

On the Existence of Monotone Pure-Strategy Perfect Bayesian Equilibrium in Games with Complementarities*

Jeffrey Mensch[†]

November 6, 2016

Abstract

Many important economic situations can be modelled as dynamic games of incomplete information with strategic complementarities of actions and types. In this paper, we extend the results of Athey (2001) and Reny (2011) from static Bayesian games to dynamic environments, providing conditions that guarantee the existence of monotone equilibria in types in such games. A feature that distinguishes this environment from those of previous results is the endogeneity of beliefs, which can complicate continuity of payoffs, needed to find a fixed point. To address this, we define an auxiliary static game which pins down beliefs while preserving continuity of payoffs. We also provide conditions which guarantee that there will exist monotone best-replies to monotone strategies of one's opponents in a dynamic environment. Applications are given to signaling games and stopping games such as auctions, wars of attrition, and joint research projects.

Keywords: Games of incomplete information, dynamic Bayesian games, pure strategy equilibrium, perfect Bayesian equilibrium, equilibrium existence, auctions, signaling games, supermodular games, single crossing property

*I am grateful to Wojciech Olszewski, Alessandro Pavan, Marciano Siniscalchi, and Bruno Strulovici for their advice and guidance during this project, as well as Eddie Dekel, Srihari Govindan, Roger Myerson, Phil Reny, Ron Siegel, and Teddy Mekonnen for fruitful conversations. All errors are my own.

[†]Northwestern University; jmensch@u.northwestern.edu

1 Introduction

Many important economic situations can be modelled as dynamic games of incomplete information with strategic complementarities of actions and types. These complementarities can be informational, in the sense that an agent may have information that tends to influence his own or others' actions in a certain direction; alternatively, the complementarities can be strategic, in that higher actions may influence other agents' actions to tend higher as well. Some well-known examples in the economic literature where these components come into play include the signaling models of Spence (1973) and Crawford and Sobel (1982); the models of bargaining with uncertainty, such as those of Grossman and Perry (1986), Gul, Sonnenschein, and Wilson (1986), and Gul and Sonnenschein (1988); reputation models, such as that of Kreps and Wilson (1982a); and various dynamic auctions, as analyzed in Milgrom and Weber (1982). Hence a general description of the equilibria of such games would be of major significance across a wide array of economic topics. This paper provides conditions for complementarities under which an equilibrium in strategies that are monotone in types within each subgame is guaranteed to exist.

A large literature has been developed to explore the equilibria of games with strategic complementarities in games with simultaneous moves. Vives (1990) and Milgrom and Roberts (1990) show that pure strategy Nash equilibrium exists in supermodular games; these results have been extended to games with other types of complementarities, such as quasisupermodularity in Milgrom and Shannon (1994). Later results by Athey (2001), McAdams (2003), Reny and Zamir (2004), Van Zandt and Vives (2007), and Reny (2011) demonstrate the existence of monotone pure-strategy equilibrium in various classes of games of incomplete information.

By contrast, there have been relatively few papers attempting to extend these results to dynamic games. In terms of games without private information, Curtat (1996), Vives (2009), and Balbus, Reffett, and Wozny (2013) consider environments with strategic complementarity and Markov payoffs. Echenique (2004) extends the lattice properties of the set of equilibria in games with strategic complementarities to a restrictive class of dynamic games. In games with private information, Athey (2001), Okuno-Fujiwara, Postlewaite, and Suzumura (1990), Van Zandt and Vives (2007), Hafalir and Krishna (2009), and Zheng (2014) consider various specific examples of games with complementarities for which they show existence of monotone equilibrium. However, none of these approaches study existence under general conditions for multi-period games. Our purpose, therefore, is to provide a straightforward existence result that can be applied off-the-shelf without resorting to complicated constructions. Moreover, even in cases where it will fail to hold, it will be instructive as to the sort of ingredients that

can guarantee existence of monotone equilibrium.

To derive sufficient conditions for monotone equilibrium, we must address the topological conditions which are needed to guarantee existence of such an equilibrium if monotone best-replies exist. To do so, we must address the issue of endogenous beliefs, which does not arise in static environments. The potential concern is that beliefs in subsequent periods will jump around due to small changes in players' strategies. This, in turn, will drastically change the incentives in those periods, and so lead to failures of upper-hemicontinuity of best-replies. By contrast, in static games, any such jump would be "smoothed" under integration, thereby not affecting other players' payoffs much; hence upper-hemicontinuity of best-replies would be preserved.

To circumvent this difficulty, we exploit an additional feature of monotone strategies that is only relevant to dynamic games: the posterior beliefs will be restrictions of the prior to a product of intervals of types. Thus one can always rescale this interval of types to the unit interval. This allows for a novel construction of an auxiliary static game, in which types are endogenously translated to the unit interval at each subgame, and players optimize over continuation strategies.¹ We thus are able to break down the strategies of the players by subgame, and show that a small perturbation of the strategies of the players leads to a continuous perturbation of the beliefs of other players. We then show that equilibrium exists in this static transformation of the game by the existence results of Reny (2011). Finally, we use the equilibrium strategies from the auxiliary game to derive monotone strategies that form a perfect Bayesian equilibrium in the original game.

An interesting feature of the derivation of the existence of equilibrium is that it uniquely pins down the beliefs that must be held upon observing some off-path action. Specifically, in the constructed equilibrium, all other players must place probability one on the highest type to choose a lower action. In tandem with the monotonicity of strategies among on-path actions, this generates beliefs that are "monotone" in the sense that the support of types conditional on observing a higher action is "higher," consisting of intervals that can only overlap at the endpoints. This lends credence to the intuitive notion that a higher type is more likely to have deviated to a higher action, even if off-path, as actions and types are viewed as complementary in monotone equilibria.

Another point of difficulty is the characterization of single-crossing conditions in dynamic games. The existence results described above assumes the existence of monotone best-replies. To guarantee that such best-replies exist, one needs a single-crossing condition. As we show

¹While one might be tempted instead to use "backward induction" in an agent-normal form translation of the game, this will lead to complications in that the strategies in later subgames will not be sufficiently well-behaved; we discuss this in Section 3.

in an example, not even supermodularity of all variables and independence of types, which is sufficient to guarantee existence of monotone equilibrium in static games, is sufficient to guarantee complementarity in dynamic games with at least three periods under the imposition of sequential rationality. This is related to the failure to generate higher beliefs from higher actions in the first two periods: the choice of actions in period 1 affects the choice of actions in period 2, and so can affect what players learn about the types of other players going into period 3. Hence it may be optimal for a higher type to choose a lower action in order to lead to a subgame with more favorable beliefs held by the other players.

Nevertheless, single-crossing can be shown in some more specialized environments that are still of economic interest. Specifically, we show that in the case of two-period games, one-dimensional types, and finite, one-dimensional actions in period 1, a monotone equilibrium exists in the following sense. In the first period, each player's actions are weakly increasing in one's own type. Moreover, holding all other players' actions fixed, each player chooses an action in the second period that is (a) weakly increasing in the actions chosen in the first period, and (b) weakly increasing in one's own type, showing that the best replies of all players are monotonic in both of these senses. We apply this result to show existence of monotone equilibrium in signaling games under fairly general conditions, including a sender with a supermodular payoff in the message and his type, and multiple receivers with private information.

While, as mentioned earlier, single-crossing conditions do not generalize as easily to games with at least three periods, we nevertheless provide some conditions under which these results can be extended. Specifically, at any period in which a player's action set is not a singleton, we restrict the payoff relevance of the continuation game for any path of play for that player to the current period for all but (at most) one choice of action by that player. Despite the strong sufficient conditions that we invoke, these results will still apply to a wide variety of economic environments, including (but not limited to) games with short-lived players, and stopping games such as auctions.

The rest of the paper proceeds as follows. Section 2 describes the model of the games considered in this paper. Section 3 proves existence of a monotone PBE assuming that the best-replies to monotone strategies are increasing in the strong-set order. Section 4 closes the loop, providing primitive conditions under which the best replies to monotone strategies by the other players are also monotone, so that the criteria of the existence theorem will hold. Section 5 provides several applications of the various results found throughout this paper to signaling games and to stopping games. Section 6 concludes.

2 The Model

Consider any arbitrary set S endowed with a partial order \geq_S . For any two elements $s, s' \in S$, the *join* of s and s' , written as $s \vee s'$, is the unique least upper bound of s and s' under \geq_S , i.e. the smallest \hat{s} such that $\hat{s} \geq s$ and $\hat{s} \geq s'$. Conversely, the *meet* of s and s' , written as $s \wedge s'$, is the unique greatest lower bound of s and s' under \geq_S , i.e. the largest \hat{s} such that $\hat{s} \leq s$ and $\hat{s} \leq s'$. The set S is called a *lattice* if for all $s, s' \in S$, we have $s \vee s' \in S$ and $s \wedge s' \in S$. A *sublattice* is a subset $S' \subset S$ that is also a lattice.

Let the game Γ have N players and last T periods. Each player has a type $\theta_i \in \Theta_i \equiv [\underline{\theta}_i, \bar{\theta}_i] \subset \mathbb{R}$, which is private information. In each period t , each player chooses an action $x_t^i \in X_t^i$, where $X_t^i \subset \mathbb{R}$ has a finite number of elements for all $t < T$ and is compact in period T .^{2,3} Define $X = \prod_{t,i} X_t^i$ and $\Theta = \prod_i \Theta_i$, and endow (X, Θ) with the Euclidean partial order. The joint density over types is given by $f(\cdot)$, which we assume (a) is bounded (b) has full support on Θ , and (c) continuous in θ .

The actions taken in periods $1 \leq \tau \leq t$ induce the history, $H^t \in \mathcal{H}^t \equiv \prod_{\tau=1}^{t-1} \prod_{i=1}^N X_\tau^i$. We define $H \in \mathcal{H} \equiv \mathcal{H}^{T+1} \equiv X$ as the full history of the game. Histories are endowed with the partial ordering such that, if $x_\tau \geq \hat{x}_\tau$ for all $\tau < t$, then $H^t \equiv (x_1, \dots, x_{t-1}) \geq (\hat{x}_1, \dots, \hat{x}_{t-1}) \equiv \hat{H}^t$. Similarly, we can define the actions chosen in the continuation game from any period t as $C^t \equiv \prod_{\tau=t+1}^T \prod_{i=1}^N X_\tau^i$ with the corresponding partial order; the realized path is then $C^t \in C^t$. Players perfectly observe past play, and so after any history H^t , they form some belief over players' types from the set $\mathcal{M}_t^i \equiv \Delta(\Theta_{-i})$. Players are endowed with prior beliefs restricted to Θ_{-i} as given by the prior distribution of types f , conditional on observing their private information, namely θ_i . Let \mathcal{Y} be the Borel σ -algebra of measurable subsets $Y \subset \Theta$. We denote conditional beliefs in each period for each player i by $\mu_t^i \in \mathcal{M}_t^i$, such that

$$\mu_t^i : \mathcal{Y} \times \mathcal{H}^t \times \Theta_i \rightarrow [0, 1]$$

$$(Y, H^t, \theta_i) \rightarrow \mu_t^i(Y|H^t, \theta_i)$$

We let $\mu_t(Y|H^t, \theta) = (\mu_t^1, \dots, \mu_t^N)$ and $\mu(Y|H, \theta) = (\mu_1, \dots, \mu_T)$.

We now define behavioral strategies for player i . Define the conditional probability that player

²Throughout this paper, the script i means that the variable in question refers to player i , while the script $-i$ means that the variable refers to all players other than i . If there is no such script, then the variable can be taken to refer to all players. Similarly, if there is no script for the period t , the variable will refer to the vector over all periods.

³Note that we have defined the action sets at each t to be history-independent. However, this is without loss of generality since one can always define the size set of actions to be the maximum over all possible histories, and then define the payoffs at the extraneous actions to be very low in order to ensure that they are never chosen in equilibrium.

i chooses $x_t^i \in X_t^i$ by $\rho_t^i(x_t^i|H^t, \theta_i) \in \Delta(X_t^i)$. We restrict our attention to such probabilities that are measurable with respect to Θ_i . This induces a *strategy correspondence* $\mathbf{x}_t^i : \mathcal{H}^t \times \Theta_i \rightarrow X_t^i$ represents the actions chosen with positive probability out of X_t^i . Note that this is not inherently a best reply, as the definition of a strategy merely states what the player chooses, not whether it maximizes his payoff.

Player i 's ex-post payoff is given by the function $u_i : X \times \Theta \rightarrow \mathbb{R}$. Assume that u_i is bounded and continuous in X and Θ . The interim payoff is defined as follows. For any belief μ_T^i , the interim payoff for player i in period T (i.e. the last period of the game) from choosing x_T^i is given by

$$U_i^T(H^T, x_T^i, \theta_i) \equiv \int \sum_{x_T^{-i}} u_i(H^T, x_T^i, x_T^{-i}, \theta_i, \theta_{-i}) \prod_{j \neq i} \rho_T^j(x_T^j|H^T, \theta_j) d\mu_T^i(\theta_{-i}|H^T, \theta_i)$$

Inductively, the interim payoffs for earlier periods given any μ_t^i from choosing action x_t^i is given by

$$U_i^t(H^t, x_t^i, \theta_i) \equiv \int \sum_{x_{t+1}^i} \sum_{x_t^{-i}} U_i^{t+1}(H^{t+1}, x_{t+1}^i, \theta_i) \rho_t^i(x_{t+1}^i|H^{t+1}, \theta_i) \prod_{j \neq i} \rho_t^j(x_t^j|H^t, \theta_j) d\mu_t^i(\theta_{-i}|H^t, \theta_i)$$

where $H^{t+1} = \{H^t, x_t^i, x_t^{-i}\}$. The objective of player i in each period is to maximize $U_i^t(H^t, \cdot, \theta_i)$ with respect to x_t^i . To indicate the set of actions that maximize U_i^t (but that are not necessarily chosen), we define the *best-reply correspondence* $BR_t^i : H^t \times \Theta_i \rightarrow X_t^i$ as the subset of actions such that, given μ_t^i and θ_i , $U_i^t(H^t, x_t^i, \theta_i) \geq U_i^t(H^t, \hat{x}_t^i, \theta_i)$, $\forall \hat{x}_t^i \in X_t^i$.

We must also define what we mean when we say a strategy is ‘‘monotonic.’’ We say that $\mathbf{x}_t^i(\cdot, \cdot)$ *monotonic in pure strategies within/across subgames* (respectively) if $\mathbf{x}_t^i(H^t, \theta_i)$ is a singleton and

$$\hat{\theta}_i > \theta_i \implies \mathbf{x}_t^i(H^t, \hat{\theta}_i) \geq \mathbf{x}_t^i(H^t, \theta_i)$$

$$\hat{H}^t \geq H^t \implies \mathbf{x}_t^i(\hat{H}^t, \theta_i) \geq \mathbf{x}_t^i(H^t, \theta_i)$$

While our results focus on pure strategies, it will be useful to define monotone mixed strategies for the purposes of the proofs.⁴ $\mathbf{x}_t^i(\cdot, \cdot)$ is *monotonic in mixed strategies within subgames* if⁵

$$\hat{\theta}_i > \theta_i \implies \inf\{x_t^i \in \mathbf{x}_t^i(H^t, \hat{\theta}_i)\} \geq \sup\{x_t^i \in \mathbf{x}_t^i(H^t, \theta_i)\}$$

$\mathbf{x}_t^i(\cdot, \cdot)$ is *monotonic in mixed strategies across subgames* if the induced distribution of play over $x_t^i \in \mathbf{x}_t^i(\hat{H}^t, \theta_i)$, given by ρ_t^i , first-order stochastically dominates (FOSD) that over

⁴As we shall see, the main results do not guarantee pure strategies off-path.

⁵Note that this is stronger than the strong-set order.

$\mathbf{x}_t^i(H^t, \theta_i)$ for $\hat{H}^t \geq H^t$. For the rest of this paper, the term “monotone” will be assumed to refer to pure strategies unless otherwise specified.

We now turn to our equilibrium concept. As the definition of perfect Bayesian equilibrium can be elusive,⁶ we define precisely what we mean by this. We first define what restrictions on beliefs must hold at each subgame. As is standard, Bayes’ rule will be used to generate the conditional distributions for any subgame that is reached with positive probability from a given strategy profile. Thus, for any measurable $\Psi_{-i} \subset \Theta_{-i}$

$$\mu_t^i(\Psi_{-i}|H^t, \theta_i) = \frac{\int_{\Psi_{-i}} \prod_{j \neq i} \rho_{t-1}^j(x_{t-1}^j|H^{t-1}, \theta_j) d\mu_t^i(\theta_{-i}|H^{t-1}, \theta_i)}{\int_{\Theta_{-i}} \prod_{j \neq i} \rho_{t-1}^j(x_{t-1}^j|H^{t-1}, \theta_j) d\mu_t^i(\theta_{-i}|H^{t-1}, \theta_i)}$$

where $H^t = \{H^{t-1}, x_{t-1}^i, x_{t-1}^{-i}\}$. One can also, by Bayes’ Theorem, look at the conditional distribution of types given the history of play, $F(\theta|H^t)$; this can be viewed as an “objective” distribution over types as would be seen by an outside observer who can directly see only the past histories, but not the types of the players. We can in turn condition this distribution on types θ_{-i} to generate $F_i(\theta_i|H^t, \theta_{-i})$, as well as on θ_i to generate $F_{-i}(\theta_{-i}|H^t, \theta_i)$.

To extend this to off-path histories, suppose that the conditional distribution of types at H^{t-1} is $F(\theta|H^{t-1})$. Player i then has some conditional belief $\mu_{t-1}^i(\theta_{-i}|H^{t-1}, \theta_i)$. Suppose that player j deviates to some off-path action x_{t-1}^j in period $t-1$. As one cannot use Bayes’ rule at probability-0 events, the beliefs are unconstrained by that criterion. However, intuitively, beliefs at such events should be reasonable, in that they could arise from potential strategies of the players. Thus the following properties must be satisfied:

(a) Beliefs on any player’s type must be within the support of types of that player in period $t-1$. That is, the support of $F_j(\cdot|\{H^{t-1}, x_{t-1}^j\})$ must be a subset of that of $F_j(\cdot|H^{t-1})$.

(b) For any subset of players \mathcal{I} , the distribution over the vector of the other players’ types, $\theta_{-\mathcal{I}}$, is independent of the actions taken by players $i \in \mathcal{I}$, holding their types fixed:

$$F_{-\mathcal{I}}(\theta_{-\mathcal{I}}|\{H^{t-1}, \{x_{t-1}^i\}_{i \in \mathcal{I}}\}, \{\theta_i\}_{i \in \mathcal{I}}) = F_{-\mathcal{I}}(\theta_{-\mathcal{I}}|H^{t-1}, \{\theta_i\}_{i \in \mathcal{I}})$$

(c) Conditional on type, the actions chosen by all players must be done so independently. Thus for $\Psi \subset \Theta_{-\{i,j\}}$, if we interpret $\frac{\partial \mu_t^i}{\partial \theta_j}$ as the density of μ_t^i with respect to θ_j (if it exists), then for any x_{t-1}, \hat{x}_{t-1} ,

$$\frac{\mu_t^i(\theta_j, \Psi|\{H^{t-1}, x_{t-1}^{-j}, x_{t-1}^j\}, \theta_i)}{\mu_t^i(\theta_j, \Psi|\{H^{t-1}, x_{t-1}^{-j}, \hat{x}_{t-1}^j\}, \theta_i)} = \frac{\mu_t^i(\theta_j, \Psi|\{H^{t-1}, \hat{x}_{t-1}^{-j}, x_{t-1}^j\}, \theta_i)}{\mu_t^i(\theta_j, \Psi|\{H^{t-1}, \hat{x}_{t-1}^{-j}, \hat{x}_{t-1}^j\}, \theta_i)}$$

⁶Although an attempt at a definition exists in Fudenberg and Tirole (1991), their definition has been critiqued in papers such as Battigalli (1996) and Kohlberg and Reny (1997).

whenever these ratios are well-defined; if positive probability is put on some θ_j at H^t , then we replace $\frac{\partial \mu_t^i}{\partial \theta_j}$ with μ_t^i . In either case, the interpretation is that the relative probability of x_{t-1}^j being chosen to \hat{x}_{t-1}^j must be the same for θ_j , regardless of what other players do. The last two conditions are analogous to the “no signalling what you don’t know” condition of Fudenberg and Tirole (1991), except with a continuum of types that may not be independent. A *perfect Bayesian equilibrium* (henceforth PBE), then, is a vector $(\mathbf{x}(\cdot, \cdot), \rho((\cdot|\cdot, \cdot), \mu))$ in which beliefs satisfy Bayes’ rule on-path, properties (a)-(c) hold at all subgames, and at each H^t , the continuation strategies form a Bayesian Nash equilibrium given beliefs μ_t , i.e. if $\hat{x}_t^i \in \mathbf{x}_t^i(H^t, \theta_i) \subset X_t^i$, then $\hat{x}_t^i \in \arg \max_{x_t^i \in X_t^i} U_t^i(H^t, \cdot, \theta_i)$.

3 Existence Theorem

3.1 Main results

The approach we use in this paper proceeds in two steps. In this section, we restrict attention to the case where all players use monotone (mixed) strategies within subgames, and best-replies to such strategies are increasing in the strong-set order in each player’s own type; the latter condition ensures that some monotone strategy is optimal. We provide single-crossing conditions under which players actually wish to take such strategies in the following section. By checking the primitives of the model for single-crossing, one can then conclude by our results that there exists a monotone PBE.

We present the main theorems of the section here, and then outline a sketch of the proof, the formal details of which can be found in Appendix A.

Theorem 3.1: *Under the conditions of the model, if best-replies to monotone strategies are increasing in the strong-set order, then there exists a PBE in monotone mixed strategies within subgames, and in monotone pure strategies within subgames at any on-path history.*

Theorem 3.1 guarantees that, with probability 1 along the actual path of play, players use monotone pure strategies within subgames. Off-path, strategies will still be monotone, with the possibility of mixing in the sense defined above.

An additional appealing feature of the construction is that we are able to pin down the precise beliefs that players have at every subgame, even those that are off-path.

Theorem 3.2: *Suppose that x_τ^j is off-path according to player j ’s strategy at H^τ . Then at any subsequent subgame H^t ,*

$$\mu_t^i(\sup\{\theta_j : \mathbf{x}_\tau^j(H^\tau, \theta_j) < x_\tau^j\}, \Theta_{-\{i,j\}}|H^t, \theta_i) = 1 \tag{1}$$

A point of interest regarding the beliefs as found in Theorem 3.2 is that, for a given H^τ , if $\hat{x}_\tau^j > x_\tau^j$, then the induced beliefs over types conditional on $\{H^\tau, \{\hat{x}_\tau^j, x_\tau^{-j}\}\}$ must be (weakly) greater than those induced conditional on $\{H^\tau, \{x_\tau^j, x_\tau^{-j}\}\}$ in the sense that the lowest type θ_j in the support of the beliefs in the former must be greater than the highest θ_j in the latter. This corresponds to what one might intuitively anticipate in a monotone equilibrium: actions and types are complements. Therefore, higher types are more likely to take higher actions, even when comparing the support of types for actions that are off-path and hence unexpected. This will in turn aid in establishing the optimality of monotone best replies in many cases, as we shall see in Section 4.

We now turn to the proof of the main theorems. The general approach will be to convert the game to an appropriate static Bayesian game, but the key lies in precisely stating how that game is constructed, so that an equilibrium in this static auxiliary game translates into a perfect Bayesian equilibrium of the original game. A crucial difficulty is that, when attempting to prove existence of equilibrium in games with multiple periods, one must address the issue of beliefs, which will endogenously determine the distribution of types and the subsequent actions at all subgames. As highlighted in the following subsection, a naive agent-normal form approach runs into topological issues which can lead to discontinuities of payoffs in other players' strategies. Since continuity is essential for guaranteeing that payoffs behave well enough to find a fixed point over players' strategies, this would potentially preclude using existence results for static Bayesian game. Thus we must come up with an alternative construction that preserves continuity and so guarantees that, in any game in which the action set is finite, if the set of best-replies to monotone strategies is increasing in the strong set order in own type, a monotone equilibrium will exist.

We divide this section as follows. First, we provide a heuristic approach to the proof, showing where the potential problems arise and how the construction in this paper circumvents those issues. We provide a more formal construction and proof in Appendix A. Lastly, we will provide extensions to symmetric games and infinite-horizon games, and discuss extensions to a continuum of actions as well as multidimensional actions and types.

3.2 Heuristics of proof

In this subsection, we restrict our attention to a scenario with two players, two periods, and no more than three actions for each player in each period. This will suffice to illustrate the main issues and the approach that is used in this paper, which is described more rigorously for the general case in Appendix A.




Figure 1: Inference over types for on-path actions

As alluded to above, in order to show that equilibrium exists, one needs to ensure that players' payoffs are continuous in their opponents' strategies, so that the set of best-replies will then be upper-hemicontinuous in their opponents' strategies. This allows for the application of known fixed-point theorems to the best-reply correspondences, which gives equilibrium existence. In a dynamic environment, one has to check for continuity of players' second-period beliefs with respect to their opponents' strategies. Specifically, one must be concerned about the beliefs one has in the second period both with respect to their opponents' types, as well as their opponents' continuation strategies.

We first address the issue of continuity of beliefs over types. The key observation here is that a monotone strategy partitions each player's types into intervals, the conditional distribution over whom is the prior restricted to the product of such intervals. We illustrate this in the following example.

Example 3.1: Consider a scenario in which both players 1 and 2 can choose between three potential actions in period 1. If we place their types on the x - and y -axes, respectively, then any monotone strategy must look like those displayed in Figure 1(a) and 1(b), respectively. So, if one observes that player 1 chooses $x_1^1 = 2$, then one can conclude that θ_1 is between




Figure 2: Continuity of beliefs for off-path actions

the middle two vertical lines in Figure 1(a). Similarly, if one observe that player 2 chooses $x_1^2 = 1$, then one can conclude that θ_2 is below the bottom horizontal line in Figure 1(b). If one observes these two choices in tandem, then one concludes that the joint vector of player 1 and 2's types is in the shaded box in Figure 1(c). Moreover, one would then infer that the conditional distribution of types is the prior restricted to this box. \square

It is easy to show that the beliefs that players have about their opponent's types are continuous in their opponent's strategy for any action that is on-path according to a given strategy. This is seen by perturbing the strategies that the players take around a given one. Notice that any such perturbation is equivalent to a perturbation of the endpoints of the intervals of types that take given actions. Since the players' types are distributed according to a continuous density function, the beliefs over types that player 2 has conditional on perturbing the endpoints of an interval of types of player 1 will not change much.

We now turn to the beliefs that players may have if they observe an unexpected action in period 1, i.e. one that was anticipated to occur with probability 0 given the strategy of their opponent. This is unrestricted by Bayes' rule, since one cannot condition probability on a probability-0 event. However, it turns out that if beliefs are to be continuous in the strategies of one's opponent in period 1, then one can use a perturbation argument (in the spirit of Kreps and Wilson (1982b)) to pin down the beliefs based on the strategy of one's opponent. We illustrate the intuition for this in the following example.

Example 3.2: Suppose that player 1 has $\theta_1 \in [0, 1]$, and chooses a strategy in period 1 such that all types $\theta_1 < 0.5$ take action $x_1^1 = 1$, and all types $\theta_1 > 0.5$ take action $x_1^1 = 3$. Notice that according to this strategy, action $x_1^1 = 2$ is a probability-0 event (Figure 2(a)). In this scenario, what should player 2 believe about player 1's type if he sees $x_1^1 = 2$?

To answer this, we perturb player 1's strategy so that a small interval of types now chooses $x_1^1 = 2$ (Figure 2(b)). Then $x_1^1 = 2$ is no longer off-path, and so player 2 would correctly infer that player 1's type is close to $\theta_1 = 0.5$. Recall that any perturbation of strategies reduces to a perturbation of the endpoints of the interval of types, so any perturbation of player 1's




Figure 3: Discontinuity of continuation strategies in agent-normal form

strategy must only include types near $\theta_1 = 0.5$. Taking the limit of such strategies, if beliefs of player 2 are to be continuous in the strategy of player 1, then in the case where $x_1^1 = 2$ is off-path, player 2 must believe with probability 1 that player 1's type is exactly equal to $\theta_1 = 0.5$. \square

We have now illustrated that one can ensure that beliefs over types in period 2 are continuous in the strategies that players choose in period 1. At this point, one may think one is close to being done, by reducing the game to its agent-normal form and using backward induction from the strategies chosen in period 2 to see which strategies are optimal in period 1. However, such an approach runs into a continuity problem: the beliefs of player 2 in period 2 about what player 1 will do in period 2 may not be continuous in the strategy that player 1 takes in period 1. This is illustrated in the following example.

Example 3.3: Suppose that player 1, with types between 0 and 1, moves in both periods 1 and 2. Consider the strategy given in Figure 3, where types between 0 and 0.5 take action 1 in period 1, and types between 0.5 and 1 take action 3. Meanwhile, in period 2, player 1 chooses the same strategy no matter what happened in period 1: types $\theta_1 < 0.5$ take action $x_2^1 = 1$, while $\theta > 0.5$ takes action $x_2^1 = 2$.

Consider a subgame where player 1 has chosen $x_1^1 = 2$. By the same logic as in Example 3.2, player 2 must believe that player 1's type is exactly $\theta_1 = 0.5$. The question is now this: if beliefs are to be continuous in player 1's first-period strategy, what should player 2 think that player 1 will do in period 2?

Suppose that we perturb the strategy that player 1 takes in period 1, so a small interval of types $\theta_1 < 0.5$ now chooses $x_1^1 = 2$. Then player 2 correctly infers that $\theta_1 < 0.5$. Since these types take action $x_2^1 = 1$ in period 2, player 2 will believe that action 1 will occur with probability 1. Taking the limits of such perturbations in player 1's first-period strategy should lead player 2 to believe that, in the case where player 1 was observed to have deviated to 2 in period 1, then player 1 will play 1 in period 2.

Now let us look at a different perturbation, where a small interval of types $\theta_1 > 0.5$ now chooses $x_1^1 = 2$. By the same reasoning as in the previous paragraph, player 2 must believe that player 1 will choose $x_2^1 = 2$ with probability 1, and so must believe this as well in the limit case where $x_1^1 = 2$ was off-path. But this contradicts the result of the previous paragraph, where we deduced that player 2 must believe that player 1 will choose $x_2^1 = 1$ with probability 1. Thus the belief that player 2 will have in period 2 about what player 1 will do in period 2 is discontinuous in the strategy that player 1 chooses in period 1. \square

As shown in the previous example, the agent-normal form runs into continuity problems. To circumvent these issues, we must use an alternative construction, which works as follows. As in the agent-normal form, we consider each player at each subgame as separate players; however, this is where the similarity in the construction ends. As seen in Examples 3.1 and 3.2, the set of types present at any subgame is a subinterval of the original interval of types, and the distribution is the prior restricted to the product of intervals (in Example 3.2, this interval could be degenerate). The crucial insight is this: one can always rescale this interval to the interval $[0, 1]$, so that at either period, each player's type is some α_i in that interval; we then keep track of what this α_i "means" in terms of real types. So, if, as in Figure 2, types $\theta_1 < 0.5$ take action $x_1^1 = 1$, then in the subgame following that choice of action in period, 1, one interprets $\alpha_1 = 0$ as $\theta_1 = 0$, and $\alpha_1 = 1$ as $\theta_1 = 0.5$.

Instead of optimizing over the current action, taking as given what happens in the future as in the agent-normal form, players indexed at each subgame now optimize over *continuation* strategies. Thus in period 1, players 1 and 2 see their translated type, which is just a rescaled value of θ_i to the unit interval; they then optimize over what they will do in period 1, as well as what they would do at each subgame in period 2 conditional on reaching that subgame. In addition, the versions of players 1 and 2 indexed at each subgame in period 2 optimize given their type α_i , taking into account what true types θ_i their value of α_i corresponds to, from their conjectures of the strategies chosen in period 1.⁷

This construction solves the continuity issue in period 2 that we found in Example 3.3. Returning to the case described in Figure 3, suppose that player 1 indexed at the subgame in period 2 where $x_1^1 = 2$ was chosen in period 1, follows the strategy

$$\tilde{x}_{2,2}^1(H^2, \alpha_1) = \begin{cases} 1, & \alpha_1 < 0.3 \\ 2, & \alpha_1 > 0.3 \end{cases}$$

⁷Recall that the version of player i in period 1 is considered for the purposes of the construction a different player from the version in period 2, and so the players are optimizing with respect to the conjecture of what their earlier self did. Obviously, as in any Bayes Nash equilibrium, this conjecture must be correct in equilibrium.

Let us examine what happens now when we perturb the strategy of player 1 in period 1. If we have a small interval of types (of size ϵ) below $\theta_1 = 0.5$ choose $x_1^1 = 2$, then in period 2, $\alpha_1 = 0$ corresponds to $\theta_1 = 0.5 - \epsilon$, and $\alpha_1 = 1$ corresponds to $\theta_1 = 0.5$. Translating back the types from α_1 to θ_1 , we find that $\theta_1 \in [0.5 - \epsilon, 0.5 - 0.7\epsilon)$ choose $x_2^1 = 1$, and types $\theta_1 \in (0.5 - 0.7\epsilon, 0.5]$ choose $x_2^1 = 2$. Similarly, if we perturb the strategy above $\theta_1 = 0.5$, then in period 2, $\alpha_1 = 0$ corresponds to $\theta_1 = 0.5$, and $\alpha_1 = 1$ corresponds to $\theta_1 = 0.5 + \epsilon$. In the limit, both $\alpha_1 = 0$ and $\alpha_1 = 1$ correspond to $\theta_1 = 0.5$. In any case, the range of types doesn't vary much; since payoffs are continuous in types, and the density function over types is continuous, what player 2 expects in period 2 does not change much from a perturbation of player 1's strategy in period 1, as player 1 is choosing a strategy as a function of α_1 in period 2, not θ_1 .

There is still a potential issue with the perfection of the strategies of each player. Since players are indexed at each subgame, and using continuation strategies, there is a version of each player planning to do something in period 2 as of period 1, and another version of the same player who actually chooses an action in period 2. Since these two versions of the players indexed at periods 1 and 2, respectively, are being treated as separate players, what player 1 plans to do as of period 1 may not be what his future self actually chooses to do in period 2. If this occurs at a history that is not reached given player 2's first-period strategy, then player 1 is indifferent as of period 1. However, this planned action may deter player 2 from choosing an action to get to this history, while the version of player 1 at this subgame would take a different choice. In the same way as in a complete information game, player 1's strategy is then a non-credible threat. As this is incompatible with sequential rationality, we must rule this out in any PBE construction.

We get around the perfection issue by having players choose continuation strategies that are best-replies to "perfected" versions of their opponent's strategy in period 1, which we will illustrate momentarily in an example. This "perfected" version of the strategy will be a mechanical construction from the choices of their opponent's continuation strategies from both periods, differing from the agent-normal form. We then find an equilibrium over continuation strategies as best-replies to the perfected versions of players' opponent's continuation strategies (which, again, is a function of their continuation strategies). We then show that in this equilibrium, one can replace the continuation strategies with the perfected continuation strategies, and still maintain an equilibrium in this Bayesian game; since the strategy is now perfect, it will form a perfect Bayesian equilibrium.

Example 3.4: To illustrate what the perfected continuation strategy of player 1 (which we




Figure 4: Construction of perfected continuation strategy

denote by $\hat{\mathbf{x}}_{\tau,t}^1$ looks like, consider Figure 4. For this construction, $\alpha_1 \in [0, 1]$ is the rescaled type for player 1 from the perspective of period 1, while $\hat{\alpha}_1 \in [0, 1]$ is the rescaled type from the perspective of period 2.

In period 1, there is no issue of perfection from the perspective of period 1, since one is actually present to execute one's strategy; there is no future self that may choose something else. Thus the chosen continuation strategy, $\tilde{\mathbf{x}}_{1,1}^1$, is the same as the perfected continuation strategy, $\hat{\mathbf{x}}_{1,1}^1$ as given by Figure 4(a):

$$\hat{\mathbf{x}}_{1,1}^1(H^1, \alpha_1) = \tilde{\mathbf{x}}_{1,1}^1(H^1, \alpha_1) = \begin{cases} 1, & \alpha_1 \in [0, 0.4) \\ 2, & \alpha_1 \in [0.4, 0.6) \\ 3, & \alpha_1 \in [0.6, 1] \end{cases}$$

Consider the subgame H^2 can only be reached by a choice of $x_1^1 = 1$ in period 1. From the perspective of period 1, we are given that if his current type α_1 were to reach H^2 , then (as seen in Figure 4(b)) he would plan to choose

$$\tilde{x}_{2,1}^1(H^2, \alpha_1) = \begin{cases} 1, & \alpha_1 \in [0, 0.2) \\ 2, & \alpha_1 \in [0.2, 0.8) \\ 3, & \alpha_1 \in [0.8, 1] \end{cases}$$

Notice that the only types of player 1 that reach this subgame are those $\alpha_1 \in [0, 0.4]$, since they are the only ones to choose $x_1^1 = 1$. Thus the only relevant portion of the strategy from the perspective of period 1 is $\tilde{\mathbf{x}}_{1,1}^1(H^2, \cdot)$ is really for $\alpha_1 \in [0, 0.4]$.

We do not yet know that $\tilde{\mathbf{x}}_{2,1}^1$ (what he plans to do as of period 1) and $\tilde{\mathbf{x}}_{2,2}^1$ (what he will actually do in period 2) are consistent, as the versions of player 1 indexed at periods 1 and 2 are treated as separate players. It may be that the version of player 1 in period 2, with rescaled type $\hat{\alpha}_1$, will actually choose the strategy in Figure 4(c):

$$\tilde{\mathbf{x}}_{2,2}^1(H^2, \hat{\alpha}_1) = \begin{cases} 2, & \hat{\alpha}_1 \in [0, 0.5) \\ 3, & \hat{\alpha}_1 \in [0.5, 1] \end{cases}$$

Notice that, given the strategy that player 1 chooses in period 1, type $\hat{\alpha}_1 = 0$ in period 2 corresponds to $\alpha_1 = 0$ in period 1, and type $\hat{\alpha}_1 = 1$ in period 2 corresponds to $\alpha_1 = 0.4$ in period 1, as these are the types who choose $x_1^1 = 1$ in period 1, as seen in Figure 4(a). Thus to construct a “perfect” version of player 1’s strategy, we replace the portion of player 1’s plans from period 1 for this subgame in period 2 with what those types actually do. So, we “shrink” the strategy in Figure 4(c) to the interval $[0, 0.4]$, yielding the modified, “perfected” continuation strategy from the perspective of period 1 given in Figure 4(d):

$$\hat{\mathbf{x}}_{2,1}^1(H^2, \alpha_1) = \begin{cases} 2, & \alpha_1 \in [0, 0.2) \cup [0.4, 0.8) \\ 3, & \alpha_1 \in [0.2, 0.4) \cup [0.8, 1] \end{cases}$$

One may be concerned that this is no longer a monotone strategy function. However, the portion that is observed on path from H^1 is only that following a choice of $x_1^1 = 1$, and the set of types $\{\alpha_1\}$ which choose this is $[0, 0.4)$. When we restrict our attention to this set $\{\alpha_1\}$, the perfected continuation strategy is still monotonic in α_1 over the interval $[0, 0.4)$. Hence it is monotonic on path, and so any best-reply by other players to the perfected continuation strategy treats it as if it were a monotone strategy; this implies that there will be a monotone best-reply as before. \square

Once we have modified the payoffs so that players choose optimal (unperfected) continuation strategies (given by $\tilde{\mathbf{x}}^i$) as best-replies to the perfected continuation strategies of their opponents (given by $\hat{\mathbf{x}}^{-i}$), we are able to invoke the existence theorem of Reny (2011) for this static game. It turns out that once we have found this equilibrium, the perfected continuation strategies will also be best-replies from the perspective of period 1. Returning to Figure 4, suppose that the continuation strategies in Figure 4(a)-(c) are equilibrium strategies. Then the range of rescaled types of player 1 in period 2 corresponds to the interval $[0, 0.4]$ of rescaled types in period 1. Since the strategy of player 1 in period 2 (given in Figure 4(c)) is optimal from the perspective of period 2, it will be optimal for the same types θ_1 from the perspective of period 1 contingent on reaching that subgame. Thus the strategy




Figure 5: Monotonization of $\hat{x}_{2,1}^1(H^2, \alpha_1)$

in Figure 4(d) must also be optimal for player 1 from the perspective of period 1, contingent on reaching that subgame in period 2.⁸

Since players are now using perfected strategies $\hat{\mathbf{x}}^i$ as best-replies to perfected strategies of their opponent, $\hat{\mathbf{x}}^{-i}$, this will form a PBE. The only potential remaining concern is that the perfected strategy is non-monotone for types that are off-path, i.e. those that do not actually reach that subgame in period 2. Fortunately, since best-replies are increasing in the strong set order (by assumption in this section), we can find a monotone best-reply for these off-path types. Since these types are off-path, this modification does not affect the other player's payoff, and so preserves the PBE. We illustrate this in the context of Example 3.4.

Example 3.4, Revisited: Recall that as constructed in Figure 4(d), $\hat{x}_{2,1}^1(H^2, \alpha_1)$ is non-monotone: types in the interval $[0.4, 0.8]$ take actions that are lower than those of types in the interval $[0.2, 0.4]$ (Figure 5(a)). However, the former interval is off path, since only the types in the interval $[0, 0.4]$ take the action in period 1 that could lead to this subgame. Therefore, changing the strategy for the types in the interval $[0.4, 0.8]$ will only affect the payoff of player 1. Since best-replies are increasing in the strong-set order, the strategy of choosing $x_2^1 = 3$ is also optimal for these types, since it is optimal for both lower and higher types (Figure 5(b)). Notice that this strategy is now monotone even for these off-path types. \square

Since we have found a PBE in monotone strategies for all types at all subgames, including off-path types, we are done.

⁸Note that with the current construction, for the purposes of player i 's optimization, we cannot simply take as given the strategy that player i actually uses in period 2 before we have solved for equilibrium, since the values of θ_i that correspond to given values of α_i in period 2 depend on the strategy function chosen in period 1. Since player i knows his own type in period 1, he is optimizing in period 1 with respect to his own type, irrespective of what the rest of the strategy for other types of player i is. Thus player i cannot determine what type $\hat{\alpha}_i$ in period 2 will correspond to his current type, and so cannot use backward induction on what he will do in the next period using the perfection of the strategies described in this paragraph. This is, however, possible for his opponent, who takes all versions of player i 's strategies as given when optimizing.

3.3 Extensions

3.3.1 Symmetric games

Our approach can be extended to games which are symmetric. The analysis here follows that of Reny (2011). Consider a subset of players I , and associate with this subset the set of possible permutations of the players, given by $\{\pi(I)\}$, so that player i is permuted to player $\pi(i)$.⁹ We indicate a permutation of the vector of actions x by these players by $(x_{\pi(I)}, x_{-I})$. Let $u(\cdot)$ be the vector of payoffs for all players.

Definition 3.1: For players $i \in \mathcal{I}$, Γ is *symmetric* if the following conditions hold for all $i, j \in \mathcal{I}$:

1. $\Theta_i = \Theta_j$ and the marginals given by F over θ_i and θ_j , respectively, are identical;
2. $X_t^i = X_t^j$ for all t ;
3. Payoffs remain the same from switching the labels of players $i \in \mathcal{I}$ over their actions, types, and payoffs:

$$u(x^{\pi(\mathcal{I})}, x^{-\mathcal{I}}, \theta_{\pi(\mathcal{I})}, \theta_{-\mathcal{I}}) = u_{\pi(\mathcal{I})}(x^{\mathcal{I}}, x^{-\mathcal{I}}, \theta_{\mathcal{I}}, \theta_{-\mathcal{I}})$$

These conditions correspond to the conditions of Theorem 4.5 of Reny (2011), which guarantees the existence of a symmetric monotone equilibrium in symmetric static games which satisfy the conditions of Theorem 3.1 in this paper.

Theorem 3.3: *If players $i \in I$ are symmetric, and the conditions of Theorem 3.1 are satisfied, then there exists a symmetric monotone PBE.*

3.3.2 Infinite-horizon games

We can use our results for finite T to extend the results to infinite T and N with discounted payoffs. To do so, we must construct a metric over strategies, which we do by weighting the metric δ_t^i over strategies in period t by a factor of $(\frac{1}{2})^t$. We take a sequence of truncations of the game to T' periods. Thus we will be able to meaningfully define convergence of the strategies as $T' \rightarrow \infty$.

We require some additional regularity assumptions to ensure that we can extend our results to infinite periods.

⁹Though Reny (2011) considers the scenario where all players are symmetric, the result extends to a subset of players by the same reasoning: namely, if all players in the subset choose the same strategies, then the set of best-replies is symmetric for all players in that subset.

Assumption 3.1: *Continuity at infinity:* for any player i in period t , then for any $\epsilon > 0$, there exists T_ϵ such that for all $t > T_\epsilon$, for any history H^{t+1} , for any θ , and for any continuations C^t, \hat{C}^t , $|u_i(H^{t+1}, C^t, \theta) - u_i(H^{t+1}, \hat{C}^t, \theta)| < \epsilon$.

Assumption 3.2: At any period t , there is a finite number of players N_t who have non-empty action sets in any period $\tau \leq t$.

Assumption 3.1 is very much in the spirit of Fudenberg and Levine (1983), who use the condition of continuity at infinity to show that the subgame-perfect equilibria of infinite-horizon games arise as the limits of ϵ -equilibria of finite-horizon truncations of games that satisfy continuity at infinity. In a similar spirit, we will use this assumption to show that there is a monotone equilibrium in the infinite-horizon game which is the limit of equilibria of the finite-horizon truncations, each of which has a finite number of players by Assumption 3.2. However, we cannot use their result directly, as they only derive their results for games with finitely many players and types.

Theorem 3.4: *Suppose that $T = \infty$, and that the conditions of Theorem 3.1, as well as Assumptions 3.1 and 3.2, are satisfied. Then there exists a PBE that is monotone within subgames. Furthermore, if the game is symmetric, then there exists a symmetric monotone PBE.*

3.3.3 Continuum of actions

Many applications of Bayesian games involve a continuum of actions, and so the literature has often attempted to extend the existence results to such environments as well. The standard approach has been to approximate such environments by games with finite action spaces, and use Helly’s selection theorem (Kolmogorov and Fomin, p. 373) to find that monotone equilibrium is preserved in the continuous-action limit; this approach was used, for instance, by Athey (2001) and McAdams (2003). However, in the dynamic environment, a naive application of this approach leads to “belief entanglement,” analogous to what Myerson and Reny (2015) refer to as “strategic entanglement” in the limit in periods $t \leq T - 1$, which may preclude the generation of equilibrium by this method. Thus the limit of such a sequence of monotone strategies may generate different beliefs from the limit of the beliefs in the sequence of approximation games. In the online Appendix, we illustrate this effect with an example. We show that in some cases, it is possible to circumvent this issue, essentially showing conditions under which such “ties” between types either do not occur in the limit, or do not entangle beliefs; however, the conditions that we have found are very strong. Nevertheless, this will be useful for some applications, such as that presented in Section 5 on signaling games.

3.3.4 Multidimensional types and actions

Several papers in the literature on monotone equilibrium (such as McAdams (2003) and Reny (2011)) show that it is possible to prove existence of equilibrium in a variety of environments with multidimensional types and/or actions. As our main results and applications are to one-dimensional types and actions, our primary setup is for such environments. However, it is natural to ask whether the results could extend to multidimensional environments.

To examine the case of K -dimensional actions, one can interpret such an environment as one in which the same player takes K separate one-dimensional actions in successive periods, which are unobserved by other players until all K such actions have been taken. In an earlier working paper version (Mensch, 2015), we considered environments in which previous actions may be unobservable, and provided conditions under which the continuity of beliefs condition as in Theorem 3.2 would still hold. The case of multidimensional actions, thus reinterpreted, would fall under these conditions, and so such an extension is possible.

On the other hand, it is not possible to use our method to generalize to multidimensional types. A key feature of our proof is that the conditional distribution at each period is the prior restricted to an interval, and so the boundary of the support will be a (unique) point. However, with multidimensional types, the boundary will be more complicated, and so there would be multiple ways to perturb the boundary while maintaining monotonicity of strategies. This would preclude a result like Theorem 3.2, which exploited the boundary conditions to find beliefs that were continuous in the strategies taken.

4 Conditions for Monotone Best-Replies

We have demonstrated earlier that, with finite actions sets and single-dimensional types, one can guarantee the existence of a monotone PBE if there exists a monotone best-reply to monotone strategy profiles by the other players in the current period, as well as monotone strategies by all players in other periods. We now explore sufficient conditions to guarantee the existence of such monotone best-replies. It turns out that monotonicity of best-replies is considerably more difficult to guarantee in a dynamic environment compared to a static one. Nevertheless, while the conditions presented here may seem restrictive at first glance, it turns out that they are satisfied in a wide variety of environments of economic interest. It should be noted that these conditions are not exhaustive; indeed, in Section 5, the conditions of this section will not be satisfied in some of the applications, yet there exist monotone best-replies for each player.

In static games, it is possible to guarantee the existence of monotone equilibrium with very

general conditions on payoffs; for instance, Athey (2001) and McAdams (2003) show that best responses to monotone strategies are monotone themselves if u_i is either supermodular or log-supermodular in (x, θ) , and types are affiliated. Quah and Strulovici (2012) extend these conditions to any preferences that, in addition to single-crossing, satisfy what they refer to as *signed-ratio monotonicity*. Yet in a dynamic environment, the additional imposition of sequential rationality frequently negates the effects of the presence of such complementarities in the ex-post utility function, as seen in the example below.¹⁰

Example 4.1: We show that even when utility functions are supermodular in all arguments, players' best replies need not be monotone. Suppose that θ_1 and θ_2 are independently distributed, where that of θ_1 is uniform over $[0, 2]$, and that of θ_2 is a compound lottery which places probability 0.5 on $\theta_2 = 0$, 0.49 on $\theta_2 = 1$, and with probability 0.01 is distributed uniformly over $[0, 1]$. In period 1, player 1 chooses $x_1^1 \in \{1, 2\}$; in period 2, player 1 chooses $x_2^1 \in \{0.5, 1.5\}$ and player 2 chooses $x_2^2 \in \{1, 2\}$; in period 3, player 3 chooses $x_3^3 \in \{0, 1\}$, and player 2 chooses $x_3^2 \in \{0, 1\}$. The payoff for player 1 is $u_1(x_1, x_2, x_3, \theta) = x_1^1(\theta_1 - 0.5) - (x_2^1 - \theta_1)^2 + 0.1x_3^3(\theta_1)^6$, while for player 2, it is $u_2(x_1, x_2, x_3, \theta) = -(x_2^1 + 0.6 - x_2^2)^2 + x_2^2(\theta_2)^2 - (x_3^2 - \theta_2)^2$, and for player 3, it is $u_3(x_1, x_2, x_3, \theta) = -(x_3^3 - x_3^2)^2$. Note that payoffs are supermodular in (x, θ) . We do not solve fully for an equilibrium, but merely show that a monotone equilibrium cannot exist by contradiction.

Proposition 4.1: *There does not exist a monotone PBE of the game described in Example 4.1.*

The proof of Proposition 4.1 is located in the online Appendix. The intuition for the failure of monotonicity stems from a failure of beliefs to be monotone (in the sense of FOSD) after period 2. If a monotone equilibrium were to exist, it would have to be that the conditional beliefs over θ_2 that ensure after observing $x_1^1 = 2$ and $x_2^2 = 2$ must be lower than upon observing $x_1^1 = 1$ and $x_2^2 = 2$. These lower beliefs lead to lower actions by player 3 in period 3. Since later actions by other players are no longer higher in response to player 1 choosing higher actions in period 1, this in turn reduces the incentive for high types of player 1 so much as to lead them to deviate to a lower action in period 1. These effects on beliefs induce a tradeoff between action-action complementarities and action-type complementarities, precluding a monotone strategy from being a best-reply. We will therefore need stronger conditions than those sufficient in static environments to enforce monotonicity in dynamic ones.

To define complementarities in dynamic Bayesian games, we will make reference to several

¹⁰Echenique (2004) discusses similar failures of sufficient conditions for strategic complementarities in static games to translate into strategic complementarities in extensive-form games without private information, concluding that the set of subgame-perfect Nash equilibria do not form a lattice under many such conditions.

notions of complementarity in utility functions. When we speak of monotone strategies, this connotes complementarities between the players' types and the actions they take. A general version of this property is given in the following definition.

Definition 4.1: A function g satisfies Milgrom and Shannon's (1994) *single-crossing property* (SCP) in (y, s) if, for $y' \geq y$ and $s' \geq s$, $g(y', s) - g(y, s) \geq 0 \implies g(y', s') - g(y, s') \geq 0$

For several cases, we will require stronger notions of complementarity. Supermodularity provides a definition of stronger complementarities between actions, so higher actions grouped together, along with lower actions grouped together, yields higher payoffs than mixing between such actions. Increasing differences provides a stronger notion of complementarities between actions and types, so that the difference in payoffs from choosing a high action instead of a low one is larger for higher types than for lower types.

Definition 4.2: A function g is *supermodular* in y if for any two y, y' , $g(y' \vee y, s) + g(y' \wedge y, s) \geq g(y', s) + g(y, s)$.

Definition 4.3: A function g satisfies *increasing differences* (ID) in (y, s) if for $s' \geq s$ and $y' \geq y$, $g(y', s') - g(y, s') \geq g(y', s) - g(y, s)$.

Lastly, our environment deals with games of incomplete information, as players do not know what their opponents' types are. In order to have single-crossing at the interim stage, players need to be able to aggregate the single-crossing properties from the ex-post utility function under uncertainty. A property that allows one to do so is given by Quah and Strulovici (2012), defined below.

Definition 4.4: Two functions g, h obey *signed-ratio monotonicity* (SRM) if, for $s' \geq s$,

(a) Whenever $g(s) < 0$ and $h(s) > 0$, then $-\frac{g(s)}{h(s)} \geq -\frac{g(s')}{h(s')}$

(b) Whenever $h(s) < 0$ and $g(s) > 0$, then $-\frac{h(s)}{g(s)} \geq -\frac{h(s')}{g(s')}$

Using the above definitions, it will be possible to generate conditions that ensure single-crossing, and therefore existence of monotone PBE. We introduce the following lemma which will be useful for several of our results. When combined in this way, these definitions encompass several better-known formulations of strategic complementarity, such as supermodularity and log-supermodularity.

Lemma 4.1: (a) (Milgrom and Shannon (1994), Theorem 4) u_i satisfies SCP in (x_t^i, θ) if and only if $(x_t^i)^*(\theta_i) \equiv \arg \max_{x_t^i \in X_t^i} u_i^t(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i})$ is increasing in the strong set order (SSO) in θ_i .

(b) (Quah and Strulovici (2012), Theorem 1) If u_i is a function of (x, θ) that satisfies SCP in (x_t^i, θ_i) , then so is $\int u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) d\mu(\theta_{-i}, x_t^{-i}, C^t)$ for some measure μ , if for any

pair of vectors $(\hat{x}_t^{-i}, \hat{C}^t, \hat{\theta}_{-i}), (x_t^{-i}, C^t, \theta_{-i})$, the pair of functions of θ_i given by

$$u_i(H^t, \hat{x}_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) - u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i})$$

and

$$u_i(H^t, \hat{x}_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i}) - u_i(H^t, x_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i})$$

satisfy SRM.

(c) (Quah and Strulovici (2012), Theorem 2) If u_i is a function of (x, θ) that satisfies SCP between x_t^i and all other variables, then $\int u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) d\mu(\theta_{-i}, x_t^{-i}, C^t)$ satisfies SCP in (x_t^i, θ_i) for some measure μ , if for $\hat{x}_t^i > x_t^i$ and $(\hat{x}_t^{-i}, \hat{C}^t, \hat{\theta}_{-i}) > (x_t^{-i}, C^t, \theta_{-i})$, both

(i) Both

$$u_i(H^t, \hat{x}_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) - u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i})$$

and

$$u_i(H^t, \hat{x}_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i}) - u_i(H^t, x_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i})$$

satisfy SRM, and

(ii) In addition, SRM still holds for any pair of functions in (i) whenever we condition on one additional variable (e.g. for θ_j , SRM is satisfied for

$$u_i(H^t, \hat{x}_t^i, x_t^{-i}, C^t, \theta_i, \theta_j, \theta_{-\{i,j\}}) - u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_j, \theta_{-\{i,j\}})$$

and

$$u_i(H^t, \hat{x}_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \theta_j, \hat{\theta}_{-\{i,j\}}) - u_i(H^t, x_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \theta_j, \hat{\theta}_{-\{i,j\}})$$

Note that by the remark following Theorem 2 in Quah and Strulovici (2012), (c) can be extended to environments in which $(x_t^{-i}, C^t, \theta_{-i})$ is affiliated with θ_i , i.e. μ depends on θ_i .

For ease of use in the results in this section, we say that if the conditions of both (a) and either (b) or (c) hold, then u_i satisfies SCP and SRM in (x_t^i, θ_i) .

Recall that in Example 4.1, the failure of monotonicity of best-replies stemmed from a failure of beliefs to be monotone. This led to player 3 taking an action in period 3 that was deleterious to high types of player 1. However, this only manifested itself through an action in period 3, suggesting that if future actions are discounted enough, then this should no longer be an issue. We formalize this idea in the following definition.

Definition 4.5: At period t with history H^t , we say that future play C^τ (where $\tau \geq t$), is *irrelevant* for player i and action x_t^i if, for all continuations $C, C' \in \mathcal{C}^\tau$, $u_i(H^{\tau+1}, C, \theta) =$

$u_i(H^{\tau+1}, C', \theta)$, where $\{H^t, x_t^i\} \subset H^{\tau+1}$.

The condition of future-play irrelevance guarantees that what occurs in the future does not interfere with the single-crossing conditions in a given period conditional on choosing a given action.¹¹ This will aid in the existence of monotone best-responses within subgames.

We are now able to show existence of monotone best-replies in several environments. We assume throughout that all players $-i$ use consistent monotone strategies. Note that for monotone equilibrium to exist, it is not necessary that we use the same guarantor of existence of monotone best-replies for all players; we can mix-and-match as needed.

The first case that we present is that of short-lived players whose payoffs only depend on what happens before they choose an action.

Proposition 4.2 (Short-run players): *Suppose that*

(i) u_i satisfies SCP and SRM in (x_t^i, θ_i) ;

(ii) (a) θ is independently distributed, or (b) when specifically u_i satisfies the conditions of Lemma 4.1(c), θ is affiliated; and

(iii) C^t is irrelevant for player i at all H^t whenever $X_t^i \neq \emptyset$.

Then there exists a best-reply of player i in period t that is monotone within subgames when monotone (mixed) strategies within subgames are used by all other players and in all other periods.

The intuition for Proposition 4.2 is that, due to the condition of future irrelevance, the decision problem faced by any two types $\theta_i, \hat{\theta}_i$ in period t will essentially be static. By Lemma 4.1, we can aggregate the single-crossing property via integration, so the best-replies must be increasing in the strong set order.

We are able to relax the irrelevance of future play in the following proposition.

Proposition 4.3: *Suppose that*

(i) u_i satisfies SCP and SRM in (x_t^i, θ_i) ;

(ii) θ is independently distributed; and

(iii) At all H^t , C^t is irrelevant for player i except following at most a unique choice $(x_t^i)^*$.

Then there exists a best-reply of player i in period t that is monotone within subgames when monotone (mixed) strategies within subgames are used by all other players and in all other periods.

Intuitively, as in Proposition 4.2, the distribution of types and actions for players $-i$ for all

¹¹Note that the future can be irrelevant for some players, but not others. For instance, if a bidder drops out of an English auction, then the future is irrelevant for the bidder who dropped out, but not the bidders who remain.

relevant periods through t must be the same for θ_i and $\hat{\theta}_i$. Moreover, by future irrelevance, one can consider the same future play for both types, since if one ever chooses $x_t^i \neq (x_t^i)^*$, player i is indifferent between all future choices by all players. This simplifies the comparison for the choice between x_t^i and \hat{x}_t^i , and allows for the invocation of standard single-crossing arguments.

Though the conditions of Proposition 4.3 may at first seem arcane, they hold in a large number of environments of economic interest. For example, the conditions hold for any stopping game in which payoffs do not depend on what happens in periods after one stops. Such games include wars of attrition and auctions, as payoffs there only depend on what has happened before the period t^* in which one drops out. More specifically, each period consists of the choice of whether to stay in or stop, which can be set as a choice between $x_t^i = 1$ and $x_t^i = 0$. Future play C^τ for $\tau \geq t$ is irrelevant to player i at all t unless player i chooses $x_t^i = 1$. Moreover, under appropriate complementarity conditions, the choice of some player j to exit in a given period t may make it more appealing for other players i to exit in a given period τ . These situations will be examined in more detail in the applications in Section 5.

It is possible to obtain even stronger results when $T = 2$: one need no longer assume that the future is irrelevant. Moreover, not only will actions be monotone within subgames in equilibrium, but higher actions in period 1 will induce higher actions in period 2. Such equilibria naturally arise in signaling games, in which a higher type sends a higher message to induce a higher response by one or more receivers.

Proposition 4.4 (Two periods): *Suppose that, for $T = 2$,*

- (i) u_i satisfies ID in (x, θ) and is supermodular in x ; and*
- (ii) θ is independently distributed.*

Then there exists a best-reply of player i in period t that is monotone within and across subgames when monotone (mixed) strategies within and across subgames are used by all other players and in all other periods.

The intuition for Proposition 4.4 is that the effective distribution of types and actions is higher in period 2 when higher actions are taken in period 1, and so higher types are more inclined to take higher actions in period 1 when preferences are supermodular. Similarly, in period 2, the higher distribution of types and actions across subgames will induce high actions in higher subgames. Putting this all together, we find that there exist monotone best-replies to monotone strategies within and across subgames.

Remark: Proposition 4.4 allows for even stronger single-crossing conditions to be imposed in equilibrium than in most of our results. It is easy to verify that the set of strategies that

are monotonic best-replies within and across subgames will be non-empty and join-closed.¹² By restricting attention to this subset of best-replies that are monotone in this stronger sense, one can generate equilibria in which strategies are monotone within and across subgames by the same arguments as in Section 3.

With the results of Sections 3 and 4 in hand, we are now able to examine some applications to some economic questions.

5 Applications

5.1 Generalized games of strategic communication

In many economic environments, an agent wishes to convince other players to take a high action by some sort of communication. However, the communication must be credible to be efficacious; otherwise, there may at best only exist “babbling” equilibria, in which all types choose the same strategy in the first period, and so the receivers’ conditional beliefs over the distribution of senders’ types upon receiving the message will be identical to their priors. On the other hand, if the incentives to separate are sufficiently strong (for instance, if low types always want to send low messages, and high types always want to send a high one), a babbling equilibrium may not exist either. Non-babbling equilibria are of special interest, since they allow credible communication to take place.

It is an immediate corollary of Proposition 4.4 to show existence of monotone equilibrium with multiple senders/receivers and a finite number of possible messages when payoffs are supermodular in all actions and satisfy increasing differences between actions and type, and types are independently distributed. For the rest of the section, we focus on the case where there is one receiver who wishes to induce a higher action by multiple receivers. This will allow us to explore incentives of the sender to pool types by sending the same message, or to separate.

The application we present here generalizes prior similar results of Okuno-Fujiwara et al. (1990), Kartik et al. (2007), Van Zandt and Vives (2007), and Kartik (2009) in the following ways:

1. There may be multiple senders/receivers;
2. The preferences of the receivers may be uncertain, so that there will be some uncertainty from the perspective of the sender as to how the beliefs will influence the choice of actions by the receivers;

¹²See Reny (2011), Lemmas A.13 and A.16.

3. We do not need exogenously require full separation or convex loss functions; and
4. We provide weaker single-crossing conditions than those previously found (i.e. supermodularity/increasing differences).

We will thus provide sufficient conditions under which not only a monotone equilibrium exists, but it will also be non-babbling.

Formally, consider a two-period game with N players, in which player 1 (the sender) has type $\theta_1 \in \Theta_1 = [\underline{\theta}_1, \bar{\theta}_1]$ and chooses an action $x_1^1 \in X_1^1 = [\underline{x}_1, \bar{x}_1]$ in period 1, while all other players j (the receivers) have types $\theta_j \in \Theta_j$, and choose actions $x_2^j \in X_2^j$ (where both $\Theta_j \subset \mathbb{R}$ and $X_2^j \subset \mathbb{R}$) in period 2. Types are distributed independently across players. Payoffs for all players are continuous, and for player 1 are given by $u_1(x_1, x_2, \theta) \equiv u_1^1(x_1^1, \theta_1) + u_1^2(x_2, \theta)$, while for players $j \in \{2, \dots, N\}$, they are given by $u_j(x_1, x_2, \theta) \equiv u_j^2(x_2, \theta)$. We assume that u_1^2 is strictly increasing in x_2 . We also assume that u_1^1 satisfies ID in (x_1, θ_1) , and $u_k^2(x_2, \theta)$ is supermodular in x_2 and satisfies ID in (x_2, θ) .¹³ These conditions are weaker than those of Van Zandt and Vives (2007), who make many additional assumptions,¹⁴ as well as those of Kartik et al. (2007) and Kartik (2009), who assume convexity of the loss function.

There are many possible economic interpretations of this application; the reader is directed to Kartik (2009) or Okuno-Fujiwara et al. (1990). One possible interpretation is of a salesman trying to make a “sales pitch” to a diverse group of investors, in order to attempt to convince them to invest as much as possible in their projects. The salesman has private information on the quality of his project, and it is relatively easier to signal that the project is of higher quality if it is indeed of higher quality. If messages and investment are complementary, in the sense that it becomes relatively easier to send a higher message when the project is of higher quality, and investors want to invest more in higher quality projects and if other investors are investing more, then this scenario will fit under this application.

Another possible interpretation is that of an interested party making a recommendation of a candidate for hire to a committee of prospective employers. All things being equal, the recommender wants the candidate to do as well as possible; however, if the recommender exaggerates too much about the quality of the candidate, then it will hurt her prestige among the prospective employers. Therefore, the recommender will be more inclined to write a better recommendation for the candidate if he is of better quality.

Proposition 5.1: *There exists a monotone PBE in the generalized game of strategic communication.*

¹³If this complementarity is sufficiently strong, then babbling equilibria may not exist, as high θ will always want to choose high x_1 , while low θ will want to choose low x_1 , regardless of the posteriors they induce.

¹⁴See Proposition 20 in their paper.

All proofs in this section are presented in the online Appendix.

Now that we have shown existence of monotone PBE, we are able to use this information to derive properties that must be true of any monotone equilibrium in this model. We provide general conditions which prove sufficient for results analogous to those found in earlier works on games with strategic communication, but under much weaker assumptions. We therefore now additionally assume that for some $j \geq 2$ (but not necessarily all), u_j^2 is differentiable in x_2^j and has strictly increasing differences in (x_2^j, θ_1) . Moreover, we assume that (for the same j) $\arg \max u_2^j(x_2^j, \bar{x}_2^{-j}, \theta)$ is in the interior of X_t^j . The idea will be that, if there is pooling at a particular signal, then some type can strictly increase his payoff by sending a slightly higher signal, thereby inducing the players in period 2 to choose a discretely higher action. While this insight is not new in the context of signaling games, it is important in illustrating how our existence results can be applied to infer additional properties of equilibria in dynamic games without assuming the very strong properties in the models in the previous literature. First, we have yet to show that such an equilibrium is not just a babbling equilibrium. There is a potential for more interesting off-path signaling effects as one must then consider off-path beliefs. This may induce additional pooling, as the beliefs conditional on observing some of the actions may be sufficiently adverse so that no types wish to choose them, and instead pool on the same signal.

Lemma 5.1: *In any monotone PBE with beliefs defined as in Theorem 3.2,¹⁵ there can only be pooling in period 1 at \bar{x}_1^1 .*

Note that this property is exactly that which was found in the equilibrium described in Kartik (2009), in which low types separate, and high types pool. However, we have shown the existence of such an equilibrium under much more general conditions, as, among other things, we have not assumed convexity of the loss function.

It is now apparent how to guarantee a non-babbling equilibrium. By ensuring that the lowest type $\underline{\theta}_1$ has an incentive to deviate to some action other than \bar{x}_1^1 when all other types choose \bar{x}_1^1 , it will necessarily follow that there cannot be any equilibrium in which all types choose \bar{x}_1^1 , and so there cannot be a completely pooling (babbling) equilibrium at all.

Proposition 5.2: *There exists a non-babbling monotone PBE of the generalized strategic communication game if*

$$u_1^1(\bar{x}_1^1, \underline{\theta}_1) + \int u_1^2(\mathbf{x}_2(\{\bar{x}_1^1\}, \theta_{-1}), \underline{\theta}_1, \theta_{-1}) f_{-1}(\theta_{-1}) d\theta_{-1}$$

¹⁵That is, beliefs conditional on observing off-path actions place probability 1 on the supremum of types to choose a lower action.

$$< \sup_{x_1^1 \in X_1^1 \setminus \{\bar{x}_1^1\}} u_1^1(x_1^1, \underline{\theta}_1) + \int u_1^2(\mathbf{x}_2(\{x_1^1\}, \theta_{-1}), \underline{\theta}_1, \theta_{-1}) f_{-1}(\theta_{-1}) d\theta_{-1}$$

To illustrate Proposition 5.2, let the greatest BNE action profile in the subgame in period 2 given the prior beliefs $d\mu_2^j(\theta_1, x_1|x_1) = f_1(\theta_1)$ be $\bar{\mathbf{x}}_2^*(\theta_{-1})$, and the smallest BNE action profile in the subgame in period 2 given by the beliefs $\mu_2^j(\underline{\theta}_1, x_1|x_1) = 1$ be $\underline{\mathbf{x}}_2^*(\theta_{-1})$. From Proposition 16 of Van Zandt and Vives (2007), since the former beliefs first-order stochastically dominate those of the latter, it follows that $\bar{\mathbf{x}}_2^*(\theta_{-1}) \geq \underline{\mathbf{x}}_2^*(\theta_{-1})$. This is the starkest possible set of alternatives that $\underline{\theta}_1$ can face. Thus it will be the case that there is a non-babbling monotone equilibrium if $\theta_1 = \underline{\theta}_1$ prefers to choose some $x_1^1 < \bar{x}_1^1$ and be believed to have $\theta_1 = \underline{\theta}_1$ with probability 1 (and so induce $\underline{\mathbf{x}}_2^*(\theta_2)$) rather than choose \bar{x}_1^1 and be believed to have the prior distribution over θ_1 , and thereby induce $\bar{\mathbf{x}}_2^*(\theta_2)$.

Proposition 5.2 thus gives relatively straightforward conditions to check whether non-degenerate strategic communication is possible: all we have to do is check whether the lowest type of sender would prefer to send the highest message and be believed to have type drawn from the prior distribution, or choose some other message and be known to be the lowest type. In the context of investment, in order for a babbling equilibrium to be possible, it must be optimal for every salesman to pretend that the project is of the highest quality, and therefore have their communication be meaningless, rather than be believed to have the lowest possible quality project. If this is not the case, then we know that a non-babbling monotone equilibrium exists, in which higher messages correspond to higher types, and therefore induce higher actions in period 2.

We can strengthen Proposition 5.2 to provide conditions under which there is complete separation. Intuitively, if it is too costly for any type to send the highest possible message \bar{x}_1^1 , then they all must choose some $x_1^1 < \bar{x}_1^1$. Since pooling can only occur at the top, there must be complete separation.

Proposition 5.3: *Suppose that $\arg \max_{x_1^1} U_1^1(x_1^1, \bar{\theta}_1) < \bar{x}_1^1$ regardless of the choice of monotone $\mathbf{x}_2^j(\{x_1^1\}, \theta_j)$ for all $j \geq 2$. Then there exists a monotone equilibrium with complete separation.*

Thus, by our existence result, we are able to provide much more general conditions that guarantee the existence of a fully separating equilibrium, thereby weakening the assumptions necessary for the existence result found in Proposition 20 of Van Zandt and Vives (2007), or in Theorem 1 of Kartik et al. (2007).

5.2 Stopping games

As noted earlier in Section 4, stopping games with strategic complementarities will satisfy the conditions for existence of monotone equilibrium. We now explore the details of this analysis to demonstrate that this is indeed the case.

Consider a game of T periods in which each player chooses between $x_t^i = 1$ and $x_t^i = 0$ in each period. The payoff for choosing $x_1^i = 0$ is normalized to 0, regardless of other players' strategies, so there is free exit. If $x_t^i = 1$ is chosen, then the player stays in, and may choose $x_{t+1}^i \in \{0, 1\}$ in period $t + 1$. Otherwise, player i has exited permanently, and so the game is over for player i . When the game ends, the remaining players are declared the winners.

This can be interpreted as a game with strategic complementarities. Upon exit, the payoff for player i will be the same regardless of whether $x_\tau^i = 1$ or $x_\tau^i = 0$ for $\tau > t$, and so the game will satisfy future irrelevance. So, the only item that remains is to guarantee that the complementarity conditions of Proposition 4.3 hold within each period t .

Proposition 5.4: *Consider any discrete-time stopping game in which types are independent and payoffs satisfy SCP and SRM in (x_t^i, θ_i) . Then there exists a monotone PBE.*

We now present several applications of the previous proposition to auctions. Previous existence results for dynamic auctions can be found in Milgrom and Weber (1982), Maskin (1992), Lizzeri and Persico (2000), Krishna (2003), Dubra, Echenique, and Manelli (2009), and Birulin and Izmalkov (2011). However, these have mostly focused on conditions for efficiency, and as far as the author is aware, no general existence result is available for auctions with asymmetric, interdependent values and (asymmetrically) affiliated types when efficiency is not guaranteed. Previous existence results for wars of attrition include Milgrom and Weber (1985), Fudenberg and Tirole (1986), Amann and Leininger (1996), Krishna and Morgan (1997), Bulow and Klemperer (1999), and Myatt (2005), though most of these focus on symmetric agents and/or two agents.

Proposition 5.5: *Let $v_i(\theta) \geq 0$ be a continuous, weakly increasing function, c_i and b_i be positive functions of t , with c_i and $-b_i$ increasing (at least one strictly so), and*

$$t^* = \arg \max_t \{ \exists i \neq j : x_t^j = 1 = x_t^i \}$$

$$t_i = \arg \max_t \{ x_t^i = 1 \}$$

$$W = \sum_{i=1}^N 1[x_{t^*}^i = 1]$$

The following games have a monotone PBE:

(i) English auctions with affiliated types, where payoffs are given by:

$$u_i(x, \theta) = [v_i(\theta) - c_i(t^*)] \cdot \frac{1[x_{t^*}^i = 1]}{W}$$

(ii) All-pay auctions with independent types, where

$$u_i(x, \theta) = b_i(t^*)v_i(\theta) \cdot \frac{1[x_{t^*}^i = 1]}{W} - c_i(t)$$

(iii) Auctions with costly bidding with independent types,¹⁶ where

$$u_i(x, \theta) = [v_i(\theta) - t^*] \cdot \frac{1[x_{t^*}^i = 1]}{W} - c_i(t_i)$$

Moreover, when players are symmetric, there exists a symmetric monotone PBE.

We also present an application to joint research projects, in which the probability of success is increasing in type and the cost of doing research is decreasing in type and the number of other firms participating. One can model the probability of success as a choice by nature, thereby treating it as a strategy for the purpose of incorporation into equilibrium analysis. Conditional on success, all remaining participating firms share in the benefits from discovery. While the literature mostly considers rival technological discovery, some papers on this subject include Lee and Wilde (1980), Fudenberg and Tirole (1986), Harris and Vickers (1987), and Leininger (1991).

Formally, for any firm that is still participating at time t , let the probability of success at period t be $P(x, \theta, t)$, which is increasing in x and θ . The cost of research in period t is $c_i(\theta_i, t)$, which is decreasing in θ_i . Thus each player's payoff is, if the discovery is made in period t^* ,

$$u_i(x, \theta) = v_i(\theta) \cdot 1[x_{t^*}^i = 1] - \sum_{t=1}^{t^*} c_i(\theta_i, t)$$

while if they drop out at some t^* , their payoff is

$$u_i(x, \theta) = - \sum_{t=1}^{t^*} c_i(\theta_i, t)$$

Lastly, we assume that types are affiliated.

Proposition 5.6: *There exists a monotone PBE in the joint research project game. More-*

¹⁶A common environment in which such payoffs appear is internet auctions, where one must pay a small fee per bid placed, and the bidding proceeds in small increments (often one cent, giving them the nickname "penny auctions").

over, when players are symmetric, there exists a symmetric equilibrium.

6 Conclusion

This paper shows how the results from static games with incomplete information and complementarities between actions and types can be extended to dynamic games to generate a perfect Bayesian equilibrium. While single-crossing is tricky in dynamic environments, as seen in Example 4.1, we have shown that single-crossing still holds in a wide variety of environments of economic interest, and so the existence results presented here have bite for important economic questions. They also illustrate many potential issues that one must address when attempting to show existence of monotone equilibrium.

A natural avenue for extension of our results is to develop other single-crossing conditions for environments not considered in this paper; once single-crossing has been shown, it will be possible to use our results to immediately show that a monotone PBE will exist. Indeed, some of the applications in Section 5 invoked alternative single-crossing conditions from those in Section 4; combined with the existence results of Section 3, one is nonetheless able to state that a monotone PBE exists. There is therefore room for demonstration of such alternative conditions, especially for particular applications.

Similarly, the methodology here for showing continuity of beliefs (the key ingredient in the existence result) may prove useful in other environments where the basic model is slightly different. For instance, in many applications, players do not observe the entire history of actions, but only some statistic derived from them. One such example that has appeared in the literature is that of auctions with resale, in which one might not observe all of the bids placed in the auction stage, but only the price at which the item was sold. Even though the results here do not directly apply, this paper gives a technique for looking at the issues in such problems. Combined with an appropriate single-crossing condition, one could state that if beliefs are continuous, a monotone PBE will exist. It would be of interest to show continuity of beliefs in other environments in future work.

References

- Amann, E., and Leininger, W. (1996): "Asymmetric all-pay auctions with incomplete information: the two-player case." *Games and Economic Behavior*, 14, 1-18.
- Athey, S. (2001): "Single crossing properties and the existence of pure strategy equilibria in games of incomplete information." *Econometrica*, 69, 861-889.

- (2002): “Monotone comparative statics under uncertainty.” *Quarterly Journal of Economics*, 117, 187-223.
- Balbus, L., Reffett, K., and Wozny, L. (2013): “A constructive study of Markov equilibria in stochastic games with strategic complementarities.” *Journal of Economic Theory*, 150, 815-840.
- Battigalli, P. (1996): “Strategic independence and perfect Bayesian equilibria.” *Journal of Economic Theory*, 70, 201-234.
- Birulin, O., and Izmalkov, S. (2011): “On efficiency of the English auction.” *Journal of Economic Theory*, 146, 1398-1417.
- Bulow, J., and Klemperer, P. (1999): “The generalized war of attrition.” *American Economic Review*, 89, 175-189.
- Crawford, V. P., and Sobel, J. (1982): “Strategic information transmission.” *Econometrica*, 50, 1431-1451.
- Curtat, L. O. (1996): “Markov equilibria of stochastic games with complementarities.” *Games and Economic Behavior*, 17, 177-199.
- Dubra, J., Echenique, F., and Manelli, A. M. (2009): “English auctions and the Stolper-Samuelson theorem.” *Journal of Economic Theory*, 144, 825-849.
- Echenique, F. (2004): “Extensive-form games and strategic complementarities.” *Games and Economic Behavior*, 46, 348-364.
- Eilenberg, S., and Montgomery, D. (1946): “Fixed point theorems for multi-valued transformations.” *American Journal of Mathematics*, 68, 214-222.
- Fudenberg, D., and Levine, D. (1983): “Subgame-perfect equilibria of finite-and infinite-horizon games.” *Journal of Economic Theory*, 31, 251-268.
- Fudenberg, D., and Tirole, J. (1986): “A theory of exit in duopoly.” *Econometrica*, 54, 943-960.
- (1991): “Perfect Bayesian equilibrium and sequential equilibrium.” *Journal of Economic Theory*, 53, 236-260.
- Grossman, S. J., and Perry, M. (1986). “Sequential bargaining under asymmetric information.” *Journal of Economic Theory*, 39, 120-154.
- Gul, F., and Sonnenschein, H. (1988): “On delay in bargaining with one-sided uncertainty.” *Econometrica*, 56, 601-611.
- Gul, F., Sonnenschein, H., and Wilson, R. (1986): “Foundations of dynamic monopoly and the Coase conjecture.” *Journal of Economic Theory*, 39, 155-190.
- Hafalir, I., and Krishna, V. (2009): “Revenue and efficiency effects of resale in first-price

- auctions.” *Journal of Mathematical Economics*, 45, 589-602.
- Harris, C., and Vickers, J. (1987): “Racing with uncertainty.” *Review of Economic Studies*, 54, 1-21.
- Kartik, N. (2009): “Strategic communication with lying costs.” *Review of Economic Studies*, 76, 1359-1395.
- Kartik, N., Ottaviani, M., and Squintani, F. (2007): “Credulity, lies, and costly talk.” *Journal of Economic Theory*, 134, 93-116.
- Kohlberg, E., and Reny, P. J. (1997): “Independence on relative probability spaces and consistent assessments in game trees.” *Journal of Economic Theory*, 75, 280-313.
- Kolmogorov, A. N., and Fomin, S. V. (1970): *Introductory Real Analysis*. New York: Dover Publications.
- Kreps, D., and Wilson, R. (1982): “Reputation and imperfect information.” *Journal of Economic Theory*, 27, 253-279.
- (1982): “Sequential equilibria.” *Econometrica*, 50, 863-894.
- Krishna, V. (2003) “Asymmetric English auctions.” *Journal of Economic Theory*, 112, 261-288.
- Krishna, V., and Morgan, J. (1997): “An analysis of the war of attrition and the all-pay auction.” *Journal of Economic Theory*, 72, 343-362.
- Lee, T., and Wilde, L. L. (1980): “Market structure and innovation: a reformulation.” *Quarterly Journal of Economics*, 94, 429-436.
- Leininger, W. (1991): “Patent competition, rent dissipation, and the persistence of monopoly: the role of research budgets.” *Journal of Economic Theory*, 53, 146-172.
- Maskin, E. (1992): “Auctions and Privatization.” In H. Siebert (ed): *Privatization*. Institut für Weltwirtschaften der Universität Kiel, Kiel, 115-136.
- McAdams, D. (2003): “Isotone equilibrium in games of incomplete information.” *Econometrica*, 71, 1191-1214.
- Mensch, J. (2015): “On the existence of monotone pure-strategy perfect Bayesian equilibrium in games with complementarities.” Working paper.
- Milgrom, P. (1981): “Good news and bad news: Representation theorems and applications.” *Bell Journal of Economics*, 12, 380-391.
- Milgrom, P., and Roberts, J. (1990): “Rationalizability, learning, and equilibrium in games with strategic complementarities.” *Econometrica*, 58, 1255-1277.
- Milgrom, P., and Shannon, C. (1994): “Monotone comparative statics.” *Econometrica*, 62, 157-180.

- Milgrom, P. R., and Weber, R. J. (1982): “A theory of auctions and competitive bidding.” *Econometrica*, 52, 1089-1122.
- (1985): “Distributional strategies for games with incomplete information.” *Mathematics of Operations Research*, 10, 619-632.
- Myatt, D. P. (2005): “Instant exit from the asymmetric war of attrition.” Manuscript.
- Myerson, R. B., and Reny, P. J. (2015): “Open sequential equilibria of multi-stage games with infinite sets of types and actions.” Mimeo.
- Okuno-Fujiwara, M., Postlewaite, A., and Suzumura, K. (1990): “Strategic information revelation.” *Review of Economic Studies*, 57, 25-47.
- Quah, J. and Strulovici, B. (2012): “Aggregating the single crossing property.” *Econometrica*, 80, 2333-2348.
- Reny, P. J. (2011): “On the existence of monotone pure strategy equilibria in Bayesian games.” *Econometrica*, 79, 499-553.
- Reny, P.J. and Zamir, S. (2004): “On the existence of pure strategy monotone equilibria in asymmetric first-price auctions.” *Econometrica*, 72, 1105-1125.
- Spence, M. (1973). “Job market signaling.” *Quarterly Journal of Economics*, 87, 355-374.
- Van Zandt, T., and Vives, X. (2007): “Monotone equilibria in Bayesian games of strategic complementarities.” *Journal of Economic Theory*, 134, 339-360.
- Vives, X. (1990): “Nash equilibrium with strategic complementarities.” *Journal of Mathematical Economics*, 19, 305-321.
- (2009): “Strategic complementarity in multi-stage games.” *Economic Theory*, 40, 151-171.
- Zheng, C. Z. (2014): “Existence of monotone equilibria in first-price auctions with resale.” Manuscript.

Appendix A

As mentioned in the introduction, much work has been done to demonstrate existence of monotone equilibrium in static Bayesian games, i.e. where $T = 1$. In particular, Reny (2011) established the following theorem.

Theorem A.1 (Reny, Proposition 4.4): *Suppose that the following four conditions hold:*

- (i) $X \times \Theta$ forms a Euclidean sublattice;
- (ii) F is atomless;
- (iii) Payoffs u_i are bounded, measurable in x and θ , and continuous in x ;¹⁷ and that
- (iv) The set of monotone strategies that are best-replies to monotone strategies by one's opponents is non-empty and join-closed (i.e. if x_i, \hat{x}_i are optimal for type θ_i , then so is $x_i \vee \hat{x}_i$).

Then there exists a monotone pure-strategy equilibrium.

To apply his results, we translate the dynamic game into an appropriate static game. We then verify that the topological conditions as defined in Reny (2011) are satisfied.

In order to reinterpret the game as a static one, will need to break down the players by H^t . Thus we define the auxiliary game Γ^1 in which player i at each H^t is considered a distinct player. Player i at each H^t will then choose not only what he does at H^t , but what he plans to do at all subsequent subgames; thus the action space for player i is $\prod_{\tau=t}^T (X_\tau^i)^{|H^\tau: H^t \subset H^\tau|}$. Note that this is not a reduction of the original dynamic game to its agent-normal form, but rather a description of continuation strategies from a given subgame.

Consider player i 's problem in period t . Our approach will necessitate the description of the type space of player i in period t to be independent of the types that actually appear at H^t , which will depend endogenously on which strategy is chosen in earlier periods. That is, θ_i has some distribution conditional on H^t ; moreover, there will be a joint distribution of θ_{-i} conditional on (H^t, θ_i) .

The restriction that the players' strategies be monotone allows us to further restrict the beliefs that are generated by Bayes' rule. Specifically, the set of types of player i that choose any action x_t^i in a given period will be a subinterval of the set of types in the support at period t . The distribution of θ_i conditional on choosing x_t^i and given θ_{-i} will then just be the prior restricted to this interval. We formalize this in the following lemma.

Definition A.1: The distribution over types θ_i is *completely atomic* if it places probability

¹⁷Reny notes that this condition is solely to ensure that best-replies are upper-hemicontinuous in the strategies of the other players. It is therefore possible to relax this condition as long as this upper-hemicontinuity still holds, which (as we shall show) it will when we look at strategy profiles (under the \mathcal{L}^1 topology) instead of actions taken by every type of each player.

1 on $\theta_i = \theta_i^*$ for some $\theta_i^* \in \Theta_i$.

Lemma A.1: *Suppose that each player i chooses a fixed monotone strategy at each subgame. Then for any H^t that is on-path, the conditional distribution $F(\cdot|H^t)$ is completely atomic over some subset \mathcal{I} of players and absolutely continuous for $i \notin \mathcal{I}$, with conditional density equal to the prior restricted to an interval $[\theta_i^1, \theta_i^2] \subset \Theta_i$.¹⁸*

As noted in Lemma A.1, the conditional distribution of types θ at any subgame H^t are the prior restricted to a product of subintervals. This enables the rescaling of types used in the construction outlined in the heuristic argument in Section 3.

Proof of Lemma A.1: We show this inductively. In period 1, the conditional density is just the prior f , which as given is absolutely continuous. Now suppose that in period $t - 1$ with history H^{t-1} , the distribution is completely atomic or absolutely continuous. In either case, the support of θ_i is an interval (possibly degenerate). Since strategies are monotone, the support of θ_i conditional on choosing x_{t-1}^i at H^{t-1} is a subinterval of the set of types who have chosen actions $(x_1^i, \dots, x_{t-2}^i)$, which again must be an interval (again, possibly degenerate). Thus the conditional distribution at H^t must be either completely atomic over some subset \mathcal{I} of players and is absolutely continuous for all other players. \square

For actions that are off-path, we can assume that players' beliefs place the conditional distribution over the deviating player's type in accordance with Lemma A.1 without loss of generality since beliefs are not otherwise specified from on-path play. Though this may affect the set of potential equilibria, we will show that it is possible to find a perfect Bayesian equilibrium which satisfies this restriction. For the non-deviating players, the interval is defined as before, thereby ensuring that the deviating player does not "signal what he does not know." It is therefore possible to assume that the conditional distribution over θ_i will always be a Cartesian product of intervals. Therefore, by extension, in subsequent periods beliefs will also be generated in the same manner as in Lemma A.1, i.e. the prior restricted to subintervals.

We are now able to construct the transformation of the players' problems in period t to a static one. Formally, assume that all players choose monotone (mixed) strategies within subgames. From the perspective of an outsider, by Lemma A.1, at any H^t , the distribution over the types of all players will have support over a Cartesian product of intervals. Thus, if the conditional support of θ_i is $[\theta_i^1, \theta_i^2]$, then any θ_i can be written as $\alpha_i \theta_i^2 + (1 - \alpha_i) \theta_i^1$ for some choice of $\alpha_i \in [0, 1]$. The interval $[0, 1]$ now serves as the type space for player i indexed at H^t in Γ^1 . We therefore are able to transform the support to an N -dimensional Euclidean unit hypercube, so that each player i has type α_i . To translate types back from $[0, 1]$ to Θ_i ,

¹⁸Note that if $\theta_i^1 = \theta_i^2$, then the distribution is completely atomic, so the former is a special case of the latter.

we define

$$\begin{aligned}\tilde{\theta}_i^t &: \mathcal{H}^t \times [0, 1] \rightarrow \Theta_i \\ (H^t, \alpha_i) &\rightarrow \theta_i\end{aligned}$$

to set the type that satisfies $\theta_i = \alpha_i \theta_i^2 + (1 - \alpha_i) \theta_i^1$. If the support of types given H^t is A , we can therefore express the distribution over θ as a distribution over $\alpha \in [0, 1]^N$; i.e. if the conditional distribution is absolutely continuous with respect to the prior (which itself has full support), then the distribution of $G_t(\alpha|H^t)$ is given by the density function

$$g_t(\alpha|H^t) = \frac{f(\tilde{\theta}^t(H^t, \alpha)) \int_A d\theta}{\int_A f(\theta) d\theta}$$

and so for any two $\alpha, \alpha' \in [0, 1]^N$,

$$\frac{g_t(\alpha'|H^t)}{g_t(\alpha|H^t)} = \frac{f(\tilde{\theta}^t(H^t, \alpha'))}{f(\tilde{\theta}^t(H^t, \alpha))} \quad (2)$$

Recall that the prior density f of θ_{-I} is continuous in θ_I for any subset of players I . Thus the extension to the case where the conditional distribution of θ_j is completely atomic at θ_j^* can be found by taking the limit of equation (2) as the set of θ_j in A converges to the singleton at $\{\theta_j^*\}$; this will just place the uniform distribution over α_j conditional on any values of α_{-j} . Moreover, G_t will be atomless due to the fact that the prior F was also atomless. It will therefore be absolutely continuous with respect to the uniform distribution over the Euclidean unit hypercube, i.e. that given by the Lebesgue measure. As we will see, this will ensure that at any given subgame, one can consider a conditional distribution over types (i.e. the uniform distribution) that is atomless over α and does not vary with the actual path of play.¹⁹

To see how the actions chosen at each subgame translate from the original game to those in Γ^1 , suppose that the support of types θ_i that is believed to occur at H^t is $[\theta_i^1, \theta_i^2]$, so that $\tilde{\theta}_i^t(H^t, 0) = \theta_i^1$ and $\tilde{\theta}_i^t(H^t, 1) = \theta_i^2$. As higher types θ_i choose higher actions in any monotone strategy, it will follow that if $\hat{\alpha}_i > \alpha_i$, then the type θ_i corresponding to $\hat{\alpha}_i$ chooses a (weakly) higher action for all H^τ , where $\tau \geq t$. Since we have broken down each player i according to each H^t , one can represent the strategies of each player i in each period $\tau \geq t$ (conditional on reaching H^τ) as monotone functions of α_i . We define the actions chosen in period $\tau \geq t$ from the perspective of player i in period t (i.e. what he will do if these subgames are reached)

¹⁹This will be important because Reny's theorem applies to atomless type spaces, and so it will be useful to ensure that the translated distribution over α is indeed atomless.

according to this monotone strategy by the function

$$\tilde{\mathbf{x}}_{\tau,t}^i : \mathcal{H}^\tau \times [0, 1] \rightarrow X_\tau^i$$

$$(H^\tau, \alpha_i) \rightarrow x_\tau^i$$

At each H^τ , with players choosing strategies according to $\tilde{\mathbf{x}}_{\tau,t}^i$, we generate a new history $H^{\tau+1} = \{H^\tau, \{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i)\}_i\}$. Thus, inductively, players choose their period $\tau + 1$ action according to their strategy $\tilde{\mathbf{x}}_{\tau+1,t}^i$. Indicating the collection of $\{\tilde{\mathbf{x}}_{\tau,t}^j, \tilde{\theta}_j^t\}_{\tau \geq t, j \neq i}$ by $\{\tilde{\mathbf{x}}_{\tau,t}^{-i}, \tilde{\theta}_{-i}^t\}$, if we were to consider the payoffs (as of now) based on continuation strategies, the expression for the interim payoff of player i conditional on being type α_i and choosing actions $\{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i)\}_{\tau=t}^T$ can now be written as (suppressing arguments for $\tilde{\theta}_i^t$ and $\tilde{\theta}_{-i}^t$)²⁰

$$\int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i), \tilde{\mathbf{x}}_{\tau,t}^{-i}(H^\tau, \alpha_{-i})\}_{\tau=t}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \quad (3)$$

With the payoffs as given in (3), one can show that one's payoffs will be continuous in the strategies of one's opponents. However, as alluded to in the heuristics sections, we now have a perfection issue. Since we are considering player i at each H^t as a separate player, we need to ensure that the actions as described by $\tilde{\mathbf{x}}_{\tau,t}$ are *consistent* with those chosen by $\tilde{\mathbf{x}}_{\tau,\tau}$, so that the action that player i with type θ_i plans to take (as of period t) in period τ if H^τ is reached will be the actual action taken by player i with type θ_i indexed by H^τ .

In order to get around this issue, we will modify the payoffs so that each player will be best-replying to a “perfect” version of their opponents strategies, which we will denote as $\hat{\mathbf{x}}_{\tau,t}^{-i}$, which will be constructed mechanically from their opponents' continuation strategies, $\tilde{\mathbf{x}}_{\tau,t}^{-i}$. This construction will ensure that players are, in a certain sense (to be formalized shortly) that they are responding to what their opponents “actually do” in later periods. We then find an equilibrium among choices of $\tilde{\mathbf{x}}_{\tau,t}^i$ with these modified payoffs, i.e. where players are replying to $\hat{\mathbf{x}}_{\tau,t}^{-i}$. Once we do this, it will turn out that the “perfect” versions of all players' strategies, $\hat{\mathbf{x}}_{\tau,t}^i$, will also be best-replies, and so form a PBE in the game with payoffs given by (3).

To proceed in this manner, we first need to define a notion of reachability.

Definition A.2: H^τ is *reachable* from H^t for some sequence of actions by player i , $\{x_{t'}^i\}_{t'=t}^{\tau-1}$, if $H^\tau = \{H^t, \{x_{t'}^i\}_{t'=t}^{\tau-1}\}$ for some sequence of actions $\{x_{t'}^{-i}\}_{t'=t}^{\tau-1}$.

²⁰To understand this formula, we note that the player's decision in period t , given his information, is to maximize his payoff over his expected payoff over all possible terminal histories. Hence we look at the payoff at each individual history, given by u_i , weighted by the probability given by g_t . The arguments $H^\tau \in \mathcal{H}^\tau$ for each $\tilde{\mathbf{x}}_{\tau,t}^i$ and $\tilde{\mathbf{x}}_{\tau,t}^{-i}$ are simply the histories generated by $\{H^t, \{\tilde{\mathbf{x}}_{t',t}^i, \tilde{\mathbf{x}}_{t',t}^{-i}\}_{t'=t}^{\tau-1}\}$, i.e. by the actions up to period τ .

Thus, if H^τ is reachable from H^t , then for the strategy profiles to be consistent, the joint distributions of actions and types as generated from both $(\tilde{\mathbf{x}}_{\tau,t}^i, \tilde{\theta}_i^t)$ and $(\tilde{\mathbf{x}}_{\tau,\tau}^i, \tilde{\theta}_i^\tau)$ must coincide, i.e. for any H^t and $A \subset \Theta$,

$$\begin{aligned} & \frac{\int \mathbb{1}_{\{(\tilde{\theta}_i^t, \tilde{\theta}_{-i}^t, H^t, \{\tilde{\mathbf{x}}_{t',t}^i, \tilde{\mathbf{x}}_{t',t}^{-i}\}_{t'=t}^{\tau-1}, \tilde{\mathbf{x}}_{\tau,t}^i)\} \in (A, H^\tau, \{x_\tau^i\})}(\alpha) g_t(\alpha | H^t) d\alpha}{\int \mathbb{1}_{\{(\tilde{\theta}_i^t, \tilde{\theta}_{-i}^t, H^t, \{\tilde{\mathbf{x}}_{t',t}^i, \tilde{\mathbf{x}}_{t',t}^{-i}\}_{t'=t}^\tau)\} \in (\Theta, H^\tau)}(\alpha) g_t(\alpha | H^t) d\alpha} \\ &= \int \mathbb{1}_{\{(\tilde{\theta}_i^\tau, \tilde{\theta}_{-i}^\tau, H^\tau, \tilde{\mathbf{x}}_{\tau,\tau}^i)\} \in (A, H^\tau, \{x_\tau^i\})}(\hat{\alpha}) g_\tau(\hat{\alpha} | H^\tau) d\hat{\alpha} \end{aligned} \quad (4)$$

where the term on the left-hand side is the conditional probability that x_τ^i is chosen by $\tilde{\theta}_i^t(H^t, \alpha_i) = \theta_i = \tilde{\theta}_i^\tau(H^\tau, \hat{\alpha}_i)$ after history H^τ according to the planned action as of period $t < \tau$, and the right hand side is the probability that x_τ^i is chosen when H^τ is actually realized.

We now define an alternative auxiliary game Γ^2 , in which type and action spaces are the same as in Γ^1 , but the payoff function differs as follows. Suppose that we fix monotone strategy $(\{\tilde{\mathbf{x}}_{t',t}^j\}_{t'=t}^T, \tilde{\theta}_j^t)$ for all j , and consider the problem from the perspective of player i at history H^t . Suppose that type α_i chooses strategy $\{\tilde{\mathbf{x}}_{t',t}^i\}_{t'=t}^T$. We define the following alternative strategy from the perspective of period t , $\hat{\mathbf{x}}_{\tau,t}^i$, inductively, starting from period t and working forward to period T , and setting $\hat{\mathbf{x}}_{t,t}^i = \tilde{\mathbf{x}}_{t,t}^i$. Suppose that $H^t \subset H^\tau$; there will be a unique vector $\{x_{t'}^i\}_{t'=t}^{\tau-1} \in \prod_{t'=t}^{\tau-1} X_{t'}^i$ which can generate H^τ conditional on $\{H^{t'}\}_{t'=t}^{\tau-1}$ being reached. Let $\underline{\alpha}_i(H^\tau, H^t)$ be the infimum of the set of types α_i such that α_i chooses $\{x_{t'}^i\}_{t'=t}^{\tau-1}$ according to $\{\tilde{\mathbf{x}}_{t',t}^i\}_{t'=t}^{\tau-1}$ at the respective $\{H^{t'}\}_{t'=t}^{\tau-1}$, and $\bar{\alpha}_i(H^\tau, H^t)$ be the supremum of that set of types. Then

$$\hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i) \equiv \begin{cases} \tilde{\mathbf{x}}_{\tau,\tau}^i(H^\tau, \hat{\alpha}_i), & \alpha_i \in (\underline{\alpha}_i(H^\tau, H^t), \bar{\alpha}_i(H^\tau, H^t)) \text{ and } \hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i(H^\tau, H^t)}{\bar{\alpha}_i(H^\tau, H^t) - \underline{\alpha}_i(H^\tau, H^t)} \\ \tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i), & \text{otherwise} \end{cases}$$

Similarly, we define the vector $\{\hat{\mathbf{x}}_{\tau,t}^j\}_{j \neq i}$. We now define the payoff function for Γ^2 so that the interim payoff function given H^t , where H^τ is generated the same way as in (3), remains the same modulo replacing $\tilde{\mathbf{x}}_{\tau,t}^{-i}(H^\tau, \alpha_{-i})$ with $\hat{\mathbf{x}}_{\tau,t}^{-i}(H^\tau, \alpha_{-i})$. Player i 's interim payoff will therefore be given by (suppressing the arguments for $\tilde{\mathbf{x}}_{\tau,t}^i$, $\hat{\mathbf{x}}_{\tau,t}^{-i}$, and $\tilde{\theta}^t$)

$$\int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \quad (5)$$

This is done to ensure that the payoff of player i as indexed by H^t will be based on what players $-i$ will *actually* do when H^τ is realized, for $\tau > t$, rather than what they plan to do as of period t . The goal will be to show that there exists an equilibrium in strategies

$\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ in the static interpretation of the game when the payoffs are given by (5). From this, we construct an Bayesian equilibrium of the original game, in which payoffs are given by (3), which involves consistent strategies, and so will be a PBE. Note that, while $\hat{\mathbf{x}}_{\tau,t}^j$ is not necessarily monotone for all values of α_j , it will be for the relevant values of α_j , i.e. those at which H^τ is reached with positive probability, by construction (as seen in Example 3.4). Hence the set of monotone best-replies of to $\{\hat{\mathbf{x}}_{\tau,t}^{-i}\}$ will still be non-empty and join-closed.

An important feature of $\hat{\mathbf{x}}_{\tau,t}^i$ is that, as defined, it still may not be optimal for player i given a particular profile of $\tilde{\mathbf{x}}_{\tau,t}^i$, as players are choosing their strategies on the basis of their translated type α_i . If a different monotone strategy is used in period t , this could lead to a different interval of types θ_i being present in period τ , and so the translation between θ_i and α_i will change. This could lead to the action chosen by a particular α_i in period τ being suboptimal from the perspective of period t . Since players in period t have their own conjecture of their “true” type θ_i and optimize with respect to this, this precludes the player in period t from using backward induction when optimizing, and so must use continuation strategies. Unsurprisingly, it will turn out in equilibrium that $\hat{\mathbf{x}}_{\tau,t}^i$ will be consistent and optimal, as the way that $\tilde{\theta}_i^\tau$ is determined will ensure that it aligns with $\tilde{\theta}_i^t$ correctly from the strategies chosen in earlier periods. Since $\tilde{\mathbf{x}}_{\tau,\tau}^i$ is optimal from the perspective of period τ , it will also be optimal from the perspective of period t . Thus we will be able to replace $\tilde{\mathbf{x}}_{\tau,t}^i$ with $\hat{\mathbf{x}}_{\tau,t}^i$ once we prove existence of equilibrium with payoffs given by (5).

This approach has several advantages. First, it allows for monotone mixed strategies to be treated as monotone pure strategies. Suppose that the conditional distribution of θ_j is completely atomic at H^t , so that $\mu_t^i(\theta_j^*, \Theta_{-j} | H^t, \theta_i) = 1$ for some $\theta_j^* \in \Theta_j$. Then for some H^τ , it will be possible that $\tilde{\mathbf{x}}_{\tau,t}^j$ will vary with α_j , while $\tilde{\theta}_j^t$ will not. Second, this approach allows us to separate the strategy profile from the specific beliefs over the types of players in period t , which will depend endogenously on the actions chosen in previous periods. By doing so, we can treat each j at H^t as a distinct player in an essentially static game, with the set of types for each player distributed over $[0, 1]$.

We now verify that the constructions of $\tilde{\theta}_i^t$, $\tilde{\mathbf{x}}_t^i$, and $\hat{\mathbf{x}}_t^i$, and payoffs as given by (5), allow for the invocation of Reny’s theorem. Having gone through the above transformation of strategies and payoffs, it will be possible to view the collection $(\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T, \tilde{\theta}_i^t)$ as the strategy chosen by player i at history H^t . It should be emphasized that $\tilde{\theta}_i^t$ is now being viewed, for the purpose of the invocation of fixed-point theorems, as a choice by type α_i in period t , which as mentioned above will be determined by the beliefs. We combine these into one function, $\sigma_t^i = (\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{\tau \geq t}, \tilde{\theta}_i^t)$, which takes as arguments $(\{H^\tau\}_{\tau \geq t})$. Define the space of such functions σ_t^i as Σ_t^i . Indicating the Euclidean metric over $\prod_{\tau=t}^T (X_\tau^i)^{|\{H^\tau: H^t \subset H^\tau\}|} \times \Theta_i$ by d_t^i ,²¹ define a

²¹Note that we must count X_τ^i once for each possible subgame H^τ that is reachable from H^t .

metric over Σ_t^i by (for fixed H^t)

$$\delta_t^i(\sigma_t^i, \hat{\sigma}_t^i) = \int d_t^i(\sigma_t^i(H^t, \alpha_i), \hat{\sigma}_t^i(H^t, \alpha_i)) d\alpha_i$$

Define d_t^{-i} and δ_t^{-i} over the strategy spaces of other players $-i$ analogously.²²

In order to ensure the existence of a fixed-point (as needed to prove existence of equilibrium), we must ensure that the payoffs as defined in (5) will be continuous in the strategy profile $\{\sigma_\tau^j\}_{\tau,j}$ for player i indexed by H^t .²³ We first show that $\hat{\mathbf{x}}_{\tau,t}^j$ is continuous in $\{\tilde{\mathbf{x}}_{\tau,t}^j\}_{\tau=1}^T$.

Lemma A.2: *Consider any sequence of strategies $\{\{\tilde{\mathbf{x}}_{\tau,t,m}^j\}_{\tau \geq t}\}_{m=1}^\infty$ such that $\tilde{\mathbf{x}}_{\tau,t,m}^j \rightarrow \tilde{\mathbf{x}}_{\tau,t}^j$. Then for all j, τ, t , $\hat{\mathbf{x}}_{\tau,t,m}^j \rightarrow \hat{\mathbf{x}}_{\tau,t}^j$.*

Proof: Note that $\hat{\mathbf{x}}_{\tau,t}^i$ (respectively, $\hat{\mathbf{x}}_{\tau,t,m}^i$) divides the strategy of player i at H^τ into three subintervals of $[0, 1]$, over each of which the strategies are monotone: over two of them, the strategy is defined by $\tilde{\mathbf{x}}_{\tau,t}^i$, while over the third (which may be between the other two), it is defined by $\tilde{\mathbf{x}}_{\tau,\tau}^i$, with the interval given by $(\underline{\alpha}_i(H^\tau, H^t), \bar{\alpha}_i(H^\tau, H^t))$. We inductively show that $\hat{\mathbf{x}}_{\tau,t,m}^j \rightarrow \hat{\mathbf{x}}_{\tau,t}^j$, $\underline{\alpha}_{i,m}(H^\tau, H^t) \rightarrow \underline{\alpha}_i(H^\tau, H^t)$, and $\bar{\alpha}_{i,m}(H^\tau, H^t) \rightarrow \bar{\alpha}_i(H^\tau, H^t)$. For $\tau = t$, this is trivial because $\hat{\mathbf{x}}_{t,t}^i = \hat{\mathbf{x}}_{t,t}^i$, $\underline{\alpha}_{i,m}(H^\tau, H^t) = \underline{\alpha}_i(H^\tau, H^t) = 0$ and $\bar{\alpha}_{i,m}(H^\tau, H^t) = \bar{\alpha}_i(H^\tau, H^t) = 1$.

Given that the result is true for τ , we show that it is true for $\tau + 1$. Suppose that $H^{\tau+1}$ is only reachable from H^τ by the choice of $x_\tau^i = (x_\tau^i)^*$. Since $\hat{\mathbf{x}}_{\tau,t,m}^i \rightarrow \hat{\mathbf{x}}_{\tau,t}^i$, the set of types $\alpha_i \in (\underline{\alpha}_{i,m}(H^\tau, H^t), \bar{\alpha}_{i,m}(H^\tau, H^t))$ that can reach $H^{\tau+1}$ under $\hat{\mathbf{x}}_{\tau,t}^i$ must converge. To see this, suppose that $\hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i) < (x_\tau^i)^*$. Then for almost all such α_i , there must exist some M such that for all $m > M$, $\hat{\mathbf{x}}_{\tau,t,m}^i(H^\tau, \alpha_i) < (x_\tau^i)^*$. Thus $\lim_{m \rightarrow \infty} \underline{\alpha}_{i,m}(H^{\tau+1}, H^t) \geq \underline{\alpha}_i(H^{\tau+1}, H^t)$. A similar argument shows that $\lim_{m \rightarrow \infty} \bar{\alpha}_{i,m}(H^{\tau+1}, H^t) \leq \bar{\alpha}_i(H^{\tau+1}, H^t)$. Conversely, if $\hat{\mathbf{x}}_{\tau,t}^i(H^{\tau+1}, \alpha_i) = (x_\tau^i)^*$, then for almost all α_i , there exists some M such that for all $m > M$, $\hat{\mathbf{x}}_{\tau,t,m}^i(H^{\tau+1}, \alpha_i) = (x_\tau^i)^*$. Thus $\lim_{m \rightarrow \infty} \underline{\alpha}_{i,m}(H^{\tau+1}, H^t) \leq \underline{\alpha}_i(H^{\tau+1}, H^t)$ and $\lim_{m \rightarrow \infty} \bar{\alpha}_{i,m}(H^{\tau+1}, H^t) \geq \bar{\alpha}_i(H^{\tau+1}, H^t)$. Combining these implies that $\underline{\alpha}_{i,m}(H^{\tau+1}, H^t) \rightarrow \underline{\alpha}_i(H^{\tau+1}, H^t)$ and $\bar{\alpha}_{i,m}(H^{\tau+1}, H^t) \rightarrow \bar{\alpha}_i(H^{\tau+1}, H^t)$.

Now look at $\alpha_i \in (\underline{\alpha}_i(H^{\tau+1}, H^t), \bar{\alpha}_i(H^{\tau+1}, H^t))$. Suppose that $\hat{\mathbf{x}}_{\tau+1,t}^i(H^{\tau+1}, \alpha_i) = x_{\tau+1}^i$. For almost all such α_i , there exists $\epsilon > 0$ such that $\hat{\mathbf{x}}_{\tau+1,t}^i(H^{\tau+1}, \alpha_i - \epsilon) = \hat{\mathbf{x}}_{\tau+1,t}^i(H^{\tau+1}, \alpha_i + \epsilon) = x_{\tau+1}^i$. Let $\hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i(H^{\tau+1}, H^t)}{\bar{\alpha}_i(H^{\tau+1}, H^t) - \underline{\alpha}_i(H^{\tau+1}, H^t)}$ and define $\hat{\alpha}_{i,m}$ analogously. Note that $\hat{\alpha}_{i,m} \rightarrow$

²²I.e. the \mathcal{L}^1 metric.

²³As discussed in the footnote by Theorem A.1, this continuity condition will be sufficient to guarantee existence of equilibrium in the translated game, as it will ensure that the set of best-replies for type θ_i is upper-hemicontinuous in the other players' strategies. Reny (2011) showed that (Σ_t^i, δ_t^i) forms a compact absolute retract, and so he invoked the fixed-point theorem of Eilenberg and Montgomery (1946) to prove that an equilibrium existed over such strategy functions as long as payoffs are continuous in this metric. See Section 6 of his paper for details.

$\hat{\alpha}_i$. For almost all $\hat{\alpha}_i \in [0, 1]$, it will be the case that if $\tilde{\mathbf{x}}_{\tau+1, \tau+1}^i(H^{\tau+1}, \hat{\alpha}_i) = x_{\tau+1}^i$, then $\tilde{\mathbf{x}}_{\tau+1, \tau+1}^i(H^{\tau+1}, \hat{\alpha}_i - \epsilon) = \tilde{\mathbf{x}}_{\tau+1, \tau+1}^i(H^{\tau+1}, \hat{\alpha}_i + \epsilon) = x_{\tau+1}^i$. Since $\tilde{\mathbf{x}}_{\tau+1, \tau+1, m}^i \rightarrow \tilde{\mathbf{x}}_{\tau+1, \tau+1}^i$, it follows that there exists M such that for all $m > M$, $\tilde{\mathbf{x}}_{\tau+1, \tau+1, m}^i(H^{\tau+1}, \hat{\alpha}_i - \frac{\epsilon}{2}) = \tilde{\mathbf{x}}_{\tau+1, \tau+1, m}^i(H^{\tau+1}, \hat{\alpha}_i + \frac{\epsilon}{2}) = x_{\tau+1}^i$. Moreover, for sufficiently high M , since $\hat{\alpha}_{i, m} \rightarrow \hat{\alpha}_i$, it follows that for $m > M$, $\hat{\alpha}_{i, m} \in (\hat{\alpha}_i - \frac{\epsilon}{2}, \hat{\alpha}_i + \frac{\epsilon}{2})$. Since $\hat{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i) = \tilde{\mathbf{x}}_{\tau+1, \tau+1}^i(H^{\tau+1}, \hat{\alpha}_i)$, it follows that $\hat{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i) \rightarrow \hat{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i)$ for $\alpha_i \in (\underline{\alpha}_i(H^{\tau+1}, H^t), \bar{\alpha}_i(H^{\tau+1}, H^t))$.

For $\alpha_i \notin [\underline{\alpha}_i(H^{\tau+1}, H^t), \bar{\alpha}_i(H^{\tau+1}, H^t)]$, it will be the case that for some M , if $m > M$, then $\alpha_i \notin [\underline{\alpha}_{i, m}(H^{\tau+1}, H^t), \bar{\alpha}_{i, m}(H^{\tau+1}, H^t)]$, and so $\hat{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i)$ will be defined by $\tilde{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i)$ for all such m . The argument is similar to that for $\alpha_i \in (\underline{\alpha}_i(H^{\tau+1}, H^t), \bar{\alpha}_i(H^{\tau+1}, H^t))$: for almost all such α_i , there will exist $\epsilon > 0$ such that $\tilde{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i - \epsilon) = \tilde{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i + \epsilon) = x_{\tau+1}^i$. Since $\tilde{\mathbf{x}}_{\tau+1, t, m}^i \rightarrow \tilde{\mathbf{x}}_{\tau+1, t}^i$, then for sufficiently high M , if $m > M$, then $\tilde{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i - \frac{\epsilon}{2}) = \tilde{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i + \frac{\epsilon}{2}) = x_{\tau+1}^i$, and so $\hat{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i) \rightarrow \hat{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i)$.

Since $\hat{\mathbf{x}}_{\tau+1, t, m}^i(H^{\tau+1}, \alpha_i) \rightarrow \hat{\mathbf{x}}_{\tau+1, t}^i(H^{\tau+1}, \alpha_i)$ pointwise almost-everywhere, it follows that $\hat{\mathbf{x}}_{\tau+1, t, m}^i \rightarrow \hat{\mathbf{x}}_{\tau+1, t}^i$ in the topology given by δ_t^i . \square

We must also argue that the beliefs vary continuously as well in the strategies chosen by all players indexed by H^t , as this will affect both $\tilde{\theta}_i^t$ and g_t . We proceed by first showing that $\tilde{\theta}_i^t$ and g_t are continuous in the beliefs held by the other players in the dynamic game.

Lemma A.3: *Consider a sequence of beliefs conditional on (H^t, α_i) , $\{\mu_{t, k}^i\}_{k=1}^\infty$ that converges to some μ_t^i in the weak-* topology, where all $\mu_{t, k}^i$ and μ_t^i are the prior restricted to a Cartesian product of subintervals. Let $\tilde{\theta}_{j, k}^t$, $\tilde{\theta}_j^t$, $g_{t, k}$, and g_t be the corresponding functions describing the distribution of types of all players j (including i). Then $\tilde{\theta}_{j, k}^t(H^t, \cdot) \rightarrow \tilde{\theta}_j^t(H^t, \cdot)$ and $g_{t, k}(\cdot | H^t, \cdot) \rightarrow g_t(\cdot | H^t, \cdot)$ everywhere.*

Proof: By Lemma A.1, the conditional distributions given H^t over θ_{-i} are those defined by the prior restricted to a Cartesian product of subintervals $[\theta_{j, m}^1, \theta_{j, m}^2]$. If $\mu_{t, m}^i \rightarrow \mu_t^i$ in the weak-* topology, then it must be that $\theta_{j, m}^1 \rightarrow \theta_j^1$ and $\theta_{j, m}^2 \rightarrow \theta_j^2$; hence $\tilde{\theta}_{j, m}^t \rightarrow \tilde{\theta}_j^t$. The proof for $g_{t, m}$ then follows from the fact that $f(\theta_{-i} | \theta_i)$ is continuous in θ_i , and $\tilde{\theta}_{i, m}^t \rightarrow \tilde{\theta}_i^t$ everywhere. \square

We now show the other direction, i.e. that the beliefs at H^t vary continuously in the strategies of all players. Define $\sigma = \{\sigma_t^i\}_{i, t}$. A *belief mapping* ψ is a function from monotone strategy profiles to the collection of beliefs in all periods; that is,

$$\psi : \Sigma \rightarrow \mathcal{M}$$

$$\sigma \rightarrow \mu$$

such that the beliefs for any H^t coincide with those generated by the strategy profile σ

(specifically, $\{\tilde{\mathbf{x}}_{\tau,\tau}\}_{\tau=1}^{t-1}$) via Bayesian consistency whenever possible, as defined in Section 2). Specifically, the beliefs at any H^t over θ will be those inductively pinned down by restricting the prior over θ_j to the corresponding intervals $\{(\underline{\alpha}_j(H^t, H^1), \bar{\alpha}_j(H^t, H^1))\}$ as given from when we originally defined $\{\{\hat{\mathbf{x}}_{\tau,1}^j\}_{\tau=1}^{t-1}\}_j$, which were themselves defined (on-path) by $\{\tilde{\mathbf{x}}_{\tau,\tau}^j\}_{j,\tau}$.

For actions by player j that generate off-path H^t , beliefs will not be pinned down by the above strategies. Hence we must additionally restrict ψ to generate beliefs over Θ that are restrictions of the prior to a Cartesian product of intervals in the same sense as Lemma A.1.²⁴

Definition A.3: A belief mapping ψ is *continuous* if player i 's belief at H^t , given by $(\psi(\sigma_m(\cdot)))_t^i(\cdot|H^t, \theta_i)$, converges in the weak-* topology to $(\psi(\sigma(\cdot)))_t^i(\cdot|H^t, \theta_i)$ for any sequence of strategies $\{\sigma_m(\cdot)\}_{m=1}^\infty$ that converge to $\sigma(\cdot)$.

The continuity condition of ψ is to ensure that the beliefs (and hence the payoffs) in period t are continuous in the strategies chosen in periods $\tau < t$. This will in turn ensure that the incentive to “jump back” to play an action that is off-path does not exist, as we can generate the beliefs at any off-path H^t as the limit of beliefs generated from strategies that place it on-path. It will thus enable the continuity of payoffs in the strategy profile σ .

Theorem 3.2 claimed that we were able to pin down the beliefs that must be held given a deviation by player j . We restate this theorem here, reworded slightly, as it turns out that these beliefs must necessarily arise in any continuous ψ . Since these are the beliefs that are used in our equilibrium construction, Theorem 3.2 as stated in Section 3.1 is an immediate corollary of Theorem 3.2* once we show equilibrium existence.

Theorem 3.2*: *There exists a unique continuous ψ , which assigns, for any $H^\tau \subset H^t$,*

$$\mu_t^i(\{\sup\{\theta_j : \mathbf{x}_\tau^j(H^\tau, \theta_j) < x_\tau^j\}\}, \Theta_{-\{i,j\}}|H^t, \theta_i) = 1$$

if x_τ^j is off-path.

Proof: For H^t that is on path and any open $A \subset \Theta$, by the construction of ψ and $\hat{x}_{\tau,1}^i$ (suppressing arguments),

$$\mu_t^i(A|H^t, \tilde{\theta}_i^1(H^1, \alpha_i)) = \frac{\int \mathbf{1}_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\hat{\mathbf{x}}_{\tau,1}^i, \hat{\mathbf{x}}_{\tau,1}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i}|H^1, \alpha_i) d\alpha_{-i}}{\int \mathbf{1}_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\hat{\mathbf{x}}_{\tau,1}^i, \hat{\mathbf{x}}_{\tau,1}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i}|H^1, \alpha_i) d\alpha_{-i}} \quad (6)$$

These integrals converge when $\sigma_m \rightarrow \sigma$ since $\tilde{\theta}_i^1$ is independent of the strategy chosen (as it is given by the prior), so ψ is continuous on-path.

To extend to the cases where H^t is off-path, note that the denominator in (6) is 0. Let R be the set of pairs $\{(j, \tau)\}$ representing the set of players whose period- τ actions are off-path

²⁴As alluded to earlier, this will actually hold without loss of generality for all continuous ψ .

given the strategy profile σ . Consider a sequence $\{\sigma_m\}$ of monotone strategies in which H^t is on-path which converge to σ . Fixing j , let t_j be the first period τ in which $x_\tau^j \in H^t$ is off-path according to $\tilde{\mathbf{x}}_{\tau,\tau}^j$. Then for any $x_{t_j}^j \in X_{t_j}^j$ and almost all α_j such that $\tilde{\mathbf{x}}_{t_j,t_j}^j(H^{t_j}, \alpha_j) < x_{t_j}^j$, there must exist some M such that for all $m > M$, $\tilde{\mathbf{x}}_{t_j,t_j,m}^j(H^{t_j}, \alpha_j) < x_{t_j}^j$. A similar argument holds for $\tilde{\mathbf{x}}_{t_j,t_j}^j(H^{t_j}, \alpha_j) > x_{t_j}^j$.

Let $\underline{\theta}_{j,m} = \sup\{\theta_j : \tilde{\mathbf{x}}_{t_j,t_j,m}^j(H^{t_j}, \alpha_j) < x_{t_j}^j \text{ and } \tilde{\theta}_j^{t_j}(H^{t_j}, \alpha_j) = \theta_j\}$ and $\bar{\theta}_{j,m} = \inf\{\theta_j : \tilde{\mathbf{x}}_{t_j,t_j,m}^j(H^{t_j}, \alpha_j) > x_{t_j}^j \text{ and } \tilde{\theta}_j^{t_j}(H^{t_j}, \alpha_j) = \theta_j\}$. Then, as shown in Lemma A.1, the support of types who choose action x_{t-1}^j will be contained in $[\underline{\theta}_{j,m}, \bar{\theta}_{j,m}]$. By the argument from the previous paragraph, $\lim_{m \rightarrow \infty} \underline{\theta}_{j,m} = \lim_{m \rightarrow \infty} \bar{\theta}_{j,m} \equiv \theta_j^*$. Hence μ_t^i as generated by ψ must place probability 1 on θ_j^* conditional on observing H^t . For subsequent τ , because θ_j must be contained in the set of types who chose action x_{t-1}^j for each σ_m , it follows that the beliefs over θ_j for the subsequent τ converge to the same θ_j^* .

For H^t that is off-path, for any given τ , any deviation by j such that $(j, \tau) \in R$ cannot affect the choice of strategies by i such that $(i, \tau) \notin R$ due to conditional independence of strategies. Hence for any open set $A \subset \Theta$ containing some element with $\theta_j = \theta_j^*$ for all j that is off-path, and $H^t \supset \{x_\tau^j\}_{(j,\tau) \in R}$, μ_t^i is uniquely determined to set

$$\mu_t^i(A|H^t, \tilde{\theta}_i^1(H^1, \alpha_i)) = \lim_{m \rightarrow \infty} \frac{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\tilde{\mathbf{x}}_{\tau,1,m}^i, \tilde{\mathbf{x}}_{\tau,1,m}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i}|H^1, \alpha_i) d\alpha_{-i}}{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\tilde{\mathbf{x}}_{\tau,1,m}^i, \tilde{\mathbf{x}}_{\tau,1,m}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i}|H^1, \alpha_i) d\alpha_{-i}}$$

as the probability that $\tilde{\theta}_1^i$ assigns $\theta_j \in (\theta_j^* - \epsilon, \theta_j^* + \epsilon)$ conditional on H^t being reached approaches 1 for any $\epsilon > 0$. \square

We have now defined the sense in which player i 's interim payoffs are continuous in the choice of strategies of all other players (via $\hat{\mathbf{x}}_{\tau,t}^{-i}$), and in one's own past strategies (via $(\tilde{\theta}_i^t, \tilde{\theta}_{-i}^t, g_t)$). The construction of ψ shows that, by indexing each player i according to each history H^t , and consider each as separate players (albeit with the same preferences), if we perturb the strategy profiles of all players in previous periods continuously, then we change the beliefs continuously as well. Since beliefs are continuous in the strategies chosen, it therefore follows that payoffs are continuous in the strategy profiles chosen in all periods as well, and so the best-reply correspondence will be upper-hemicontinuous. Thus we can invoke Theorem A.1 to generate existence of equilibrium in the static game in which payoffs are given by (5).

Lemma A.4: *Suppose that the conditions of the model are satisfied in the dynamic game, and best-plies to monotone strategies are increasing in the strong-set order. Then there exists a monotone pure-strategy equilibrium in the static game Γ^2 in which payoffs are given by (5).*

Proof: We must show that conditions (i)-(iv) of Theorem A.1 hold in the static interpre-

tation of the game. Our translation of the dynamic game to a static environment involves interpreting the types of opposing players as corresponding to some value of $\alpha_j \in [0, 1]$, while the action space in each period remains the same. One must take care with preserving the lattice property. In the original game, this is trivial, as actions are one-dimensional. In the static game, we define the following partial order. Fix player i 's opponents actions at each subgame. The partial order for the action space in the static game operates lexicographically over periods, so that if player i has two potential vectors of actions that he considers from $\prod_{t=1}^T (X_t^i)^{|H^t|}$ in Γ^2 that agree up to period τ , then the join of the two picks the continuation from the vector with the higher action at subgame H^τ , and the meet picks the lower. Thus, if given opponents' actions, the two vectors of actions would induce the sequence of actions $x^i, \hat{x}^i \in \prod_{t=1}^T X_t^i$, such that x^i and \hat{x}^i agree upto period τ , and $x_\tau^i > \hat{x}_\tau^i$, then $x^i \vee \hat{x}^i = x^i$.²⁵ This forms a well-defined lattice for the action space in Γ^2 , so (i) is satisfied. Moreover, the set of best-replies for any type is automatically join-closed because each X_t^i is finite.

Since g_t is absolutely continuous, the conditional distribution given by the measure over α_{-i} is uniform, so it satisfies condition (ii). Condition (iii) is satisfied because the interim payoff given H^t from choosing x_t^i can be rewritten as (suppressing arguments for $\tilde{\mathbf{x}}_{\tau,t}, \hat{\mathbf{x}}_{\tau,t}$, and $\tilde{\theta}^t$),

$$\int u_i(H^t, x_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t+1}^t; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i}$$

By Theorem 3.2*, this will be continuous in $\{\sigma_{t'}^j\} = \{\tilde{\mathbf{x}}_{\tau,t'}^j, \tilde{\theta}_j^{t'}\}_{t',\tau,j}$, and so the best-reply correspondence will be upper-hemicontinuous. Lastly, condition (iv) is assumed in the dynamic game, and so automatically holds in the static interpretation because preferences as given by u_i are the same, and $\hat{\mathbf{x}}_{\tau,t}^{-i}$ is monotone over the relevant intervals (i.e. the portion that is reachable from the perspective of period t). Hence the conditions of Theorem A.1 are satisfied for the static interpretation as well, and so a monotone equilibrium will exist. \square

We now use the equilibrium that we have constructed in Lemma A.4 to construct an equilibrium in the original dynamic game, in which payoffs were given by (3). We use $\tilde{\mathbf{x}}_{\tau,t}^i$ and $\hat{\mathbf{x}}_{\tau,t}^i$ to construct a consistent strategy by player i that is both a best-reply and preserves the payoffs of the other players. First, we show that the types who choose an action $x_\tau^i \in X_\tau^i$ according to $\tilde{\mathbf{x}}_{\tau,\tau}^i$ and those who choose $x_\tau^i \in X_\tau^i$ via $\hat{\mathbf{x}}_{\tau,t}^i$ must align on path from $\{\hat{\mathbf{x}}_{t',t}^i\}$, and thus be consistent.

Lemma A.5: *Define $\underline{\alpha}_i(H^\tau, H^t)$ and $\bar{\alpha}_i(H^\tau, H^t)$ as in the definition of $\hat{x}_{\tau,t}^i$. Then in any equilibrium found by Lemma A.4, $\tilde{\theta}_i^t(H^t, \underline{\alpha}_i(H^\tau, H^t)) = \tilde{\theta}_i^\tau(H^\tau, 0)$ and $\tilde{\theta}_i^t(H^t, \bar{\alpha}_i(H^\tau, H^t)) =$*

²⁵Unlike the Euclidean partial order, this avoids the potential issue that a certain choice $x_{t'}^i$, where $t' > \tau$, is suboptimal conditional on reaching subgame $H^{t'}$, but is admissible within a best-reply since $H^{t'}$ is off-path due to some action by i in period τ .

$\tilde{\theta}_i^\tau(H^\tau, 1)$.

Proof: By construction of ψ , if the distribution of θ_i is completely atomic at H^t , then

$$\tilde{\theta}_i^t(H^t, \underline{\alpha}_i(H^\tau, H^t)) = \tilde{\theta}_i^\tau(H^\tau, 0) = \tilde{\theta}_i^t(H^t, \bar{\alpha}_i(H^\tau, H^t)) = \tilde{\theta}_i^\tau(H^\tau, 1) = \theta_i^*$$

and we are done.

Otherwise, we proceed by induction on τ . By construction of ψ , the set of types θ_i defined by the interval $[\tilde{\theta}_i^t(H^t, 0), \tilde{\theta}_i^t(H^t, 1)]$ must be equal to the set of types who choose $\{\hat{\mathbf{x}}_{t',1}^i\}_{t'=1}^{t-1} = \{x_{t',t}^i\}_{t'=1}^{t-1}$ at the respective $H^{t'} \subset H^t$ for which H^t is reachable from the perspective of period 1. The same holds for $[\tilde{\theta}_i^\tau(H^\tau, 0), \tilde{\theta}_i^\tau(H^\tau, 1)]$.

Next, note that in equilibrium, $\hat{\mathbf{x}}_{t,1}^i(H^t, \alpha_i) = \tilde{\mathbf{x}}_{t,t}^i(H^t, \hat{\alpha}_i)$, where $\hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i(H^t, H^1)}{\bar{\alpha}_i(H^t, H^1) - \underline{\alpha}_i(H^t, H^1)}$. We show this by induction. This is obviously true for $t = 1$. Given that this is true for $t - 1$, then by the previous paragraph and the construction of $\hat{\mathbf{x}}_{t,1}^i$ and ψ , in period t , the set of types $\hat{\alpha}_i \in [\underline{\alpha}_i(H^{t+1}, H^t), \bar{\alpha}_i(H^{t+1}, H^t)]$ who choose $\tilde{\mathbf{x}}_{t,t}^i(H^t, \hat{\alpha}_i) = x_t^i \in X_t^i$ which can reach some H^{t+1} generates the same set of types $\{\theta_i : \tilde{\theta}_i^t(H^t, \hat{\alpha}_i) = \theta_i\}$ as $\{\theta_i : \tilde{\theta}_i^1(H^1, \alpha_i) = \theta_i\}$, where $\alpha_i \in [\underline{\alpha}_i(H^{t+1}, H^1), \bar{\alpha}_i(H^{t+1}, H^1)]$ i.e. α_i chooses (conditional on each relevant $H^{t'} \subset H^{t+1}$ being reached) $\{\hat{\mathbf{x}}_{t',1}^i\}_{t'=1}^t = \{x_{t'}^i\}_{t'=1}^t$ which can reach H^{t+1} .

Lastly, suppose that $\hat{\mathbf{x}}_{\tau,1}^i(H^\tau, \alpha_i) = \hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \hat{\alpha}_i)$. Then for $\tau + 1$, by a similar argument to that of the previous paragraph, it must be that $\hat{\mathbf{x}}_{\tau+1,1}^i(H^{\tau+1}, \alpha_i) = \hat{\mathbf{x}}_{\tau+1,t}^i(H^{\tau+1}, \hat{\alpha}_i)$, as the underlying sets of types θ_i coincide. Thus we have

$$\tilde{\theta}_i^t(H^t, \underline{\alpha}_i(H^\tau, H^t)) = \tilde{\theta}_i^1(H^1, \underline{\alpha}_i(H^\tau, H^1)) = \tilde{\theta}_i^\tau(H^\tau, 0)$$

$$\tilde{\theta}_i^t(H^t, \bar{\alpha}_i(H^\tau, H^t)) = \tilde{\theta}_i^1(H^1, \bar{\alpha}_i(H^\tau, H^1)) = \tilde{\theta}_i^\tau(H^\tau, 1)$$

□

Since $\tilde{\mathbf{x}}_{\tau,\tau}^i$ forms a best-reply, and $\{\hat{\mathbf{x}}_{\tau,t}^i\}$ is consistent in equilibrium, we can now show that, in equilibrium, $\{\hat{\mathbf{x}}_{\tau,t}^i\}$ forms a best-reply to other players' choice of strategies $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$.

Lemma A.6: *Suppose that all players i use monotone strategies $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ in an equilibrium of Γ^2 . Then $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T$ is also a best-reply for each α_i .*

Proof: We show this by backward induction on τ . This is trivial for $t = \tau = T$ since $\hat{\mathbf{x}}_{T,T}^i = \tilde{\mathbf{x}}_{T,T}^i$. Suppose that from period $\tau + 1$ onward, $\{\hat{\mathbf{x}}_{t',t}^i\}_{t'=\tau+1}^T$ is a collection of best-replies from the perspective of period t conditional on reaching $H^{\tau+1}$. We show that replacing $\tilde{\mathbf{x}}_{\tau,t}^i$ with $\hat{\mathbf{x}}_{\tau,t}^i$ does not decrease the payoff of player i . Recall that the payoff from choosing

$\{x_{t'}^i\}_{t'=t}^\tau$ through period τ and then $\{\tilde{x}_{t',t}^i\}_{t'=\tau+1}^T$ (suppressing arguments) is

$$\int u_i(H^t, \{x_{t'}^i, \hat{x}_{t',t}^{-i}\}_{t'=t}^\tau, \{\tilde{x}_{t',t}^i, \hat{x}_{t',t}^{-i}\}_{t'=\tau+1}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i}$$

Note that $\{\tilde{x}_{t',t}^i\}_{t'=\tau+1}^T$ is a best-reply conditional on $H^{\tau+1}$ being reached. By the induction hypothesis, the above equation must be equal to

$$\int u_i(H^t, \{x_{t'}^i, \hat{x}_{t',t}^{-i}\}_{t'=t}^\tau, \{\hat{x}_{t',t}^i, \hat{x}_{t',t}^{-i}\}_{t'=\tau+1}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i}$$

We now show that the conditional payoff upon reaching H^τ is maximized by replacing $\tilde{x}_{\tau,t}^i$ with $\hat{x}_{\tau,t}^i$ is not decreased. In the case where $\tilde{x}_{\tau,t}^i = \hat{x}_{\tau,t}^i$, this is true by definition since $\tilde{x}_{\tau,t}^i$ was optimal. On the other hand, by Lemma A.5, if $\hat{x}_{\tau,t}^i(H^\tau, \alpha_i) = \tilde{x}_{\tau,\tau}^i(H^\tau, \hat{\alpha}_i)$ where $\tilde{\theta}_i^{\tau+1}(H^{\tau+1}, \hat{\alpha}_i) = \tilde{\theta}_i^\tau(H^\tau, \alpha_i)$, then since $\tilde{x}_{\tau,\tau}^i(H^\tau, \hat{\alpha}_i)$ is optimal for $\hat{\alpha}_i$ from the perspective of period τ , $\hat{x}_{\tau,t}^i(H^\tau, \alpha_i)$ must be optimal as well from the perspective of period t because $\{\hat{x}_{\tau,t}^{-i}\}_{-i,\tau,t}$ forms a consistent strategy profile. Thus in either case, $\hat{x}_{\tau,t}^i(H^\tau, \alpha_i)$ is optimal. \square

Note that if player i 's strategy is consistent, the portion of $\tilde{x}_{\tau,t}^i$ that is chosen by α_i (indexed by H^t) that does not choose actions $\{x_t^i, \dots, x_{\tau-1}^i\}$ that generate H^τ with positive probability is essentially irrelevant for the purposes of players $-i$, since type α_i will never reach such H^τ in equilibrium.

The payoff of player j thus only depends (given that H^τ is reached) on $(\tilde{x}_{\tau,\tau}^i, \tilde{\theta}_i^\tau)$. Therefore, for the purposes of the payoff of player i in period t , we need only set player i 's strategy to be consistent for those types α_i that use strategies under $\hat{x}_{\tau,t}^i$ from which H^τ is reachable, and complete the strategy function with an arbitrary consistent monotone best-reply function (which will exist since the best-reply correspondence is increasing in the strong set order in type). This transformation will not affect other player's equilibrium payoffs, which for either choice just depend on $\{\hat{x}_{\tau,t}^i\}_{\tau=t}^T$ as generated from $\{\tilde{x}_{\tau,\tau}^i\}_{\tau=t}^T$.

We are therefore now able to construct the equilibrium strategies of the dynamic game. Suppose that the conditional distribution of θ_i at H^t is absolutely continuous. Let $\underline{\alpha}_i = \underline{\alpha}_i(H^\tau, H^t)$ and $\bar{\alpha}_i = \bar{\alpha}_i(H^\tau, H^t)$. Then define

$$(x_\tau^i)^*(H^\tau, \tilde{\theta}_i^t(H^t, \alpha_i)) \equiv \begin{cases} \tilde{x}_{\tau,\tau}^i(H^\tau, \hat{\alpha}_i), & \alpha_i \in (\underline{\alpha}_i, \bar{\alpha}_i) \text{ and } \hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i}{\bar{\alpha}_i - \underline{\alpha}_i} \\ \max\{x_\tau^i \in BR_\tau^i(H^\tau, \tilde{\theta}_i^t(H^t, \alpha_i))\}, & \alpha_i \geq \bar{\alpha}_i(H^\tau, H^t) \\ \min\{x_\tau^i \in BR_\tau^i(H^\tau, \tilde{\theta}_i^t(H^t, \alpha_i))\}, & \alpha_i \leq \underline{\alpha}_i(H^\tau, H^t) \end{cases}$$

Otherwise, if θ_i is completely atomic, then $\tilde{\theta}_i(H^\tau, 0) = \tilde{\theta}_i(H^\tau, 1) = \tilde{\theta}_i(H^t, 0) = \tilde{\theta}_i(H^t, 1) = \theta_i^*$.

Then define the mixed strategy over X_τ^i so that

$$\rho_\tau^i(x_\tau^i|H^\tau, \theta_i) = \bar{\alpha}_i(H^{\tau+1}, H^\tau) - \underline{\alpha}_i(H^{\tau+1}, H^\tau)$$

where $H^{\tau+1}$ is reachable only from choosing x_τ^i at H^τ , and completing the strategy for off-path values of θ_i as in the case with absolutely continuous types. It is easy to verify that such strategies are monotone in θ_i at each subgame since the set of best-replies at each subgame is increasing in the strong set order in θ_i .

We now argue that these strategies form an equilibrium in the dynamic game.

Proof of Theorem 3.1: By Lemma A.4, there exists an equilibrium in the static game in which payoffs are given by (5). As argued above in Lemma A.6, choosing $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ is also a best reply given equilibrium strategy profile $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ when payoffs are given by (5); moreover, substituting these strategies does not affect the payoffs of the other players since $\hat{\mathbf{x}}_{t,t}^i = \tilde{\mathbf{x}}_{t,t}^i$ for all i, t . By Lemma A.5, the set of types of player i that choose a given x_τ^i must align on path from both $\hat{\mathbf{x}}_{\tau,t}^i$ and $\hat{\mathbf{x}}_{\tau,\tau}^i$ on-path from $\{\hat{\mathbf{x}}_{t',t}^i\}_{t'=t}^\tau$, and so $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{\tau,t}$ form a strategy for player i that is consistent on-path.

As mentioned before, $\hat{\mathbf{x}}_{\tau,t}^i$ may not be monotonic off-path. However, since what is off-path does not affect the payoffs of players $-i$, we can set them arbitrarily as long as they form a best-reply for player i . To ensure the existence of monotone best-replies when including types that are off-path at H^τ , we must ensure that there is then a best-reply $x_\tau^i \in BR(H^\tau, \tilde{\theta}_i^t(H^t, \alpha_i))$ for $\alpha_i \geq \bar{\alpha}_i(H^\tau, H^t)$ that is at least as great as $\hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \bar{\alpha}_i(H^\tau, H^t))$. Fortunately, by the fact that best-replies are increasing in the strong-set order, it must be that $\max\{x_\tau^i \in BR(H^\tau, \tilde{\theta}_i^t(H^t, \alpha_i))\} \geq \hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \bar{\alpha}_i(H^\tau, H^t))$, and so will be monotone. An analogous argument holds for $\alpha_i \leq \underline{\alpha}_i(H^\tau, H^t)$. Choosing such values for what type α_i would do at H^τ from the perspective of H^t is therefore a monotone best-reply.

Suppose that we look at the interim payoffs in Γ^1 as given by (3), i.e. if the strategy profile is $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$, then

$$U_i^t(H^t, \tilde{\mathbf{x}}_{t,t}^i(H^t, \alpha_i), \tilde{\theta}_i^t(H^t, \alpha_i)) = \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \tilde{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t, \alpha_i) d\alpha_{-i}$$

Note now that if we use the strategies described in the above by $(x_\tau^i)^*$, the strategies are consistent, and so these coincide with the payoffs given by (5) by Lemma A.5. Therefore the strategy given by $(x_{\tau,t}^i)^*$ is optimal from the perspective of period t for type $\tilde{\theta}_i(H^t, \alpha_i)$ when the conditional distribution over θ_i is absolutely continuous at H^t ; similarly, when the distribution is completely atomic, the strategy given by $\rho_\tau^i(x_\tau^i|H^\tau, \theta_i)$ gives a correct prediction of what will be (optimally) done at H^τ from the perspective of H^t and θ_i^* . Thus

these strategies form an equilibrium of the original dynamic game.

Note that the result here only guarantees the existence of a mixed-strategy equilibrium, as there is no guarantee that the function $\tilde{\theta}_i^t$ assigns the same value to θ_i only on sets of measure 0. Thus it could be that a positive measure of values of α map to the same value of θ_i at H^t . Fortunately, without loss of generality, one can restrict attention to equilibria that are in pure strategies on-path. The reason is that, as the set of actions in any given period (except possibly period T) is finite, and strategies are monotone in equilibrium, almost all values of α must in equilibrium lead to some collection of actions x_t that a positive measure of values of $\alpha \in [0, 1]^N$ chooses. Since the original distribution F was absolutely continuous, we have shown inductively that any on-path H^t also involve absolutely continuous distributions over types by Lemma 3.1. Since strategies are pure from the perspective of type α_i , this implies that if the conditional distribution is absolutely continuous, the strategies are pure from the perspective of θ_i as well. Thus we can extend Theorem 3.1 to incorporate pure strategies whenever the conditional distribution of θ_i is not completely atomic. Thus the PBE will be monotonic in pure strategies on-path. \square

We now list the proofs of the extensions provided in Section 3.3.

Proof of Theorem 3.3: As in the proof of Theorem 3.1, we reinterpret the dynamic game as a static one. By assumption, such a reinterpretation is possible; thus the only remaining objective is to show that the symmetry of players is preserved in all subgames. We proceed inductively. Let \mathcal{I}_t be the set of players who are symmetric at H^t . In period T , we immediately have symmetry as this is essentially a static environment, since current actions do not affect the future. Hence for $i, j \in \mathcal{I}_T$, for any given $x_T^i \in X_T^i = X_T^j$,

$$\int u_i(H^T, x_T^i, \tilde{\mathbf{x}}_{T,T}^{-i}; \tilde{\theta}_i^T, \tilde{\theta}_{-i}^T) g_T(\alpha_{-i} | H^T, \alpha_i) d\alpha_{-i} = \int u_j(H^T, x_T^i, \tilde{\mathbf{x}}_{T,T}^{-j}; \tilde{\theta}_j^T, \tilde{\theta}_{-j}^T) g_T(\alpha_{-j} | H^T, \alpha_j) d\alpha_{-j}$$

Now suppose that players $i \in \mathcal{I}_t$ use symmetric monotone strategies in period t , and that we restrict our attention to equilibria in the subgames given by H^{t+1} that are symmetric in the sense of $\pi(\mathcal{I}_{t+1})$, i.e. the strategies given by $C^t(H^{t+1}, \theta)$ are permuted by $\pi(\mathcal{I}_t)$ if x_t and $\mu_{t+1}(\cdot | H^{t+1}, \theta)$ are permuted via $\pi(\mathcal{I}_t)$. Then the distribution of continuations subgames starting from period t will be symmetric, implying that the incentives in period t given by $U_i^t(H^t, x_t^i, \theta_i)$ are symmetric (in the sense of Condition (3)), i.e. for $i, j \in \pi(\mathcal{I}_t)$, choosing the strategy $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T$ at H^t yields, when payoffs are given by equation (3),

$$\int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \tilde{\mathbf{x}}_{\tau,t}^{-i}\}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} = \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \tilde{\mathbf{x}}_{\tau,t}^{-j}\}; \tilde{\theta}_j^t, \tilde{\theta}_{-j}^t) g_t(\alpha_{-j} | H^t, \alpha_j) d\alpha_{-j}$$

This implies symmetry of the subgame in period t . As symmetry is preserved in all subgames

in the sense defined above, we can invoke Theorem 4.5 of Reny (2011) to establish existence of a symmetric monotone equilibrium in the transformed static game (the proof is identical to that of Lemma A.4). To translate this into the dynamic game, we apply Lemmas A.5 and A.6 to show that we can generate consistent strategies $\{\tilde{\mathbf{x}}_{\tau,t}^i\}$ such that the payoffs as given by equation (3) match those given by equation (5). Thus there will exist a monotone PBE which is symmetric, i.e. for $i, j \in \mathcal{I}$, if in H^t , player i follows strategy $\mathbf{x}_t^i(H^t, \cdot)$ and player j follows strategy $\mathbf{x}_t^j(H^t, \cdot)$, then at $H^{\pi(\mathcal{I}),t}$ (defined as the subgame at which players' actions up to period t have been permuted according to π), i follows strategy $\mathbf{x}_t^j(H^t, \cdot)$ and j follows $\mathbf{x}_t^i(H^t, \cdot)$ (randomizing with the same probabilities ρ if necessary). \square

Proof of Theorem 3.4: Consider a sequence of truncations $\{\Gamma_m\}_{m=1}^\infty$ indicated by the stopping times, $\{T_m\}_{m=1}^\infty$, where $\lim_{m \rightarrow \infty} T_m = \infty$. The number of players in each truncation is $N_m \equiv N_{T_m}$. We modify the payoff functions accordingly to be

$$u_{i,m}(x_1, \dots, x_{T_m}, \theta_1, \dots, \theta_{N_m}) = E_{\theta_j: j > N_m} [\sup_{C^{T_m}} u_i(x_1, \dots, x_{T_m}, C^{T_m}, \theta)]$$

We index each player by H^t ; by Assumption 3.2, there are a countable number of such players in Γ_m . We define ψ as in the finite case. For each indexed player, the equilibrium function $\sigma_{t,m}^i$ (as defined in Section 3) is monotonic. We consider the sequence $\{\Gamma_m, \sigma_{1,m}^1, \dots, \sigma_{T_m,m}^{N_m}\}$; by Helly's selection theorem and Tychonoff's theorem, there exists a convergent subsequence to $\{\Gamma_m, \sigma_{n,t}^i\}_{i,n,t}$. Thus $g_{t,m} \rightarrow g_t$ by Theorem 3.2*, and $\hat{\mathbf{x}}_{\tau,t,m}^i \rightarrow \hat{\mathbf{x}}_{\tau,t}^i$ for all $\tau \geq t$ by Lemma A.3. We check that the limit strategies form an equilibrium in the static game in which payoffs are given by (5).

Without loss of generality, let the convergent subsequence of $\{\Gamma_m\}$ be the sequence itself. Note that we can subtract under the integral sign due to the uniform convergence implied by continuity at infinity.²⁶ By continuity of payoffs, we have that for any $\epsilon > 0$ and any t , there exists M such that for all $m > M$,

$$\begin{aligned} & \left\| \int u_{i,m}(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}; \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right. \\ & \quad \left. - \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^\infty; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right\| \\ & \leq \left\| \int u_{i,m}(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}; \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right. \\ & \quad \left. - \int E_{\theta_j: j > N_m} [\sup_{C^{T_m}} u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}, C^{T_m}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t)] g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right\| \end{aligned}$$

²⁶See Fudenberg and Levine (1983), Lemma 4.1 for the proof.

$$\begin{aligned}
& + \left\| \int E_{\theta_j: j > N_m} \left[\sup_{C^{T_m}} u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}, C^{T_m}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) \right] g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right. \\
& - \int u_{i,m}(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=T_m+1}^{\infty}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \left. \right\| \\
& + \left\| \int u_{i,m}(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^{T_m}, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=T_m+1}^{\infty}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right. \\
& \quad \left. - \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^{\infty}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \right\| \\
& < \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon
\end{aligned}$$

where the first inequality follows from the triangle inequality, and the second follows from (a) continuity at infinity and (b) continuity of beliefs and payoffs via ψ when $\{\sigma_{t,m}^i\}_{i,t}$ converges pointwise almost everywhere, as shown in Lemmas A.2 and A.3 and Theorem 3.2*. We then translate this into an equilibrium of the dynamic game in the manner analogous to Theorem 3.1 using Lemmas A.5 and A.6.

The demonstration of the existence of a symmetric equilibrium when $T = \infty$ follows an analogous argument, and is therefore omitted. \square

Appendix B

Proof of Proposition 4.2: At subgame H^t , suppose that $X_t^i \neq \emptyset$, $\hat{x}_t^i \geq x_t^i$, and $\hat{\theta}_i \geq \theta_i$. Suppose further that the conditional distribution over θ_{-i} at H^t is absolutely continuous, and that²⁷

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, x_t^{-i}(H^t, \theta_{-i}), C^t; \theta_i, \theta_{-i}) d\mu_t^i(\theta_{-i}|H^t, \theta_i) \\ & - \int u_i(H^t, x_t^i, x_t^{-i}(H^t, \theta_{-i}), C^t; \theta_i, \theta_{-i}) d\mu_t^i(\theta_{-i}|H^t, \theta_i) \geq 0 \end{aligned}$$

In the case that θ is affiliated, μ_t^i will be increasing in MLR in θ_i given H^t , since the conditional distribution of types θ_{-i} will be a restriction of the original distribution $f_{-i}(\cdot|\theta_i)$ to a product of intervals. In either case, by SCP and SRM, we have

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, x_t^{-i}(H^t, \theta_{-i}), C^t; \hat{\theta}_i, \theta_{-i}) d\mu_t^i(\theta_{-i}|H^t, \hat{\theta}_i) \\ & - \int u_i(H^t, x_t^i, x_t^{-i}(H^t, \theta_{-i}), C^t; \hat{\theta}_i, \theta_{-i}) d\mu_t^i(\theta_{-i}|H^t, \hat{\theta}_i) \geq 0 \end{aligned}$$

The best-reply in period t for player i will there be increasing in θ_i in the SSO by Lemma 4.1(a). The proof for the case where some type θ_j might be completely atomic is analogous²⁸, and therefore omitted. \square

Proof of Proposition 4.3: At subgame H^t , suppose that $\hat{x}_t^i \geq x_t^i$ and $\tilde{\theta}_i^t(H^t, \hat{\alpha}_i) \equiv \hat{\theta}_i > \theta_i \equiv \tilde{\theta}_i^t(H^t, \alpha_i)$.²⁹ Suppose that³⁰

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, \tilde{x}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(\hat{H}^\tau, \alpha)\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \alpha_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \\ & - \int u_i(H^t, x_t^i, \tilde{x}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \alpha)\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \alpha_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \geq 0 \end{aligned}$$

where $\{H^t, \hat{x}_t^i\} \subset \hat{H}^\tau$ and $\{H^t, x_t^i\} \subset H^\tau$ for $\tau > t$, respectively, for given action profiles by other players and in other periods. Note that for all relevant periods $\tau + 1$, it must have been that in period τ , player i played $(x_\tau^i)^*$.

Suppose that $\hat{x}_t^i \neq (x_t^i)^*$. By revealed preference, type α_i prefers to follow his continuation strategy (given by $\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i)$) after choosing x_t^i instead of that of $\hat{\alpha}_i$. Moreover,

²⁷Because C^t is irrelevant, we can substitute this into u_i without affecting the payoffs.

²⁸The only difference is that there may be some mixing by player j . Since the conditions of the proposition allow for aggregation of single-crossing as they satisfy the conditions of Lemma 4.1(b) or 4.1(c), this will not make a difference.

²⁹The proof for θ_i outside the support of $g_t(\cdot|H^t)$ is identical.

³⁰We suppress types in g_t where possible due to the independence of the distribution of θ .

by future irrelevance, α_i would have the same payoff if everyone continued by playing $\{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T$ after choosing \hat{x}_t^i . Therefore,

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \alpha)\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \alpha_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \\ & - \int u_i(H^t, x_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \alpha)\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \alpha_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \geq 0 \end{aligned}$$

By SCP and SRM, we can aggregate the single-crossing condition, yielding

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \\ & - \int u_i(H^t, x_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \geq 0 \end{aligned}$$

Lastly, by future irrelevance after \hat{x}_t^i , we replace H^τ with \hat{H}^τ , so that

$$\begin{aligned} & \int u_i(H^t, \hat{x}_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(\hat{H}^\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \\ & - \int u_i(H^t, x_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(H^t, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|H^t) d\alpha_{-i} \geq 0 \end{aligned}$$

The case where $x_t^i \neq (x_t^i)^*$ is analogous, where we first note that one can replace $\{\tilde{\mathbf{x}}_{\tau,t}(H^\tau, \alpha)\}_{\tau=t+1}^T$ with $\{\tilde{\mathbf{x}}_{\tau,t}(\hat{H}^\tau, \alpha)\}_{\tau=t+1}^T$ due to future irrelevance after x_t^i , and then invoking single-crossing in (x_t^i, θ_i) , with the argument completed by using revealed preference for the continuation after \hat{x}_t^i for $\hat{\alpha}_i$ to show that it is better than choosing x_t^i . The details are therefore omitted. \square

Proof of Proposition 4.4: Let $\hat{\theta}_i > \theta_i$. We break down our analysis by period. In period 1, we can break down the payoff of player i by H^2 , so we have

$$U_i^1(x_1^i, \theta_i) = \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(H^2, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}, \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i}$$

Since, if player i is indifferent in period 2 between various actions at a particular subgame, it does not matter which of those he chooses, we can assume without loss of optimality that $\mathbf{x}_2^i(H^2, \theta_i)$ is a singleton. By monotonicity within and across subgames in period 2 and affiliation of (x_1, θ) , it must be that $\mathbf{x}_2^i(H^2, \theta_i)$ and $\tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i})$ are increasing in all arguments. We set $\hat{H}^2 = (\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i})$ and $H^2 = (x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i})$. Suppose that

$$\int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{H}^2, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{H}^2, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i}$$

$$- \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(H^2, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \geq 0$$

For any α_{-i} , since each player's strategy is chosen independently from the others, $\hat{H}^2 > H^2$ for all α_{-i} . By monotonicity within and across subgames, it follows that $\tilde{\mathbf{x}}_{2,1}^{-i}(\hat{H}^2, \hat{\alpha}_{-i}) \geq \tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i})$ for $\hat{\alpha}_{-i} \geq \alpha_{-i}$.

A possible complication is that we do not know whether $\mathbf{x}_2^i(\hat{H}^2, \theta_i) \geq \mathbf{x}_2^i(H^2, \hat{\theta}_i)$ or vice versa. To address this, let

$$\tilde{\mathbf{x}}_2^i(\hat{H}^2, \theta_i; H^2, \hat{\theta}_i) = \mathbf{x}_2^i(\hat{H}^2, \theta_i) \vee \mathbf{x}_2^i(H^2, \hat{\theta}_i)$$

$$\hat{\mathbf{x}}_2^i(\hat{H}^2, \theta_i; H^2, \hat{\theta}_i) = \mathbf{x}_2^i(\hat{H}^2, \theta_i) \wedge \mathbf{x}_2^i(H^2, \hat{\theta}_i)$$

By revealed preference, since $\mathbf{x}_2^i(H^2, \theta_i)$ is optimal for θ_i upon reaching H^2 in period 2,

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{H}^2, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{H}^2, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \hat{\mathbf{x}}_2^i(\hat{H}^2, \theta_i; H^2, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \geq 0 \end{aligned}$$

Since one can aggregate supermodularity under integration,³¹ we have

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \tilde{\mathbf{x}}_2^i(\hat{H}^2, \theta_i; H^2, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{H}^2, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(H^2, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \geq 0 \end{aligned}$$

By revealed preference again, since $\mathbf{x}_2^i(\hat{H}^2, \hat{\theta}_i)$ is optimal for $\hat{\theta}_i$ upon reaching \hat{H}^2 ,

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{H}^2, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{H}^2, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(H^2, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(H^2, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|H^1) d\alpha_{-i} \geq 0 \end{aligned}$$

Putting all of this together, we find that

$$U_i(\hat{x}_1^i, \theta_i) - U_i(x_1^i, \theta_i) \geq 0 \implies U_i(\hat{x}_1^i, \hat{\theta}_i) - U_i(x_1^i, \theta_i) \geq 0$$

Hence the optimal action will be increasing in the strong set order in period 1, as shown in Lemma 4.1(a).

To show that best replies are monotone within and across subgames in period 2, suppose

³¹As the exogenous inequalities associated with supermodularity given by Definition 4.2 are true for every α , the endogenous inequality will be satisfied when we integrate over α .

that $\hat{x}_2^i \geq x_2^i$, and that

$$\int u_i(x_1, \hat{x}_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \theta_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i} | H^2) d\alpha_{-i} - \int u_i(x_1, x_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \theta_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i} | H^2) d\alpha_{-i} \geq 0$$

We know that beliefs are increasing in MLR in H^2 because period-1 actions are increasing in type.³² Since $\tilde{\mathbf{x}}_{2,2}^{-i}$ and $\tilde{\theta}_{-i}$ are increasing in H^2 and α (by monotonicity within and across subgames), the induced distribution of (x_2^{-i}, θ_{-i}) conditional upon observing \hat{H}^2 will first-order stochastically dominate that from upon observing H^2 . Hence we find that (by aggregating the supermodularity and ID conditions under integration)³³ that

$$\int u_i(x_1, \hat{x}_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \hat{\theta}_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i} | \hat{H}^2) d\alpha_{-i} - \int u_i(x_1, x_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \hat{\theta}_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i} | H^2) d\alpha_{-i} \geq 0$$

As in period 1, the optimal action will then be increasing in the strong set order in period 2 within and across subgames by Lemma 4.1(a). \square

³²See Milgrom (1981), Proposition 4, for the details.

³³Again, since θ_2 and $\tilde{\mathbf{x}}_{2,2}$ will be increasing in α and H^2 , for every value of α , the exogenous supermodularity inequality given in Definition 4.2 holds. Since it holds for every value of α , it must hold under integration as well.