right]. \end{aligned}$$

Using the fact that $p_n^k(s'_n) = 0$ for all $k < k'$, this simplifies to

$$r^1 \dots r^{k'-1} \left[(1 - r^{k'}) p_n^{k'}(s'_n) + r^{k'} \left[\dots + r^{K-2} \left[(1 - r^{K-1}) p_n^{K-1}(s'_n) + r^{K-1} p_n^K(s'_n) \right] \dots \right] \right].$$

On one hand, we have for all $k \geq k'$ that $p_n^k(s'_n) \leq 1$. Plugging in this bound, we obtain $\sigma_n(s'_n) \leq r^1 \dots r^{k'-1}$. On the other hand, we have $p_n^{k'}(s'_n) = m'$ and for all $k > k'$ that $p_n^k(s'_n) \geq 0$. Plugging in these bounds, we obtain $\sigma_n(s'_n) \geq r^1 \dots r^{k'-1} (1 - r^{k'}) m'$. \square

Next, the following two lemmas will be useful in proving Proposition 9, which is our main result concerning lexicographic probability systems. Lemma 12 is a version of Proposition 1 of ?. Lemma 13 uses that as an input to produce a related result, which expands focus from within-person comparisons to within-and-across-person comparisons.

Lemma 12. *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 $(\forall n \in \mathcal{N})(\forall s'_n \in S_n)(\forall s''_n \in S_n)$*

$$\sum_{s \in \bar{S}} (r \square \rho)(s) \pi_n(s/s'_n) < \sum_{s \in \bar{S}} (r \square \rho)(s) \pi_n(s/s''_n) \quad (1)$$

if and only if

$$\left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K <_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K. \quad (2)$$

Proof of Lemma 12. Fix an LPS $\rho = (p^1, \dots, p^K)$. To economize on notation, define for each $n \in \mathcal{N}$ the following function

$$H_n(k, s'_n, s''_n) = \sum_{s \in \bar{S}} p^k(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)].$$

If $H_n(k, s', s'') = 0$ for all choices of (n, k, s'_n, s''_n) , then neither (1) nor (2) ever applies, and the statement holds vacuously with any choice of δ . We therefore assume henceforth that this is not the case. Define the constants

$$W = \max_{\substack{n \in \mathcal{N} \\ s'_n, s''_n \in S_n}} |H_n(k, s'_n, s''_n)|$$

$$B = \min_{\substack{n \in \mathcal{N} \\ k \in \{1, \dots, K\} \\ s'_n, s''_n \in S_n}} \{H_n(k, s'_n, s''_n) : H_n(k, s'_n, s''_n) > 0\}.$$

As a result of the previous assumption, B is a well-defined positive number. We also obviously have $W \geq 0$. We then choose $\delta = \frac{B}{B+W} > 0$, and now show that the statement holds with this choice of δ . Suppose $r \in (0, \delta)^{K-1}$.

First, suppose further that for some $n \in \mathcal{N}$, $s'_n \in S_n$, and $s''_n \in S_n$,

$$\left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K >_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K.$$

Then there is some k such that both $H_n(k, s'_n, s''_n) \geq B > 0$ and for all $j < k$, $H_n(j, s'_n, s''_n) =$

0. For this k , let $\hat{\rho} = (p^{k+1}, \dots, p^K)$ and $\hat{r} = (r^{k+1}, \dots, r^{K-1})$. Then we have that

$$\begin{aligned}
& \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)] \\
&= r^1 \dots r^{k-1} \left[(1 - r^k) H_n(k, s'_n, s''_n) + r^k \sum_{s \in \bar{S}} (\hat{r} \square \hat{\rho})(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)] \right] \\
&\geq r^1 \dots r^{k-1} [(1 - r^k) B - r^k W] \\
&> r^1 \dots r^{k-1} [B - \delta(B + W)] \\
&= 0,
\end{aligned}$$

as desired, where we obtain the last step because $\delta = \frac{B}{B+W}$. Second, suppose instead that for some $n \in \mathcal{N}$, $s'_n \in S_n$, and $s''_n \in S_n$,

$$\left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K =_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K.$$

Then $H_n(k, s'_n, s''_n) = 0$ for all k . And because this is the case we also have

$$\sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)] = 0,$$

as desired, because it is a weighted average of $\{H_n(k, s'_n, s''_n)\}_{k=1}^K$. Combining the first and second arguments, we have established the contrapositive of (1) \implies (2). And by symmetry, the first argument serves to establish (2) \implies (1). \square

Lemma 13. *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 $(\forall l \in \mathcal{N})(\forall m \in \mathcal{N})(\forall s'_l \in S_l)(\forall s''_m \in S_m)(\forall s_l^* \in BR_l(\rho))(\forall s_m^* \in BR_m(\rho))$*

$$\begin{aligned}
& \alpha_l \max_{\hat{s}_l \in S_l} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] \\
&> \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)]
\end{aligned} \tag{3}$$

if and only if

$$\begin{aligned}
& \left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s_l^*) - \pi_l(s/s'_l)] \right]_{k=1}^K \\
&>_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s_m^*) - \pi_m(s/s''_m)] \right]_{k=1}^K.
\end{aligned} \tag{4}$$

Proof of Lemma 13. Fix an LPS $\rho = (p^1, \dots, p^K)$ and a scaling vector $\alpha \in \mathbb{R}_{++}^N$. Select for every n some $s_n^* \in BR_n(\rho)$. Because all elements of $BR_n(\rho)$ must produce identical expected payoffs against every p^k , it will not matter which element is chosen. To economize on notation, define for each $l, m \in \mathcal{N}$, possibly equal, the following function

$$H_{lm}(k, s'_l, s''_m) = \sum_{s \in \bar{S}} p^k(s) (\alpha_l [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] - \alpha_m [\pi_m(s/s''_m) - \pi_m(s/s_m^*)]).$$

If $H_{lm}(k, s'_l, s''_m) = 0$ for all choices of (l, m, k, s'_l, s''_m) , then neither (3) nor (4) ever applies, and the statement holds vacuously with any choice of δ . We therefore assume henceforth that this is not the case. Define the constants

$$W = \max_{\substack{l, m \in \mathcal{N} \\ k \in \{1, \dots, K\} \\ s'_l \in S_l, s''_m \in S_m}} |H_{lm}(k, s'_l, s''_m)|$$

$$B = \min_{\substack{l, m \in \mathcal{N} \\ k \in \{1, \dots, K\} \\ s'_l \in S_l, s''_m \in S_m}} \{H_{lm}(k, s'_l, s''_m) : H_{lm}(k, s'_l, s''_m) > 0\}.$$

As a result of the previous assumption, B is a well-defined positive number. We also obviously have $W \geq 0$. In addition, let $\hat{\delta} > 0$ be such that the conclusion of Lemma 12 applies. We then choose $\delta = \min \left\{ \hat{\delta}, \frac{B}{B+W} \right\} > 0$, and now show that the statement holds with this choice of δ . Suppose $r \in (0, \delta)^{K-1}$.

First, suppose further that for some $l, m \in \mathcal{N}$, $s'_l \in S_l$, and $s''_m \in S_m$,

$$\left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] \right]_{k=1}^K >_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s''_m) - \pi_m(s/s_m^*)] \right]_{k=1}^K.$$

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

$$\begin{aligned} & \alpha_l \max_{\hat{s}_l \in S_l} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_l(s/\hat{s}_l) - \pi_l(s/s'_l)] - \alpha_m \max_{\hat{s}_m \in S_m} \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_m(s/\hat{s}_m) - \pi_m(s/s''_m)] \\ &= \alpha_l \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] - \alpha_m \sum_{s \in \bar{S}} (r \square \rho)(s) [\pi_m(s/s''_m) - \pi_m(s/s_m^*)] \\ &= \sum_{s \in \bar{S}} (r \square \rho)(s) (\alpha_l [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] - \alpha_m [\pi_m(s/s''_m) - \pi_m(s/s_m^*)]) \\ &\geq r^1 \dots r^{k-1} \left[(1 - r^k) H_{lm}(k, s'_l, s''_m) \right. \\ &\quad \left. + r^k \sum_{s \in \bar{S}} (\hat{r} \square \hat{\rho})(s) (\alpha_l [\pi_l(s/s'_l) - \pi_l(s/s_l^*)] - \alpha_m [\pi_m(s/s''_m) - \pi_m(s/s_m^*)]) \right] \end{aligned}$$

$$\begin{aligned}
&\geq r^1 \dots r^{k-1} \left[(1 - r^k)B - r^k W \right] \\
&> r^1 \dots r^{k-1} [B - \delta(B + W)] \\
&\geq r^1 \dots r^{k-1} \left[B - \frac{B}{B + W}(B + W) \right] \\
&= 0,
\end{aligned}$$

as desired. The first step in the above follows from Lemma 12, which implies that any maximizers \hat{s}_l and \hat{s}_m must be elements of $BR_l(\rho)$ and $BR_m(\rho)$, respectively. Second, suppose instead that for some $l, m \in \mathcal{N}$, $s'_l \in S_l$, and $s''_m \in S_m$,

$$\left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s'_l) - \pi_l(s/\hat{s}_l)] \right]_{k=1}^K =_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s''_m) - \pi_m(s/\hat{s}_m)] \right]_{k=1}^K.$$

Then $H_{lm}(k, s'_l, s''_m) = 0$ for all k . And because this is the case we also have

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

which is a weighted average of $\{H_{lm}(k, s'_l, s''_m)\}_{k=1}^K$, and is therefore equal to zero, as desired. As before, the first step in the above follows from Lemma 12. Combining the first and second arguments, we have established the contrapositive of (3) \implies (4). And by symmetry, the first argument serves to establish (4) \implies (3). \square

B Omitted Proofs

B.1 Proofs Corresponding to Section 3

Proof of Theorem 3. Fix any $\alpha \in \mathbb{R}_{++}^N$ and any $\varepsilon \in (0, 1)$. Let $M = \sum_{n \in \mathcal{N}} |S_n|$ and $\delta = \frac{\varepsilon^{2M}}{M}$. 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_n^0 . We also define the correspondence $F : \prod_{n \in \mathcal{N}} \Delta_n^* \rightrightarrows \prod_{n \in \mathcal{N}} \Delta_n^*$ as

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

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 $n \in \mathcal{N}$, all $m \in \mathcal{N}$, and all $s_n \in S_n$, define $\phi_m(s_n)$ as the number of pure strategies $\hat{s}_m \in S_m$ satisfying $\alpha_m L_m(\sigma/\hat{s}_m) < \alpha_n L_n(\sigma/s_n)$. Let $\phi(s_n) = \sum_{m \in \mathcal{N}} \phi_m(s_n)$, and let $\Phi_n^0 = |\{s_n \in S_n : \phi(s_n) = 0\}|$. Finally, let

$$\sigma_n^*(s_n) = \begin{cases} \frac{\varepsilon^{\phi(s_n)}}{2M} & \text{if } \phi(s_n) > 0 \\ \frac{1}{\Phi_n^0} \left(1 - \sum_{\hat{s}_n: \phi(\hat{s}_n) > 0} \frac{\varepsilon^{\phi(\hat{s}_n)}}{2M} \right) & \text{if } \phi(s_n) = 0 \end{cases}$$

We next verify that $\sigma^* \in \prod_{n \in \mathcal{N}} \Delta_n^*$. To see this, let $n \in \mathcal{N}$ and $s_n \in S_n$. We consider two cases. First, if $\phi(s_n) > 0$, then

$$\sigma_n^*(s_n) = \frac{\varepsilon^{\phi(s_n)}}{2M} \geq \frac{\varepsilon^M}{2M} = \delta.$$

Second, if $\phi(s_n) = 0$, then

$$\sigma_n^*(s_n) = \frac{1}{\Phi_n^0} \left(1 - \sum_{s_n: \phi(s_n) > 0} \frac{\varepsilon^{\phi(s_n)}}{2M} \right) \geq \frac{1}{\Phi_n^0} \left(1 - M \frac{1}{2M} \right) = \frac{1}{2\Phi_n^0} \geq \frac{1}{2M} \geq \delta.$$

Next, we verify that $\sigma^* \in F(\sigma)$. To that end, 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 $\phi(s'_l) \geq \phi(s''_m) + 1$. Indeed, we must have $\phi_m(s'_l) \geq \phi_m(s''_m) + 1$ and $\phi_n(s'_l) \geq \phi_n(s''_m)$ for all $n \neq m$. Now we again consider two cases. First, if $\phi(s''_m) > 0$, then

$$\sigma_l^*(s'_l) = \frac{\varepsilon^{\phi(s'_l)}}{2M} \leq \frac{\varepsilon^{\phi(s''_m)+1}}{2M} = \varepsilon \cdot \sigma_m^*(s''_m).$$

Second, if $\phi(s''_m) = 0$, then

$$\sigma_l^*(s'_l) = \frac{\varepsilon^{\phi(s'_l)}}{2M} \leq \frac{\varepsilon}{2M} \leq \varepsilon \cdot \sigma_m^*(s''_m),$$

where the last step uses the fact, established above, that $\phi(s''_m) = 0$ implies $\sigma_m^*(s''_m) \geq \frac{1}{2M}$.

Claim: F is upper-hemicontinuous.

Proof of Claim: This follows from continuity of L_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 (?), 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 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 must exist a convergent subsequence $(\sigma^{t_s})_{s=0}^\infty$. The limit of this subsequence is an extended proper equilibrium. \square

Proof of Theorem 4. Let $\hat{\sigma}$ be a proper equilibrium. Using the characterization of proper equilibrium given by Proposition 8 of ? (and restated as Proposition 8 of this paper), there must exist some LPS $\rho = (p^1, \dots, p^K)$ on $S_1 \times S_2$ that satisfies strong independence, has full support, respects within-person preferences, and for which $(\rho, \hat{\sigma})$ is a lexicographic Nash equilibrium.

Let δ be as in Lemma 12. We also define

$$m_0 = \min_{\substack{n \in \{1,2\} \\ k \in \{1, \dots, K\} \\ s_n \in \text{supp } p_n^k}} p_n^k(s_n).$$

Fix any $\alpha \in \mathbb{R}_{++}^2$ and any $\varepsilon \in \left(0, \min \left\{ \delta, \frac{m_0}{m_0+1} \right\}\right)$. For any $r \in [0, 1]^{K-1}$, we use $(r \square \rho)_1$ and $(r \square \rho)_2$ to denote the marginals of $r \square \rho$ on S_1 and S_2 . Using this notation, for all $n \in \{1, 2\}$, we define the set

$$\Delta_{n,\varepsilon}^* = \left\{ \sigma_n \in \Delta_n \mid \sigma_n = (r \square \rho)_n \text{ for some } r \in \left[\varepsilon^{2K^2+K}, \varepsilon \right]^{K-1} \right\}.$$

In words, $\Delta_{n,\varepsilon}^*$ is a set of strategies for player n consisting of mixtures of his play under the different levels of ρ such that each level of ρ receives relatively less weight than the level before it, where the range of permissible relative weights is controlled by ε . Furthermore, all strategies in $\Delta_{n,\varepsilon}^*$ are totally mixed, as a consequence of ρ having full support. We also define the correspondence $F : \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^* \rightrightarrows \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ by

$$F(\sigma) = \left\{ \sigma^* \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^* \mid \begin{array}{l} (\forall l \in \{1, 2\})(\forall m \in \{1, 2\})(\forall s'_l \in S_l)(\forall s''_m \in S_m) \\ \text{if } \alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m) \\ \text{then } \sigma_l^*(s'_l) \leq \varepsilon \cdot \sigma_m^*(s''_m) \end{array} \right\}.$$

Claim: $\Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$, equipped with the Euclidean metric, is a compact metric space and an absolute retract.

Proof of Claim: It follows immediately from the definition of $\Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ that it is compact. A compact subset of a Euclidean space is an absolute retract if and only if it is contractible and locally contractible (e.g., ?, fn. 10). We argue both of these below.

Fix some $\bar{r} \in \left[\varepsilon^{2K^2+K}, \varepsilon \right]^{K-1}$. Define $H : \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^* \times [0, 1] \rightarrow \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ in the following way. For any $\sigma^* \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$, we can write $\sigma^* = ((r_1^* \square \rho)_1, (r_2^* \square \rho)_2)$ for some $r_1^*, r_2^* \in \left[\varepsilon^{2K^2+K}, \varepsilon \right]^{K-1}$. Then define $H(\sigma^*, t)$ to be the strategy profile

$$\left(([r_1^*(1-t) + \bar{r}t] \square \rho)_1, ([r_2^*(1-t) + \bar{r}t] \square \rho)_2 \right).$$

It follows from this definition that H is continuous, that $H(\sigma^*, t) \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$, that $H(\sigma^*, 0) = \sigma^*$, and that $H(\sigma^*, 1) = ((\bar{r}_1 \square \rho)_1, (\bar{r}_2 \square \rho)_2)$. Therefore, H is a contracting homotopy and $\Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ is contractible.

A similar argument establishes the local version of the property. Fix some $(\bar{r}_1, \bar{r}_2) \in [\varepsilon^{2K^2+K}, \varepsilon]^{K-1} \times [\varepsilon^{2K^2+K}, \varepsilon]^{K-1}$. Let $U \subset \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ be some neighborhood of $((\bar{r}_1 \square \rho)_1, (\bar{r}_2 \square \rho)_2)$. By continuity of the “square operator,” there must then exist some ball $B \subset [\varepsilon^{2K^2+K}, \varepsilon]^{K-1} \times [\varepsilon^{2K^2+K}, \varepsilon]^{K-1}$ around (\bar{r}_1, \bar{r}_2) such that for all $(r_1, r_2) \in B$, $((r_1 \square \rho)_1, (r_2 \square \rho)_2) \in U$. Let $V \subset U$ consist of the strategy profiles that can be formed in this way. Then define $H : V \times [0, 1] \rightarrow V$ in the following way. For any $\sigma^* \in V$, we can write $\sigma^* = ((r_1^* \square \rho)_1, (r_2^* \square \rho)_2)$ for some $r_1^*, r_2^* \in B$. Then define $H(\sigma^*, t)$ to be the strategy profile

$$\left(([r_1^*(1-t) + \bar{r}_1 t] \square \rho)_1, ([r_2^*(1-t) + \bar{r}_2 t] \square \rho)_2 \right).$$

It follows from this definition that H is continuous, that $H(\sigma^*, t) \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$, that $H(\sigma^*, 0) = \sigma^*$, and that $H(\sigma^*, 1) = ((\bar{r}_1 \square \rho)_1, (\bar{r}_2 \square \rho)_2)$. Therefore, V is contractible, and so $\Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ is locally contractible.

Claim: For all $\sigma \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$, $F(\sigma)$ is nonempty.

Proof of Claim: For $n \in \{1, 2\}$ and $k \in \{1, \dots, K\}$, define

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

Because ρ has full support, $\{\pi_n^1, \dots, \pi_n^K\}$ is a partition of S_n . Because ρ respects within-person preferences, $(\forall n \in \{1, 2\})(\forall k \in \{1, \dots, K\})(\forall s'_n \in \pi_n^k)(\forall s''_n \in \cup_{j=k}^K \pi_n^j)$:

$$\left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s'_n) \right]_{k=1}^K \geq_L \left[\sum_{s \in \bar{S}} p^k(s) \pi_n(s/s''_n) \right]_{k=1}^K.$$

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

Suppose now that $\sigma \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$. It must therefore be of the form $((r_1 \square \rho)_1, (r_2 \square \rho)_2)$, where $r_1, r_2 \in [\varepsilon^{2K^2+K}, \varepsilon]^{K-1}$. Because $\varepsilon < \delta$, we also have, by Lemma 12, that $(\forall n \in \{1, 2\})(\forall k \in \{1, \dots, K\})(\forall s'_n \in \pi_n^k)(\forall s''_n \in \cup_{j=k}^K \pi_n^j)$:

$$\sum_{s \in \bar{S}} (r_{-n} \square \rho)(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)] \geq 0.$$

Next, we observe that because this is a two-player game, each player n has just a single opponent, so that both $\sigma = ((r_1 \square \rho)_1, (r_2 \square \rho)_2)$ and $r_{-n} \square \rho$ induce the same distribution of play over the opponent’s strategies.¹⁸ Thus, $(\forall n \in \{1, 2\})(\forall s'_n \in S_n)(\forall s''_n \in S_n)$:

$$\sum_{s \in \bar{S}} (r_{-n} \square \rho)(s) [\pi_n(s/s'_n) - \pi_n(s/s''_n)] = \pi_n(\sigma/s'_n) - \pi_n(\sigma/s''_n) = L_n(\sigma/s''_n) - L_n(\sigma/s'_n).$$

¹⁸In contrast, with three or more players, r_{-n} would not even be a well-defined vector.

Thus, we have that $(\forall n \in \{1, 2\})(\forall k \in \{1, \dots, K\})(\forall s'_n \in \pi_n^k)(\forall s''_n \in \cup_{j=k}^K \pi_n^j)$:

$$L_n(\sigma/s''_n) - L_n(\sigma/s'_n) \geq 0.$$

In particular, if $s'_n, s''_n \in \pi_n^k$, then by symmetry we must have $L_n(\sigma/s'_n) = L_n(\sigma/s''_n)$. A consequence of the above is that there exist, for both players $n \in \{1, 2\}$, nondecreasing sequences $L_n^1 \leq \dots \leq L_n^K$ such that $(\forall k \in \{1, \dots, K\})(\forall s_n \in \pi_n^k) : \alpha_n L_n(\sigma/s_n) = L_n^k$.

Next, for all $n \in \{1, 2\}$, all $m \in \{1, 2\}$, and all $k \in \{1, \dots, K\}$, define $\phi_{nm}^k = |\{j \in \{1, \dots, K\} : L_m^j < L_n^k\}|$. Let $\phi_n^k = \sum_{m \in \{1, 2\}} \phi_{nm}^k$. Then for all $n \in \{1, 2\}$ and $k \in \{1, \dots, K-1\}$, let

$$r_n^k = \varepsilon^{1+(K+1)(\phi_n^{k+1} - \phi_n^k)}. \quad (5)$$

For all n and k we have $L_n^{k+1} \geq L_n^k$, which implies $\phi_n^{k+1} \geq \phi_n^k$, and so $r_n^k \leq \varepsilon < \delta$. We also have $\phi_n^{k+1} \leq 2K-1$, and so $r_n^k = \varepsilon^{1+(K+1)(\phi_n^{k+1} - \phi_n^k)} \geq \varepsilon^{2K^2+K}$. From the previous two inequalities we conclude $(r_n \square \rho)_n \in \Delta_{n, \varepsilon}^*$.

Next, we verify that $((r_1 \square \rho)_1, (r_2 \square \rho)_2) \in F(\sigma)$. To that end, suppose $l, m \in \{1, 2\}$, $s'_l \in S_l$ and $s''_m \in S_m$ are such that $\alpha_l L_l(\sigma/s'_l) > \alpha_m L_m(\sigma/s''_m)$. Let k' and k'' be the indices for which $s'_l \in \pi_l^{k'}$ and $s''_m \in \pi_m^{k''}$. We must then have $\phi_l^{k'} \geq \phi_m^{k''} + 1$. Therefore, we have the following (where to accommodate the accommodate the case of $k'' = K$, the argument remains valid if we define $r_m^K = 0$):

$$\begin{aligned} (r_l \square \rho)_l(s'_l) &\leq r_l^1 \cdots r_l^{k'-1} \\ &= \varepsilon^{k'-1+(K+1)\phi_l^{k'}} \\ &\leq \varepsilon^{k'+K+(K+1)\phi_m^{k''}} \\ &= \varepsilon^{k''+(K+1)\phi_m^{k''}} (1-\varepsilon)m_0 \cdot \frac{\varepsilon^{k'+K-k''}}{(1-\varepsilon)m_0} \\ &\leq \varepsilon^{k''+(K+1)\phi_m^{k''}} (1-\varepsilon)m_0 \cdot \frac{\varepsilon}{(1-\varepsilon)m_0} \\ &\leq \varepsilon^{k''+(K+1)\phi_m^{k''}} (1-\varepsilon)m_0 \\ &\leq \varepsilon^{k''+(K+1)\phi_m^{k''}} (1-r_m^{k''})m_0 \\ &= \varepsilon \cdot r_m^1 \cdots r_m^{k''-1} (1-r_m^{k''})m_0 \\ &\leq \varepsilon \cdot (r_m \square \rho)_m(s''_m). \end{aligned}$$

In the above, the first step follows from the definition of the “square operator” by Lemma 11. The second and eighth steps use (5) and the fact that $\phi_l^1 = \phi_m^1 = 0$. The third step uses $\phi_l^{k'} \geq \phi_m^{k''} + 1$. The fourth step is algebra. The fifth step uses $K \geq k''$ and $k' \geq 1$. The sixth step uses $\varepsilon < \frac{m_0}{1+m_0}$, which implies that $\frac{\varepsilon}{(1-\varepsilon)m_0} < 1$. The seventh step uses $r_m^{k''} \leq \varepsilon$. The ninth step uses the definitions of the “square operator” and m_0 by Lemma 11.

Claim: For all $\sigma \in \Delta_{1, \varepsilon}^* \times \Delta_{2, \varepsilon}^*$, $F(\sigma)$ is contractible.

Proof of Claim: The points in $F(\sigma)$ are those that satisfy a finite collection of linear inequalities, so $F(\sigma)$ is a convex set. Then, since $F(\sigma)$ is nonempty by the previous claim, it is contractible.

Claim: F is upper-hemicontinuous.

Proof of Claim: This follows from continuity of L_n .

Claim: $\Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ contains an (α, ε) -extended proper equilibrium.

Proof of Claim: By the above claims, F satisfies all the conditions of Lemma 10, so there exists some $\sigma^\varepsilon \in \Delta_{1,\varepsilon}^* \times \Delta_{2,\varepsilon}^*$ such that $\sigma^\varepsilon \in F(\sigma^\varepsilon)$. By the definition of F , this point is an (α, ε) -extended proper equilibrium.

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

Proof of Claim: For all $t \in \mathbb{N}$, define $\varepsilon_t = \frac{1}{t+1} \min \left\{ \delta, \frac{m_0}{m_0+1} \right\}$. 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 contained in the set $\Delta_{1,\varepsilon_t}^* \times \Delta_{2,\varepsilon_t}^*$. By the definition of $\Delta_{n,\varepsilon}^*$, we must have $\max_{n \in \{1,2\}} \max_{s_n \in S_n} |\sigma_n^\varepsilon(s_n) - p_n^1(s_n)| \leq \varepsilon$, and so $\lim_{t \rightarrow \infty} \sigma^t = (p_1^1, p_2^1)$, 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 5. 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 dramatically 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 . We also use $M = \prod_{n \in \mathcal{N}} |S_n|$.

Necessity of Nash equilibrium. Suppose that Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose that $\bar{\sigma}$ 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

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

which contradicts $(\bar{\sigma}, \dots, \bar{\sigma})$ 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 $\bar{s} = (s_1, \dots, s_N)$,

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

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

Necessity of perfect equilibrium. Suppose that Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose that $\bar{\sigma}$ 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 , $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ 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 $(M\varepsilon_t)$ -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 must then have

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

Because $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ is an ε_t -perfect equilibrium of $\bar{\Gamma}'$, we therefore have $\bar{\sigma}^t(\bar{s}/s'_n) \leq \varepsilon_t$. Since the same argument can be made for all $\bar{s} \in \bar{S}$ (of which there are M), and since σ_n^t is the marginal of $\bar{\sigma}^t$ on S_n , we conclude that $\sigma_n^t(s'_n) \leq M\varepsilon_t$, 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 , $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ is an ε_t -perfect equilibrium of $\bar{\Gamma}$. Let $\bar{s}' = (s'_1, \dots, s'_N)$ and $\bar{s}'' = (s''_1, \dots, s''_N)$ be such that

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

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 $\sigma_n^t(s'_n) \leq \varepsilon_t$. Since σ_n^t is the marginal of $\bar{\sigma}^t$ on S_n , we have $\bar{\sigma}^t(\bar{s}') \leq \sigma_n^t(s'_n)$, so that $\bar{\sigma}^t(\bar{s}') \leq \varepsilon_t$, as required.

Necessity of extended proper equilibrium. Suppose that Γ' is equivalent to Γ up to affine transformations of the payoffs, and suppose that $\bar{\sigma}$ 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 , $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ 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 $(\alpha, M\varepsilon_t)$ -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

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

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

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

Moreover, since we can make the above argument for each strategy profile $\hat{s} \in \bar{S}$ for which $\hat{s}_l = s'_l$ (of which there are at most M), and since σ_l^t is the marginal of $\bar{\sigma}^t$ on S_l , we conclude $\sigma_l^t(s'_l) \leq M \varepsilon_t \sigma_m^t(s''_m)$, 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 $(\alpha, \varepsilon_t^{M+2})$ -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 , $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ 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:

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

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 $s_n^* \in \arg \max_{s_n \in S_n} \sigma_n^t(s_n)$. Note that because σ^t is an $(\alpha, \varepsilon_t^{M+2})$ -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 single 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 $(\alpha, \varepsilon_t^{M+2})$ -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 $\varepsilon_t^{M+2} \phi^t(s^*)$. Hence, the elements of \mathcal{S}^3 must all be maximal under \succcurlyeq .

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

- [Modify s' itself:] At the step for $s' \in \mathcal{S}^2$, select some $\hat{s} \in \arg \min_{s \in \{s^*\} \cup \mathcal{S}^1 : s \succ s'} \phi^t(s)$ 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 $\bar{\sigma}^t(s^*) \leftarrow \bar{\sigma}^t(s^*) - (|\{n \in \mathcal{N} : s'_n \neq s_n^*\}| - 1) (\phi^t(s') - \bar{\sigma}^t(s'))$.

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 some arbitrary $s' \in \mathcal{S}^2$ does not alter the marginal distributions. For some arbitrary $m \in \mathcal{N}$ and some arbitrary $s_m \in S_m$, 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'), $(|\{n \in \mathcal{N} : s'_n \neq s_n^*\}| - 1) (\phi^t(s') - \bar{\sigma}^t(s'))$ 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 $(|\{n \in \mathcal{N} : s'_n \neq s_n^*\}| - 1) (\phi^t(s') - \bar{\sigma}^t(s'))$ less weight to one strategy profile involving s_m^* (namely s^*), as well as $\phi^t(s') - \bar{\sigma}^t(s')$ more weight to $|\{n \in \mathcal{N} : s'_n \neq s_n^*\}| - 1$ 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 . 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 other than the ones considered above.

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

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, we must have

$$\phi^t(s') = \sigma_m^t(s_m) \prod_{n \neq m} \sigma_n^t(s'_n) \leq \sigma_m^t(s_m) \prod_{n \neq m} \sigma_n^t(s_n^*) = \phi^t(s).$$

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

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

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 M such steps, we conclude that $\bar{\sigma}^t(s) \geq (1 - M\varepsilon_t^2)\phi^t(s)$, as desired.

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

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 M such steps, the conclusion follows.

Next, define T such that for all $t \geq T$, both $1 - NM\varepsilon_t \geq 0$ and

$$\varepsilon_t \geq \max \left\{ \frac{M^2 \varepsilon_t^{M+2}}{1 - NM\varepsilon_t}, \frac{\varepsilon_t^2}{1 - NM\varepsilon_t}, M^2 \varepsilon_t^2 \right\}.$$

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 $\bar{s} = (s_1, \dots, s_N)$ 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 $t \geq T$, $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ is an ε_t -proper equilibrium of $\bar{\Gamma}'$.¹⁹

Proof of Claim: Suppose that $\bar{s}' = (s'_1, \dots, s'_N)$ and $\bar{s}'' = (s''_1, \dots, s''_N)$ 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 for all $t \geq T$, $\bar{\sigma}^t(\bar{s}'') \leq \varepsilon_t \bar{\sigma}^t(\bar{s}')$, which in turn implies that $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ 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.

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 $(\alpha, \varepsilon_t^{M+2})$ -extended proper equilibrium of Γ , we must have that $\sigma_m^t(s''_m) \leq \varepsilon_t^{M+2} \sigma_m^t(s^*_m)$. Because σ_m^t is the marginal of ϕ^t on S_m , we must also have $\phi^t(s'') \leq \sigma_m^t(s''_m)$. And because $s_n^* \in \arg \max_{s_n \in S_n} \sigma_n^t(s_n)$ for all $n \in \mathcal{N}$, we must have $\sigma_m^t(s^*_m) \leq M\phi^t(s^*)$. Further, as

¹⁹Because $\bar{\Gamma}'$ is a symmetric game and because the strategy profile in question is symmetric, we will have also shown that $(\bar{\sigma}^t, \dots, \bar{\sigma}^t)$ is an (ι, ε_t) -extended proper equilibrium, where ι denotes a vector of ones.

derived above, $\bar{\sigma}^t(s'') \leq M\phi^t(s'')$ and $\bar{\sigma}^t(s^*) \geq (1 - NM\varepsilon_t)\phi^t(s^*)$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq M\phi^t(s'') \leq M\sigma_m^t(s''_m) \leq M\varepsilon_t^{M+2}\sigma_m^t(s^*_m) \leq M^2\varepsilon_t^{M+2}\phi^t(s^*) \leq \frac{M^2\varepsilon_t^{M+2}}{1 - NM\varepsilon_t}\bar{\sigma}^t(s^*).$$

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 $(\alpha, \varepsilon_t^{M+2})$ -extended proper equilibrium of Γ , we must have that $\sigma_m^t(s''_m) \leq \varepsilon_t^{M+2}\sigma_l^t(s'_l)$. Because σ_m^t is the marginal of ϕ^t on S_m , we must also have $\phi^t(s'') \leq \sigma_m^t(s''_m)$. And because $s_n^* \in \arg \max_{s_n \in S_n} \sigma_n^t(s_n)$ for all $n \in \mathcal{N}$, we must have $\sigma_l^t(s'_l) \leq M\phi^t(s')$. Further, as derived above, $\bar{\sigma}^t(s'') \leq M\phi^t(s'')$ and $\bar{\sigma}^t(s') \geq (1 - M\varepsilon_t^2)\phi^t(s')$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq M\phi^t(s'') \leq M\sigma_m^t(s''_m) \leq M\varepsilon_t^{M+2}\sigma_l^t(s'_l) \leq M^2\varepsilon_t^{M+2}\phi^t(s') \leq \frac{M^2\varepsilon_t^{M+2}}{1 - M\varepsilon_t^2}\bar{\sigma}^t(s').$$

Note that $\frac{M^2\varepsilon_t^{M+2}}{1 - M\varepsilon_t^2} \leq \frac{M^2\varepsilon_t^{M+2}}{1 - NM\varepsilon_t}$. 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 $s'' \in \mathcal{S}^2$. Because $s' \succ s''$, its construction implies $\bar{\sigma}^t(s'') \leq \varepsilon_t^2\phi^t(s')$. Further, as derived above, $\bar{\sigma}^t(s') \geq (1 - M\varepsilon_t^2)\phi^t(s')$. Combining these inequalities:

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

Note that $\frac{\varepsilon_t^2}{1 - M\varepsilon_t^2} \leq \frac{\varepsilon_t^2}{1 - NM\varepsilon_t}$. 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, $\bar{\sigma}^t(s') \geq \varepsilon_t^M\phi^t(\hat{s})$. We must also have $\hat{s} \succ s' \succ s''$. 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 $(\alpha, \varepsilon_t^{M+2})$ -extended proper equilibrium of Γ , we must have that $\sigma_m^t(s''_m) \leq \varepsilon_t^{M+2}\sigma_l^t(\hat{s}_l)$. Because σ_m^t is the marginal of ϕ^t on S_m , we must also have $\phi^t(s'') \leq \sigma_m^t(s''_m)$. And because $s_n^* \in \arg \max_{s_n \in S_n} \sigma_n^t(s_n)$ for all $n \in \mathcal{N}$, we must have $\sigma_l^t(\hat{s}_l) \leq M\phi^t(\hat{s})$. Further, as derived above, $\bar{\sigma}^t(s'') \leq M\phi^t(s'')$. Combining these inequalities:

$$\bar{\sigma}^t(s'') \leq M\phi^t(s'') \leq M\sigma_m^t(s''_m) \leq M\varepsilon_t^{M+2}\sigma_l^t(\hat{s}_l) \leq M^2\varepsilon_t^{M+2}\phi^t(\hat{s}) \leq M^2\varepsilon_t^2\bar{\sigma}^t(s').$$

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}', ℓ') and (\hat{s}'', ℓ'') be as specified for the steps of the loop corresponding to s' and s'' , respectively. Because $s' \succ s''$, we must have $\phi^t(\hat{s}') \geq \phi^t(\hat{s}'')$ and $\ell' \leq \ell'' + 1$. Then

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

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, $\bar{\sigma}^t(s^*) \geq (1 - NM\varepsilon_t)\phi^t(s^*)$. Combining these inequalities:

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

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 $\bar{\sigma}^t(s'') \leq M^2 \varepsilon_t^{M+2} \phi^t(s^*)$. Because $s' \in \mathcal{S}^3$, $\bar{\sigma}^t(s') = \phi^t(s') \geq \varepsilon_t \phi^t(s^*)$. Combining these inequalities:

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

Note that $M^2 \varepsilon_t^{M+1} \leq M^2 \varepsilon_t^2$. 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 $\bar{\Gamma}'$.²⁰ Furthermore, the marginals of each $\bar{\sigma}^t$ are given by the profile σ^t . Therefore, the projection of the aforementioned limit must equal $\lim_{t \rightarrow \infty} \sigma^t = \sigma$. \square

B.3 Proofs Corresponding to Section 5

Proof of Proposition 6. *Sufficiency.* Suppose (ρ, σ) is a lexicographic Nash equilibrium, where $\rho = (p^1, \dots, p^K)$ is an LPS that satisfies strong independence. As ? observe, strong independence of ρ implies that p^1 must be a product measure. Furthermore, if $p_n^1(s'_n) > 0$ for a player n , then by condition (i) of Definition 6, for all $s''_n \in S_n$,

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

Therefore (p_1^1, \dots, p_N^1) satisfies the definition of a Nash equilibrium. Thus, σ is a Nash equilibrium, since by condition (ii) of Definition 6, $\sigma = (p_1^1, \dots, p_N^1)$.

Necessity. Suppose that $\sigma = (\sigma_1, \dots, \sigma_N)$ is a Nash equilibrium. Let $K = 1$ and define $p^1 = \prod_{n \in \mathcal{N}} \sigma_n$ be 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 measure. Because $K = 1$, ρ therefore satisfies strong independence. \square

Proof of Proposition 7. This is a restatement of Proposition 7 from ?. \square

Proof of Proposition 8. This is a restatement of Proposition 8 from ?. \square

²⁰Following footnote 19, 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 9. *Sufficiency.* Suppose (ρ, σ) is a lexicographic Nash equilibrium, where $\rho = (p^1, \dots, p^K)$ is an LPS that satisfies strong independence, has full support, and respects within-and-across-person preferences. As ? point out, ρ satisfies strong independence if and only if there exists a sequence $r(t) \in (0, 1)^{K-1}$ with $r(t) \rightarrow 0$ such that $r(t) \square \rho$ is a product measure for all t . For each $n \in \mathcal{N}$, define $\sigma_n(t)$ as the marginal of $r(t) \square \rho$ on S_n , and let $\sigma(t) = (\sigma_1(t), \dots, \sigma_N(t))$. Note that $\lim_{t \rightarrow \infty} \sigma(t) = (p_1^1, \dots, p_N^1)$, which by condition (ii) of Definition 6 is equal to σ . Also define

$$m_0 = \min_{\substack{n \in \mathcal{N} \\ k \in \{1, \dots, K\} \\ s_n \in \text{supp } p_n^k}} p_n^k(s_n)$$

$$\varepsilon(t) = \max_{k \in \{1, \dots, K-1\}} \left\{ \frac{r^k(t)}{[1 - r^k(t)]m_0} \right\}$$

Because $r(t) \rightarrow 0$, we also have $\varepsilon(t) \rightarrow 0$. Letting δ be as in the statement of Lemma 13, 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-person 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 $r(t) \square \rho$ is a product measure with marginals $\sigma_1(t), \dots, \sigma_N(t)$, the above inequality is equivalent to

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

By Lemma 13, we obtain that for all $s_l^* \in BR_l(\rho)$ and all $s_m^* \in BR_m(\rho)$,

$$\left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s_l^*) - \pi_l(s/s'_l)] \right]_{k=1}^K >_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s_m^*) - \pi_m(s/s''_m)] \right]_{k=1}^K,$$

which, because ρ respects within-and-across-person 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 must then have $\sigma_m(t)(s''_m) \geq r^1(t) \dots r^{k''-1}(t) [1 - r^{k''}(t)] m_0$, by Lemma 11. Because $s'_l <_\rho s''_m$, we must also have $\sigma_l(t)(s'_l) \leq r^1(t) \dots r^{k''}(t)$ by Lemma 11. Combining this with the definition of $\varepsilon(t)$ we have

$$\begin{aligned} \varepsilon(t) \cdot \sigma_m(t)(s''_m) &\geq \varepsilon(t) \cdot r^1(t) \dots r^{k''-1}(t) [1 - r^{k''}(t)] m_0 \\ &\geq \frac{r^{k''}(t)}{[1 - r^{k''}(t)] m_0} \cdot r^1(t) \dots r^{k''-1}(t) [1 - r^{k''}(t)] m_0 \\ &= r^1(t) \dots r^{k''}(t) \\ &\geq \sigma_l(t)(s'_l). \end{aligned}$$

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 $\sigma_l(t)(s'_l) \leq \varepsilon(t) \cdot \sigma_m(t)(s''_m)$. Consequently, $\sigma(t)$ is an $(\alpha, \varepsilon(t))$ -extended proper equilibrium. Because $\varepsilon(t) \rightarrow 0$ and $\sigma(t) \rightarrow \sigma$, σ is an extended proper equilibrium.

Necessity. Suppose $\sigma = (\sigma_1, \dots, \sigma_N)$ 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 ?, there is an LPS $\rho = (p^1, \dots, p^K)$ on \bar{S} such that a subsequence $\bar{\sigma}(\tau)$ of $\bar{\sigma}(t)$ can be written as $\bar{\sigma}(\tau) = r(\tau) \square \rho$ for a sequence $r(\tau) \in (0, 1)^{K-1}$ with $r(\tau) \rightarrow 0$. We claim that ρ satisfies strong independence, has full support, respects within-and-across-person comparisons, and is such that (ρ, σ) meets condition (ii) of Definition 6.

First, condition (ii) of Definition 6 is satisfied, since $p^1 = \lim_{\tau \rightarrow \infty} r(\tau) \square \rho = \lim_{\tau \rightarrow \infty} \bar{\sigma}(\tau) = \lim_{\tau \rightarrow \infty} \prod_{n \in \mathcal{N}} \sigma_n(\tau) = \prod_{n \in \mathcal{N}} \sigma_n$. 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 , $r(\tau) \square \rho$ is a product measure; moreover, $r(\tau) \rightarrow 0$. ? point out that these facts imply that ρ satisfies strong independence. Finally, we will show that ρ respects within-and-across-person 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

$$\begin{aligned} & \left[\alpha_l \sum_{s \in \bar{S}} p^k(s) [\pi_l(s/s'_l) - \pi_l(s/s''_m)] \right]_{k=1}^K \\ & >_L \left[\alpha_m \sum_{s \in \bar{S}} p^k(s) [\pi_m(s/s'_l) - \pi_m(s/s''_m)] \right]_{k=1}^K. \end{aligned} \quad (6)$$

Letting δ be as in the statement of Lemma 13, let T be such that for all $\tau \geq T$, $r(\tau) \in (0, \delta)^{K-1}$. By Lemma 13, for all $\tau \geq T$,

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

Because $r(\tau) \square \rho$ is a product measure with marginals $\sigma_1(\tau), \dots, \sigma_N(\tau)$, 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,

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

Define $k' := \min\{k : p_l^k(s'_l) > 0\}$. 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 $r(\tau) \rightarrow 0$ and $\varepsilon(\tau) \rightarrow 0$, there must exist some $\tau^* \geq T$ for which $\varepsilon(\tau^*) < [1 - r^{k'}(\tau^*)]p_l^{k'}(s'_l)$. Using Lemma 11 and then this inequality,

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

Suppose by way of contradiction that $s''_m \leq_\rho s'_l$, so that $\min\{k : p_m^k(s''_m) > 0\} \geq k'$, and thus $\sigma_m(\tau^*)(s''_m) \leq r^1(\tau^*) \cdots r^{k'-1}(\tau^*)$ by Lemma 11. Then we would have $\sigma_l(\tau^*)(s'_l) > \varepsilon(\tau^*) \cdot \sigma_m(s''_m)$. 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-person preferences. \square

C Connections to Test-Set Equilibrium

Many of the referees' comments were requests for more clarity about how extended proper equilibrium relates to test-set equilibrium. I thought that adding an appendix might help.

Test-set equilibrium (?), 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

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

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

Definition 11. 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 12. 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.

Effectively, test-set equilibrium require that any single deviation to a strategy that is a best response to equilibrium should be treated as more likely than any single 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 test-set equilibrium also requires that any single deviation to a strategy that is a best response to equilibrium should be treated as more likely than any joint deviation. Extended proper equilibrium does not imply the latter.

On the other hand, extended proper equilibrium imposes restrictions not required by test-set equilibrium (e.g., that players tremble independently). In general, neither solution concept implies the other in games with three or more players.

C.2 Example of a Game Without a Test-Set Equilibrium

In ?, we present an example of a game—reproduced below as Figure 2—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 because $East$ weakly dominates $West$ in the test set:

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

Figure 2: A three player game[†]

[West]	Left	Center	Right		[East]	Left	Center	Right
<i>Up</i>	0, 0, 0	0, 0, 0	0, 0, 0		<i>Up</i>	0, 0, 0	0, -1, 1	0, 1, 0
<i>Down</i>	0, 0, 0	-1, 1, 0	1, -1, 0		<i>Down</i>	-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.

Despite failing the test-set condition, $(Up, Left, West)$ 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 small choice of δ :²¹

$$\begin{aligned} \sigma_{row}^\delta &= \left(\frac{2 + \delta^2}{2 + 3\delta + \delta^2}, \frac{3\delta}{2 + 3\delta + \delta^2} \right) \\ \sigma_{col}^\delta &= \left(\frac{2 - 2\delta}{2 + \delta}, \frac{3\delta^2}{2 + 3\delta + \delta^2}, \frac{3\delta}{2 + 3\delta + \delta^2} \right) \\ \sigma_{geo}^\delta &= \left(\frac{2 - 2\delta}{2 + \delta}, \frac{3\delta}{2 + \delta} \right) \end{aligned}$$

Moreover, $\lim_{\delta \rightarrow 0} \sigma^\delta = (Up, Left, West)$.

²¹To 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:

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

Note that as $\delta \rightarrow 0$, Column's tremble to *Center* becomes infinitely less likely than its tremble to *Right*, despite the fact that they are both best responses to the equilibrium (*Up, Left, West*). In that sense, these trembles might be regarded as extreme. Test-set equilibrium effectively prohibits such extreme trembling. As a result of this extreme trembling, the single deviation to *Center* is less likely than the joint deviation to (*Down, Right*), which runs contrary to the logic embedded in the test-set condition.

C.3 Generalized Second-Price Auction

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, whose values per click are 4, 1, and 1, respectively. There are two ad positions, which attract two clicks and one click, respectively. For the discretized bid set, let $m = 3$ and $M = 3$, although the arguments can easily be extended to accommodate $m \geq 3$ and $M \geq 3$.

The profile $b^* = (2, 1, 1)$ is a test-set equilibrium.²² 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 given that its opponents will be trembling independently.

Building on Proposition 7, 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 that bidder 1 contemplates raising its bid from $b_1^* = 2$ to $\frac{7}{3}$. This change makes a difference only if at least one of its opponents trembles within the set $\{2, \frac{7}{3}\}$. Thus, the relevant profiles can be partitioned into five sets:

- *Case 1(a)*: $b_2 = \frac{7}{3}$ and $b_3 = 1$
- *Case 1(b)*: $b_2 = 1$ and $b_3 = \frac{7}{3}$
- *Case 2(a)*: $b_2 = \frac{7}{3}$ and $b_3 \in \{0, \frac{1}{3}\}$
- *Case 2(b)*: $b_2 \in \{0, \frac{1}{3}\}$ and $b_3 = \frac{7}{3}$
- *Case 3*: all remaining profiles in which $b_2 \in \{2, \frac{7}{3}\}$ and/or $b_3 \in \{2, \frac{7}{3}\}$

Against the profile in case 1(a), bidder 1 is strictly better off deviating to $\frac{7}{3}$:

$$\pi_1(\frac{7}{3}, \frac{7}{3}, 1) = \frac{19}{6} > 3 = \pi_1(2, \frac{7}{3}, 1).$$

Symmetrically, bidder 1 is also strictly better off deviating to $\frac{7}{3}$ against the profile in case 1(b). Now, bidder 1 would be strictly worse off deviating to $\frac{7}{3}$ against the profiles in cases 2(a) and 2(b). For example, consider $b_2 = \frac{7}{3}$ and $b_3 = 0$:

$$\pi_1(\frac{7}{3}, \frac{7}{3}, 0) = \frac{11}{3} < 4 = \pi_1(2, \frac{7}{3}, 0).$$

²²For bidder 1, the dominated bids are those greater than 4 and those less than 2. For bidders 2 and 3, the dominated bids are those greater than 1 and those less than $\frac{1}{2}$.

Bidder 1 best responds to its opponents' equilibrium bids (1, 1) with any bid greater than 1. Each of bidders 2 and 3 best responds to its opponents' equilibrium bids (2, 1) with any bid less than 2. From these observations, the test set can be constructed and the test-set condition checked.

But these are the only profiles against which bidder 1 is made strictly worse off by deviating to $\frac{7}{3}$; in particular, bidder 1 is not worse off deviating against any of the profiles in case 3. Moreover, it follows from strong independence of ρ that each profile in case 2(a) or 2(b) is infinitely less likely under ρ than the corresponding profile in case 1(a) or 1(b). We conclude from this analysis that $\frac{7}{3}$ is a profitable deviation against ρ for bidder 1, which contradicts (ρ, b^*) being a lexicographic Nash equilibrium.

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