



CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1584

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

Jeffrey Mensch

Northwestern University

October 19th, 2015

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

Jeffrey Mensch[†]

October 19, 2015

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. To address this, we define an auxiliary static game which pins down beliefs while preserving continuity of payoffs. Difficulties arise when attempting to extend to a continuum of actions due to belief entanglement, making such extensions possible only under stronger conditions. 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 signalling games and stopping games such as auctions and wars of attrition.

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' action to tend higher as well. Some well-known examples in the economic literature where these components come into play include the cheap talk model of 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 (1982); and various dynamic auctions, as analyzed in Milgrom and Weber (1982). Recent papers including models of dynamic games with such complementarities include Pavan, Segal, and Toikka (2014), Hagenbach, Koessler, and Perez-Richet (2014), Gentry and Li (2014), Lee and Liu (2013), Aradillas-López, Gandhi, and Quint (2013), and Back and Baruch (2013). Hence a general equilibrium existence result for such games would be of major significance across a wide array of economic topics. This paper provides conditions under which an equilibrium in strategies that are monotone in types within each subgame is guaranteed to exist in dynamic games.

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), 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) and Vives (2009) 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), 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.

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.

Fortunately, the monotone structure of the strategies guarantees that the posterior beliefs will be restrictions of the prior to a product of intervals of types. This allows for the construction of a continuous transformation of beliefs into an auxiliary game in a static environment in which players choose continuation strategies from all possible future subgames (while one might be tempted to instead using backward induction, this will lead to complications in that the equilibria in the subgames will not be sufficiently well-behaved; we discuss this in Appendix B). We will thus be able to break down the strategies of the players by period and informational events, and show that a small perturbation of the strategies of the players lead to a continuous perturbation of the beliefs of other players. We then show that equilibrium exists in this static transformation of the game from the existence results of Reny (2011). Finally, we use the equilibrium strategies that were found in the static transformation of the game to derive monotone strategies that will form a perfect Bayesian equilibrium in the original dynamic game.

An interesting feature of the derivation of the existence of equilibrium is that it 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 with the support conditional on a lower action. 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.

We also consider extensions to games with a continuum of actions. Here, a new issue

arises. In static games, the standard method of extending the results to a continuum of actions, as used in Athey (2001) or McAdams (2003), is to use Helly’s selection theorem on successively finer approximations of the continuum game to show that there is a sequence of equilibrium strategy profiles which converges to an equilibrium in the limit game. This method does not work directly in a dynamic context, because there is a possibility of “belief entanglement,” leading to different information sets in later periods in the limit game from any of the finite approximations.¹ We show that this is the only possible issue, so that when belief entanglement is avoided, there will exist an equilibrium in the limit game. We also provide conditions which guarantee that belief entanglement does not occur.

A final point of difficulty is the characterization of single-crossing conditions in dynamic games. The existence result described above assumes the existence of monotone best-replies. To guarantee that such-best replies exist, one needs a single-crossing condition. Yet when one assume sequential rationality, such conditions are difficult to obtain. As we show in an example, not even supermodularity of ex-post payoffs between actions and types 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 garner better information on the types of 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 action 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 signalling games under fairly general conditions, including a sender with a supermodular payoff in the message and his type, and multiple receivers with private information, and derive additional properties about pooling and separation of types that must hold in any such equilibrium.

¹This is similar to the phenomenon of “strategic entanglement” found in Myerson and Reny (2015)

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 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 will proceed as follows. Section 2 describes the model of the games considered in this paper. Section 3 examines difficulties in extending existing results to dynamic games, providing examples that illustrate each of these issues. Section 4 provides topological conditions on the beliefs induced by monotone strategies which are sufficient to guarantee existence of monotone equilibrium. Section 5 explores extensions of the existence theorems to a continuum of actions. Section 6 provides 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 7 provides several applications of the various results found throughout this paper to signalling games and to stopping games.

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 . 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 .²³

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.

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 of the 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

We will provide an extension to a continuum of actions in Section 5; further complications may arise, so stronger conditions are necessary. Define $X = \prod_{t,i} X_t^i$ and $\Theta = \prod_i \Theta_i$. The joint density over types is given by $f(\cdot)$, which we assume (a) is bounded, (b) has full support on Θ , and (c) is 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 $C^t \in C^t$. We allow for imperfect observability, stipulating that at any given history H^t , each player i observes some event conveying information about other players' past actions. We thus define the random variable

$$Q_{it} : \mathcal{H}^t \rightarrow Q_{it}$$

$$H^t \rightarrow q_{it}$$

as a function of H^t , where $Q_{it} \subset \mathbb{R}^k$ is a finite lattice, and a given H^t generates some q_{it} with probability 1. We assume that $Q_{it}(H^t)$ perfectly reveals past play of player i , as well as $Q_{i,t-1}(H^{t-1})$, so there is no issue of imperfect recall. Let $Q_t = \prod_i Q_{it}$, and $Q = \prod_t Q_t$.

Let \mathcal{Y} be the Borel σ -algebra of measurable subsets $Y \subset \Theta \times X$. Similarly define \mathcal{Y}_t to be the Borel σ -algebra for $Y \subset \Theta \times H^t$. 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 . We denote conditional beliefs in each period for each player i by

$$\mu_t^i : \mathcal{Y}_t \times Q_{it} \times \Theta_i \rightarrow \mathcal{M}_t^i$$

$$(Y_t, q_{it}, \theta_i) \rightarrow \mu_t^i(Y|q_{it}, \theta_i)$$

for any possible information set q_{it} . We let $\mu_t(Y_t|q_t, \theta) = (\mu_t^1, \dots, \mu_t^N)$ and $\mu(Y|q, \theta) = (\mu_1, \dots, \mu_T)$.

We now define behavioral strategies for player i . Define the conditional probability that player i chooses $x_t^i \in X_t^i$ by $\rho_t^i(x_t^i|q_{it}, \theta_i) \in \Delta(X_t^i)$, where ρ_t^i is measurable with respect to Θ_i . This induces a *strategy correspondence* $\mathbf{x}_t^i : Q_{it} \times \Theta_i \rightrightarrows X_t^i$ which represents the actions chosen with positive probability out of X_t^i . Note that this is not inherently a best reply, as

never chosen in equilibrium.

this definition 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

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

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

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

where $Q_{i,t+1}$ is the informational event generated by $\{H^t, x_t^i, x_t^{-i}\}$. The objective of player i in each period is to maximize $U_i^t(q_{it}, \cdot, \theta_i)$ 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* $BR_t^i : Q_{it} \times \Theta_i \rightarrow X_t^i$ as the subset of actions such that, given μ_t^i and θ_i , $U_i^t(q_{it}, x_t^i, \theta_i) \geq U_i^t(q_{it}, \hat{x}_t^i, \theta_i)$, $\forall \hat{x}_t^i \in X_t^i$.

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 $\mathbf{x}_t^i(q_{it}, \theta_i)$ is a singleton and

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

$$\hat{q}_{it} \geq q_{it} \implies \mathbf{x}_t^i(\hat{q}_{it}, \theta_i) \geq \mathbf{x}_t^i(q_{it}, \theta_i)$$

$\mathbf{x}_t^i(\cdot, \cdot)$ is *monotonic in mixed strategies within subgames* if

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

$\mathbf{x}_t^i(\cdot, \cdot)$ is *monotonic in mixed strategies across subgames* if the induced distribution of play over $x_t^i \in \mathbf{x}_t^i(\hat{q}_{it}, \theta_i)$, given by ρ_t^i , first-order stochastically dominates (FOSD) that over $\mathbf{x}_t^i(q_{it}, \theta_i)$ for $\hat{q}_{it} \geq q_{it}$. It is immediately clear that any strategy correspondence $\mathbf{x}_t^i(\cdot, \cdot)$ that is monotonic in pure strategies is monotonic in mixed strategies. 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,⁴ 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 informational event that is reached with positive probability from a given strategy profile after H^{t-1} is reached. Thus, for any measurable $\Psi_{-i} \subset \Theta_{-i}$,

$$\mu_t^i(\Psi_{-i}, H^t | q_{it}, \theta_i) = \frac{\mu(\Theta_{-i}, H^{t-1} | q_{i,t-1}, \theta_i) \int_{\Psi_{-i}} \prod_{j \neq i} p_{t-1}^j(x_{t-1}^j | Q_{j,t-1}, \theta_j) d\mu_t^i(\theta_{-i}, H^{t-1} | q_{i,t-1}, \theta_i)}{\sum \mu(\Theta_{-i}, \hat{H}^{t-1} | q_{i,t-1}, \theta_i) \int_{\Theta_{-i}} \prod_{j \neq i} p_{t-1}^j(\hat{x}_{t-1}^j | Q_{j,t-1}, \theta_j) d\mu_t^i(\theta_{-i}, \hat{H}^{t-1} | q_{i,t-1}, \theta_i)}$$

where $H^t = \{H^{t-1}, x_{t-1}^i, x_{t-1}^{-i}\}$ and $Q_t(H^t) = q_t$ (analogously, $Q_{t-1}(H^{t-1}) = q_{t-1}$), and the denominator is summed over $\{\hat{H}^{t-1}, \hat{x}_{t-1}^{-i}\}$ such that $Q_{it}(\{\hat{H}^{t-1}, \hat{x}_{t-1}^{-i}\}) = q_{it}$. 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 $F_{-i}(\theta_{-i} | H^t, \theta_i)$.⁵

To extend to off-path informational events, suppose that the conditional distribution of types at H^{t-1} is $F(\theta | H^{t-1})$. As H^{t-1} generates some q_{t-1} , player i has conditional belief $\mu_{t-1}^i(\theta_{-i}, H^{t-1} | q_{i,t-1}, \theta_i)$. Suppose that player j deviates to some off-path action x_{t-1}^j in period $t-1$. Then the following properties must be satisfied:

- (a) The support of $F_i(\cdot | \{H^{t-1}, x_{t-1}^j\})$ must be a subset of that of $F_i(\cdot | H^{t-1})$.
- (b) For any subset of players I , the distribution over the vector of the other players' types, θ_{-I} , is independent of the actions taken by players $i \in I$, holding their types fixed:

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

- (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),

⁴Although 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).

⁵Indeed, if $Q_{it}(H^t) = H^t$, then $dF_{-i}(\theta_{-i} | H^t, \theta_i) = d\mu_t^i(\theta_{-i}, H^t | Q_{it}(H^t), \theta_i)$

then for any x_{t-1}, \hat{x}_{t-1} ,

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

whenever these ratios are well-defined; if positive probability is placed on some θ_j at q_{it} , 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 being chosen must be the same for type θ_j , regardless of what other players do.

The last two conditions above are analogous to the “no signalling what you don’t know” condition of Fudenberg and Tirole (1991). When proving our existence result, we will show conditions under which these properties can be established.

A *perfect Bayesian equilibrium* (henceforth PBE), then, is a vector $(\mathbf{x}(\cdot, \cdot), \mu)$, where $\mathbf{x}(\cdot, \cdot)$ is the strategy profile and μ is the belief profile, in which beliefs μ satisfy Bayes’ rule on-path, properties (a)-(c) hold at all informational events, and at each q_t , the continuation strategies form a Bayesian Nash equilibrium given beliefs μ_t , i.e. if $\hat{x}_t^i \in \mathbf{x}_t^i(q_{it}, \theta_i) \subset X_t^i$, then $\hat{x}_t^i \in \arg \max_{x_t^i \in X_t^i} U_t^i(q_{it}, \cdot, \theta_i)$.

3 Difficulties in Extending Previous Results

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 3.1 (Reny, Proposition 4.4): *Suppose that the following four conditions hold:*

- (i) $X \times \Theta$ forms a lattice;
- (ii) F is atomless;
- (iii) Payoffs u_i are bounded, measurable in x and θ , and continuous in x ;⁶ 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

⁶Reny 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.

$x_i \vee \hat{x}_i$).

Then there exists a monotone pure-strategy equilibrium.

However, when attempting to extend this result to games with multiple periods, one runs into three potential issues that do not appear in static games, summarized below.

3.1 Off-path beliefs

One needs to ensure that off-path beliefs are defined so as to maintain continuity of payoffs in the strategy profile chosen. Continuity is necessary to ensure upper-hemicontinuity of best-replies in the strategies of one's opponents. This will require additional restrictions on which beliefs are possible off-path. If beliefs are instead defined in such a way so that beliefs off-path “jump around” even for minor changes in the overall strategy profile, then these conditions may be violated. The intuition is simple: if there is a discrete change in the distribution over types, incentives can change discretely. These effects are illustrated in the following example.

Example 3.1: We present a dynamic game with supermodular payoffs in which, given a certain specification of off-path beliefs, the best-reply correspondence may not be upper-hemicontinuous. Let $\theta_1 \sim U[0, 1]$, $X_1^1 = \{0, 1\}$, and $X_2^1 = X_2^2 = [0, 1]$. Suppose that there are two players whose payoff functions are given by $u_1(x, \theta) = 4x_2^2\theta_1 + x_1^1(\theta_1 + \frac{1}{2}) - (x_2^1 - \theta_1)^2$ and $u_2(x, \theta) = -(x_2^2 - x_2^1)^2$, respectively. Assume that period 1 actions are perfectly observable. Note that player 1 always chooses $x_2^1 = \theta_1$, and so player 2 chooses $x_2^2 = E[\theta_1 | x_1^1]$.

In a static environment, payoffs necessarily converge in expectation when strategy profiles converge pointwise almost-everywhere, leading to upper-hemicontinuity of best-replies in the strategy profiles chosen by other players. This will not be the case in a dynamic environment without additional specification. Consider a sequence of strategy profiles $x_1^1(\theta_1)$ in which, for some sequence $\{\alpha_k\}$ such that $\alpha_k \in (0, 1)$ and $\lim_{k \rightarrow \infty} \alpha_k = 1$, $x_1^1(\theta_1) = 0$ if $\theta_1 \leq \alpha_k$ and $x_1^1(\theta_1) = 1$ otherwise. Using Bayesian updating, player 2 believes that $\theta_1 \in [0, \alpha_k]$ chooses $x_1 = 0$, while $\theta_1 \in (\alpha_k, 1]$ chooses $x_1^1 = 1$. Thus $E[\theta_1 | x_1^1 = 1] > E[\theta_1 | x_1^1 = 0]$ for each such k . However, in the limit (under pointwise almost-everywhere convergence), player 1 chooses $x_1^1 = 0$ regardless of type. Suppose that the beliefs of player 2 induced by this limit strategy profile are such that all types of θ_1 choose $x_1^1 = 0$, but conditional upon observing $x_1^1 = 1$, player 2 believes that $\theta_1 = 0$ with probability 1. This will lead to a drastic reversal of choice of actions for player 2 in period 2, as now $E[\theta_1 | x_1^1 = 1] = 0$, whereas

the limit of the sequence of the values of x_2^2 chosen will be $\frac{1}{2}$. Hence a small change in the strategy profile will lead to large changes in incentives of the players, violating upper-hemicontinuity of best-replies. \square

The above example suggests that we will require some continuity conditions on the beliefs of the players. Indeed, such conditions are provided in Section 4.

3.2 Continuum of actions

While the model as described in Section 2 posits only a finite number of actions available to each player i in each period t , many applications 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 to generalizing existence results for monotone equilibrium in static games from those with discrete actions to those with a continuum of actions has been to use Helly’s selection theorem (Kolmogorov and Fomin, p. 373). For example, both Athey (2001) and McAdams (2003) showed that, by taking successively finer partitions of the strategy space, and taking the sequence of equilibria of this sequence of games, and picking a convergent subsequence of equilibrium strategy profiles by Helly’s selection theorem we obtain strategy functions that are monotone. These limit functions turn out to be best responses to the limit strategy functions of the other players as well, and hence a Bayesian Nash equilibrium.

In our case, however, a naive application of this approach leads “belief entanglement,” analogous to what Myerson and Reny (2014) refer to as “strategic entanglement” in the limit in periods $t \leq T - 1$, which may preclude the generation of equilibrium by this method. Intuitively, this issue arises from an induced discontinuity in beliefs, as one may have some types choosing distinct actions for every element of the sequence, but the same action in the limit. To illustrate this, we consider a sequence of monotone strategy profiles in a two-period environment in which this effect occurs.

Example 3.2: We show that the limit of the beliefs in the weak-* topology from the sequence of monotone strategy profiles may not be the same as the beliefs induced by the limit of the strategy profiles under pointwise almost-everywhere convergence. Consider a sequence of games $\{\Gamma_m\}_{m=1}^\infty$ with two periods and two players in which for each m , the set of period 1 actions for player 1 is $X_{1,m}^1 = \{\frac{k}{2^m}\}_{k=0}^{2^m}$. In the limit as $m \rightarrow \infty$, $X_{1,m}^1$ is dense on the interval $[0, 1]$, so let $X_{1,\infty}^1 = [0, 1]$. Let $\theta_1 \sim U[0, 1]$. Assume that actions in period 1 are perfectly observable. The exact specification of the rest of the game is irrelevant for the




Figure 1: Belief entanglement

purposes of this example.

Consider the sequence of monotone strategy profiles (illustrated in Figure 1) in which $x_{1,m}^1(\theta_1) = \frac{2^{m-1}-1}{2^m}$ for $\theta_1 \in [0, \frac{1}{3})$, $\frac{1}{2}$ for $\theta_1 \in [\frac{1}{3}, \frac{2}{3}]$, and $\frac{2^{m-1}+1}{2^m}$ otherwise.

Upon observing $x_1^1 = \frac{1}{2}$ in any game Γ_m , player 2 concludes in period 2 that $\mu_{2,m}^2(\{\theta_1 \in [\frac{1}{3}, \frac{2}{3}]\} | x_1^1 = \frac{1}{2}) = 1$. However, if we take the limit of the strategy profiles $x_{1,m}^1(\cdot)$, we find that all types choose $x_{1,\infty}^1(\theta_1) = \frac{1}{2}$. Player 2 then concludes that $\mu_2^2(\{\theta_1 \in [\frac{1}{3}, \frac{2}{3}]\} | x_1^1 = \frac{1}{2}) = \frac{1}{3}$, which is not the limit of the beliefs (in the weak-* topology) of the beliefs in Γ_m . \square

As seen through this example, we will need a more sophisticated approach than a simple application of Helly's selection theorem to generate a monotone equilibrium in the game with a continuum of actions in order to circumvent this issue. While we cannot show that it is always possible to get around the potential problems raised here, in Section 5 we are able to provide fairly general conditions under which it will be possible to extend our results from finite to continuous actions.

3.3 Monotonicity of best replies

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.⁷

Example 3.3: 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) . Assume that all actions are observable.

Proposition 3.2: *There does not exist a monotone PBE of the game described in Example 3.3.*

All proofs are relegated to Appendix A. 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 must be that the conditional beliefs over Θ_2 that ensue 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, preventing the existence of a monotone PBE.

While it will not be possible to provide conditions that are as general as those found in static models, it will be possible to guarantee existence of monotone best-replies to monotone strategies by the other players under stronger conditions. Some such conditions are provided in Section 6.

⁷Echenique (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.

4 Topological Conditions for Finite Actions

While one would like to be able use Reny's existence results in dynamic environments, his results (as written) apply only to static games. In this section, we will show that one can apply his results by translating the dynamic game into an appropriate static game in which players choose continuation strategies. To do so, we must explore the topological conditions necessary to guarantee that, in any game in which the action set is finite, if the best-reply to monotone strategies is increasing in the strong set order in own type, a monotone equilibrium will exist. This involves imposing certain conditions on the beliefs of the players that are induced by monotone strategies, and splitting each player's decisions by period. Throughout this section, we restrict our attention to the case where all players use monotone (mixed) strategies within subgames, and best-replies are increasing in the strong set order in each player's own type; the latter condition ensures that some monotone strategy will be optimal.

We divide the proof into three major steps. In the first step, we translate the dynamic game as defined in Section 2 into an auxiliary static game. In the second, we show that the preferences, so defined, meet the criteria of Theorem 3.1, and so we can invoke Reny's theorem to show that there exists a monotone equilibrium in the auxiliary game. In the third, we use the equilibrium from the auxiliary game to construct a PBE in the dynamic game. Once we have done so, we will provide extensions to symmetric games and infinite-horizon games.

4.1 Translation to a static environment

In order to reinterpret the game as a static one, will need to break down the players by q_{it} . We define the auxiliary game Γ^1 in which player i at each q_{it} is considered a distinct player. Player i at each q_{it} will then choose not only what he does at q_{it} , but what he plans to do at all subsequent subgames; thus the action space for player i is $\prod_{\tau=t}^T (X_\tau^i)^{|\{q_{i\tau}: q_{it} \subset q_{i\tau}\}|}$. 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 $Q_{it}(H^t)$, which will depend endogenously on which strategy is chosen in earlier periods. Thus θ_i has some distribution conditional on q_{it} ; moreover, there will be a joint distribution

of θ_{-i} conditional on (q_t, θ_i) .⁸

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 4.1: The distribution over types θ_j is *completely atomic* if it places probability 1 on $\theta_j = \theta_j^*$ for some $\theta_j^* \in \Theta_j$.

Lemma 4.1: Suppose that each player i chooses a fixed monotone strategy in each period. Then for any H^t that is on-path, the conditional distribution over θ is completely atomic over some (possibly empty) subset I of players and absolutely continuous for $i \notin I$, with conditional density equal to the prior restricted to an interval $[\theta_i^1, \theta_i^2] \subset \Theta_i$.

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 4.1 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 4.1, 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 outside observer, by Lemma 4.1, at any q_t (which pins down H^t), the distribution over the types of all players will have support over a Cartesian product of intervals. Specifically, as the past actions of any player i are completely observable through q_{it} , from the point of view of player i , this interval for types θ_i is pinned down just by q_{it} due to independence of $F_i(\theta_i | \{H^{t-1}, x_{t-1}^i\}, \theta_{-i})$ from the actions of other players, x_{t-1}^{-i} . 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

⁸Since all players observe their own past actions, q_t (the vector of all players' informational events) uniquely pins down H^t , and so we need not condition on H^t .

space for player i indexed at q_{it} 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 , we define

$$\tilde{\theta}_i' : Q_{it} \times [0, 1] \rightarrow \Theta_i$$

$$(q_{it}, \alpha_i) \rightarrow \theta_i$$

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 q_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 $G_t(\alpha|q_t)$ is given by the density function

$$g_t(\alpha|q_t) = \frac{f(\tilde{\theta}'(q_t, \alpha)) \int_A d\theta}{\int_A f(\theta) d\theta}$$

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

$$\frac{g_t(\alpha'|q_t)}{g_t(\alpha|q_t)} = \frac{f(\tilde{\theta}'(q_t, \alpha'))}{f(\tilde{\theta}'(q_t, \alpha))} \quad (1)$$

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 (1) 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, and that the limiting case of completely atomic type θ_j will generate a uniform marginal distribution over α_j . 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, 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.⁹

To see how the actions chosen as a function of type 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

⁹This 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.

at q_{it} is $[\theta_i^1, \theta_i^2]$, so that $\tilde{\theta}_i^t(q_{it}, 0) = \theta_i^1$ and $\tilde{\theta}_i^t(q_{it}, 1) = \theta_i^2$. 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 $q_{i\tau}$, where $\tau \geq t$. Since we have broken down each player i according to each q_{it} , one can represent the strategies of each player i in each period $\tau \geq t$ (conditional on reaching $q_{i\tau}$) 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 informational events are reached) according to this monotone strategy by the function

$$\begin{aligned} \tilde{\mathbf{x}}_{\tau,t}^i : Q_{i\tau} \times [0, 1] &\rightarrow X_\tau^i \\ (q_{i\tau}, \alpha_i) &\rightarrow x_\tau^i \end{aligned}$$

At each H^τ , with players choosing strategies according to $\tilde{\mathbf{x}}_{\tau,t}^i$, we generate a new history $H^{\tau+1} = \{H^\tau, \{\tilde{\mathbf{x}}_{\tau,t}^i(Q_{i\tau}, \alpha_i)\}_i\}$, which in turn generates $Q_{\tau+1}$. 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 $\{\tilde{\mathbf{x}}_{\tau,t}^{-i}, \tilde{\theta}_{-i}^t\}$, the expression for the interim payoff of player i conditional on being type α_i and choosing actions $\{\tilde{\mathbf{x}}_{\tau,t}^i(Q_{i\tau}, \alpha_i)\}_{\tau=t}^T$ can now be written as (suppressing arguments for $\tilde{\theta}_i^t$ and $\tilde{\theta}_{-i}^t$)¹⁰

$$\int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i(Q_{i\tau}, \alpha_i), \tilde{\mathbf{x}}_{\tau,t}^{-i}(Q_{-i\tau}, \alpha_{-i})\}_{\tau=t}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | Q_t, \alpha_i) d\alpha_{-i} d\mu_t^i(\Theta, H^t | Q_{it}, \tilde{\theta}_i^t(Q_{it}, \alpha_i)) \quad (2)$$

We now turn to the relationship between the actions chosen by i in period τ , and what he plans to do in τ as of period t . Since we are considering player i at each q_{it} 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 q_τ is reached will be the actual action taken by player i with type θ_i indexed by $q_{i\tau}$. The potential concern is that some players may plan to do something in the future that they would never actually want to do if this subgame were actually reached in period τ ; however, if this subgame occurs with probability 0 because other players never take the actions necessary to reach that subgame, then their payoff is unaffected by this sequentially irrational choice. This, in turn, may affect the payoffs of other players, who may be

¹⁰To 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 $d\mu_t^i$. The arguments $Q_\tau \in Q_\tau$ for each $\tilde{\mathbf{x}}_{\tau,t}^i$ and $\tilde{\mathbf{x}}_{\tau,t}^{-i}$ are simply the information 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 .

concerned about reaching these subgames, and so indeed lead to these subgames occurring with probability 0. As this is incompatible with sequential rationality, we must rule this out in any PBE construction. To do this, we first need to define a notion of reachability.

Definition 4.2: The event $q_{i\tau}$ is *reachable* from q_{it} for some sequence of actions $\{x_{t'}^i\}_{t'=t}^{\tau-1}$ if there exist some history H^t and action profile $\{x_{t'}^i\}_{t'=t}^{\tau-1}$ such that $Q_{it}(H^t) = q_{it}$ and $Q_{i\tau}(H^t, \{x_{t'}^i\}_{t'=t}^{\tau-1}) = q_{i\tau}$.

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

$$\frac{\int 1_{\{(\tilde{\theta}_t^i, \tilde{\theta}_{-i}^t, H^t, \{\tilde{\mathbf{x}}_{t',t}^i, \tilde{\mathbf{x}}_{t',t}^{-i}\}_{t'=t}^{\tau-1}, \tilde{\mathbf{x}}_{\tau,t}^i) \in (A, H^\tau, x_\tau^i)\}} (\alpha_{-i}) g_t(\alpha_{-i} | q_{it}, \alpha_i) d\alpha_{-i}}{\int 1_{\{(\tilde{\theta}_t^i, \tilde{\theta}_{-i}^t, H^t, \{\tilde{\mathbf{x}}_{t',t}^i, \tilde{\mathbf{x}}_{t',t}^{-i}\}_{t'=t}^{\tau-1}) \in (\Theta, H^\tau)\}} (\alpha_{-i}) g_t(\alpha_{-i} | q_{it}, \alpha_i) d\alpha_{-i}} \\ = \int 1_{\{(\tilde{\theta}_t^i, \tilde{\theta}_{-i}^t, H^\tau, \tilde{\mathbf{x}}_{\tau,\tau}^i) \in (A, H^\tau, x_\tau^i)\}} (\alpha_{-i}) g_\tau(\alpha_{-i} | q_{i\tau}, \hat{\alpha}_i) d\alpha_{-i} \quad (3)$$

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

To approach this problem, we define an additional 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 $(\{\tilde{\mathbf{x}}_{t',t}^j\}_{t'=t}^T, \tilde{\theta}_t^j)$ for all j in Γ^1 , and consider the problem from the perspective of player i at informational event q_{it} . Suppose that type α_i chooses strategy $\{\tilde{\mathbf{x}}_{t',t}^i\}_{t'=t}^T$. We consider the following strategy instead at $q_{i\tau}$ from the perspective of period t . Suppose that $q_{i,\tau-1} \subset q_{i\tau}$; since $q_{i\tau}$ completely reveals both $\{q_{it'}\}_{t'=t}^{\tau-1}$ and $\{x_{t'}^i\}_{t'=t}^{\tau-1}$ due to perfect recall, there will be a unique vector $\{x_{t'}^i\}_{t'=t}^{\tau-1} \in X_{\tau-1}^i$ which can reach $q_{i\tau}$ conditional on $\{q_{it'}\}_{t'=t}^{\tau-1}$ being reached. We thus define the alternative strategy $\hat{\mathbf{x}}_{\tau,t}^i$ inductively, starting from period t and working forward to period T , and setting $\hat{\mathbf{x}}_{t,t}^i = \tilde{\mathbf{x}}_{t,t}^i$. Let $\underline{\alpha}_i(q_{i\tau}, q_{it})$ be the infimum of the set of types α_i such that α_i chooses $\{x_{t'}^i\}_{t'=t}^{\tau-1}$ according to $\{\hat{\mathbf{x}}_{t',t}^i\}_{t'=t}^{\tau-1}$ at the respective $\{q_{it'}\}_{t'=t}^{\tau-1}$, and $\bar{\alpha}_i(q_{i\tau}, q_{it})$ be the supremum of that set of types. Then

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




Figure 2: Construction of $\hat{x}_{\tau,t}^i$

Similarly, we define the vector $\{\hat{\mathbf{x}}_{\tau,t}^{-i}\} = \{\hat{\mathbf{x}}_{\tau,t}^j\}_{j \neq i}$. We now define the payoff function for Γ^2 so that the interim payoff function given q_{it} , where Q_τ is generated the same way as in (2), replacing $\tilde{\mathbf{x}}_{\tau,t}^{-i}(Q_{-i\tau}, \alpha_{-i})$ with $\hat{\mathbf{x}}_{\tau,t}^{-i}(Q_{-i\tau}, \alpha_{-i})$. 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 $\tilde{\theta}^t$)

$$\int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^T; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | q_t, \alpha_i) d\alpha_{-i} d\mu_t^i(\Theta, H^t | q_{it}, \tilde{\theta}_i^t(q_{it}, \alpha_i)) \quad (4)$$

This is done to ensure that the payoff of player i as indexed by q_{it} will be based on what players $-i$ will *actually* do when $q_{-i\tau}$ 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 (4), and from this to construct an equilibrium of the original game, in which payoffs are given by (2), which involves consistent strategies. 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 $q_{i\tau}$ is reached with positive probability, by construction. Hence the set of monotone best-replies of to $\{\hat{\mathbf{x}}_{\tau,t}^{-i}\}$ will still be non-empty and join-closed. We illustrate this in the following example.

Example 4.1: To illustrate what $\hat{\mathbf{x}}_{\tau,t}^i$ looks like, consider Figure 2. Suppose that $\hat{\mathbf{x}}_{\tau,t}^i$ is given

by Figure 2(a), in which

$$\hat{\mathbf{x}}_{t,t}^i(q_{it}, \alpha_i) = \tilde{\mathbf{x}}_{t,t}^i(Q_{it}, \alpha_i) = \begin{cases} 1, & \alpha_i \in [0, 0.4) \\ 2, & \alpha_i \in [0.4, 0.6) \\ 3, & \alpha_i \in [0.6, 1] \end{cases}$$

Suppose that $q_{i,t+1}$ can only be reached by a choice of $x_t^i = 1$ given q_{it} . From the perspective of period t , we are given that if type α_i were to reach $q_{i,t+1}$, then (as seen in Figure 2(b))

$$\tilde{\mathbf{x}}_{t+1,t}^i(q_{i,t+1}, \alpha_i) = \begin{cases} 1, & \alpha_i \in [0, 0.2) \\ 2, & \alpha_i \in [0.2, 0.8) \\ 3, & \alpha_i \in [0.8, 1] \end{cases}$$

However, we do not yet know that $\tilde{\mathbf{x}}_{t+1,t}^i$ and $\tilde{\mathbf{x}}_{t+1,t+1}^i$ are consistent. It may be that the best replies from the perspective of $q_{i,t+1}$ of player i would be to choose (as seen in Figure 2(c))

$$\tilde{\mathbf{x}}_{t+1,t+1}^i(q_{i,t+1}, \alpha_i) = \begin{cases} 2, & \alpha_i \in [0, 0.5) \\ 3, & \alpha_i \in [0.5, 1] \end{cases}$$

From this, we construct $\hat{\mathbf{x}}_{t+1,t}^i$ so that we “shrink” $\tilde{\mathbf{x}}_{t+1,t+1}^i$ to the interval of types α_i that could reach $q_{i,t+1}$ from their choice of $x_{t,t}^i$ via $\hat{\mathbf{x}}_{t,t}^i$. The interval of types that does so is $[0, 0.4)$, so, as seen in Figure 2(d),

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

Note that this is no longer a monotone strategy function. However, the portion that is observed on path from q_{it} is only that following a choice of $x_t^i = 1$, and the set of types $\{\alpha_i\}$ which choose this is $[0, 0.4)$. When we restrict our attention to this set $\{\alpha_i\}$, $\hat{\mathbf{x}}_{t+1,t}^i$ is still monotonic in α_i over the interval $[0, 0.4)$. Hence $\hat{\mathbf{x}}_{t+1,t}^i$ is monotonic on path, and so any best reply by other players to $\hat{\mathbf{x}}_{t+1,t}^i$ treats it as if it were a monotone strategy; this implies that the set of monotone best-replies will remain non-empty and join-closed. \square

An important feature of $\hat{\mathbf{x}}_{\tau,t}^i$ is that, as defined, it still may not be consistent, as we have constructed it independently of the values of $\tilde{\theta}_i^t$ and $\tilde{\theta}_i^\tau$; thus we do not know simply from

the definition that $\tilde{\theta}_i^\tau(q_{i\tau}, 0) = \tilde{\theta}_i^\tau(q_{i\tau}, \underline{\alpha}_i(q_{i\tau}, q_{i\tau}))$ and $\tilde{\theta}_i^\tau(Q_{i\tau}, 1) = \tilde{\theta}_i^\tau(q_{i\tau}, \bar{\alpha}_i(q_{i\tau}, q_{i\tau}))$ since we have not yet explicitly connected $\tilde{\theta}_i^\tau$ to Bayesian updating. Unsurprisingly, it will turn out in equilibrium that $\hat{\mathbf{x}}_{\tau,t}^i$ will be consistent, as the way that $\tilde{\theta}_i^\tau$ is determined will ensure that it aligns with $\tilde{\theta}_i^\tau$ as above.

This approach has several advantages. First, it allows for monotone mixed strategies. Suppose that the conditional distribution of θ_j is completely atomic at q_t , so that $\mu_t^j(\theta_j^*, \Theta_{-j}, H^t | q_{it}, \theta_i) = 1$ for some $\theta_j^* \in \Theta_j$. Then for some $q_{j\tau}$, it will be possible that $\hat{\mathbf{x}}_{\tau,t}^j$ will vary with α_j , while $\tilde{\theta}_j^\tau$ 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 j at each q_{jt} as a distinct player in an essentially static game, with the set of types for each player distributed over $[0, 1]$.

4.2 Equilibrium existence in the static environment

Having gone through the above transformation of strategies and payoffs, it will be possible to view the collection $(\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T, \tilde{\theta}_i^\tau)$ as the strategy chosen by player i at informational event q_{it} . It should be emphasized that $\tilde{\theta}_i^\tau$ 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, $\sigma_t^i = (\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{\tau \geq t}, \tilde{\theta}_i^\tau)$, which takes as arguments $(\{q_{i\tau}\}_{\tau \geq t}, \alpha_i)$. Define the space of such functions σ_t^i as Σ_t^i . Indicating the Euclidean metric over $\prod_{\tau=t}^T (X_\tau^i)^{|\{q_{i\tau}: q_{i\tau} \subset q_{i\tau}\}|} \times \Theta_i$ by d_t^i ,¹¹ define a metric over Σ_t^i by (for fixed q_{it})

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

Define d_t^{-i} and δ_t^{-i} over the strategy spaces of other players $-i$ analogously.¹²

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 (4) will be continuous in the strategy profile $\{\sigma_\tau^j\}_{\tau,j}$ for player i indexed by q_{it} .¹³ We first show that $\hat{\mathbf{x}}_{\tau,t}^j$ is continuous in

¹¹Note that we must count X_τ^i once for each possible information event $q_{i\tau}$ that is reachable from q_{it} .

¹²I.e. the \mathcal{L}^1 metric.

¹³As discussed in the footnote by Theorem 3.1, 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

$$\{\tilde{\mathbf{x}}_{\tau,t}^j\}_{\tau=1}^T.$$

Lemma 4.2: 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 , $\hat{\mathbf{x}}_{\tau,t,m}^j \rightarrow \hat{\mathbf{x}}_{\tau,t}^j$.

We must also argue that the beliefs vary continuously as well in the strategies chosen by all players indexed by q_{it} , 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 4.3: Consider a sequence of beliefs conditional on (q_{it}, α_i) , $\{\mu_{t,k}^i\}_{k=1}^\infty$ that converges to some μ_t^i (for the same restriction) in the weak-* topology, where all $\mu_{t,k}^i$ and μ_t^i satisfy the conditions of Lemma 4.1. 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 $\tilde{\theta}_{j,k}^t(q_{jt}, \cdot) \rightarrow \tilde{\theta}_j^t(q_{jt}, \cdot)$ and $g_{t,k}(\cdot|q_t, \cdot) \rightarrow g_t(\cdot|q_t, \cdot)$ everywhere.

We now show the other direction, i.e. that the beliefs in period t conditional on observing q_{it} vary continuously in the past 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 event q_{it} coincide with those generated by the strategy profile σ (specifically, $\{\tilde{\mathbf{x}}_{\tau,\tau}\}_{\tau=1}^{t-1}$) and the structure by which q_{it} is generated from H^t , via Bayesian consistency whenever possible (as defined in Section 2). Specifically, the beliefs at any H^t (which pins down Q_{jt} and Q_{it}) over θ given H^t will be those inductively pinned down by restricting the prior to corresponding θ_j the intervals $\{(\underline{\alpha}_j(q_{jt}, q_{j1}), \bar{\alpha}_j(q_{jt}, q_{j1}))\}$ as given from when we originally defined $\{\{\tilde{\mathbf{x}}_{\tau,1}^j\}_{\tau=1}^{t-1}\}_j$, which were themselves defined (on-path) by $\{\tilde{\mathbf{x}}_{\tau,\tau}^j\}_{j,\tau}$.

For actions by player j that generate off-path q_{it} , 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 4.1. It will turn out that such ψ exist in many general circumstances.

Definition 4.3: A belief mapping ψ is *continuous* if player i 's belief at q_{it} , given by $(\psi(\sigma_m(\cdot)))_i^i(\cdot|q_{it}, \theta_i)$, converges in the weak-* topology to $(\psi(\sigma(\cdot)))_i^i(\cdot|q_{it}, \theta_i)$ for any sequence of strategies $\{\sigma_m(\cdot)\}_{m=1}^\infty$ that converge to $\sigma(\cdot)$.

Section 6 of his paper for details.

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 q_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 .

The following proposition fully characterizes the set of $\{\psi\}$ that are continuous, showing in which cases such a ψ exists and what beliefs it must entail.

Theorem 4.1: *There exists continuous ψ if and only if the following two conditions hold:*¹⁴

- (i) *For each i, j and $\tau < t$, if H^τ is reached, then either (a) Q_{it} fully reveals x_τ^j for all q_{it} , or (b) it does not reveal it at all (i.e. all choices of x_τ^j are consistent with Q_{it}) for all q_{it} , regardless of what players $-j$ choose at H^τ .*
- (ii) *There do not exist i, j, k and periods $t_1 < t_2 < t_3$ such that j perfectly observes $x_{t_1}^k$ in period t_2 , and i perfectly observes $x_{t_2}^j$ in period t_3 , but i does not perfectly observe $x_{t_1}^k$.*

Moreover, ψ is unique, assigning for $H^\tau \subset H^t$ such that $Q_{j\tau}(H^\tau) = q_{j\tau}$

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

in case (i)(a) if x_τ^j is off-path, and the beliefs induced by Bayes' Rule applied by $\tilde{\mathbf{x}}_{\tau,\tau}^j$ in case (i)(b).

To prove Theorem 4.1, the following lemma will be needed.

Lemma 4.4: *Event q_{it} is off-path if and only if q_{it} fully reveals some x_τ^j for some (j, τ) such that, according to the equilibrium strategy of j , j does not choose x_τ^j .*

We illustrate the intuition behind Theorem 4.1 through the following examples. It should be noted that these examples represent only possible chains of observations within the context of a larger game, so that in the periods represented in these examples, there may be other players taking actions and observing different subsets of the actions taken in past histories.

Example 4.2: Figure 3(a) shows an example of a case where the extensive-form structure of what the players can observe allows for the existence of a monotone PBE. All actions are fully observable, and so beliefs for any off-path history are found by taking the limits

¹⁴These conditions are roughly analogous to the “observable deviators” condition of Battigalli (1996), though they are not identical. While Battigalli’s condition states that, if some player deviates, all players can detect who was the deviator, these conditions here state that, if it is possible to detect that someone has deviated, one can tell to which action they have deviated.




Figure 3: Continuous beliefs possible

from a sequence of strategy profiles for which it is on-path. \square

Example 4.3: Figure 3(b) shows another example in which the structure allows for existence of monotone PBE. Player 3 cannot distinguish between the actions of Player 2, and so they will have full support: $q_{3,3}$ is generated by (x_1^1, x_2^2) for all x_2^2 . However, he observes Player 1's action, and so each $q_{3,3}$ is reachable from a unique x_1^1 . Player 2 can, in turn, observe Player 1's actions. Player 3 can then use Bayes' rule to form conditional beliefs over the play of Player 2 conditional on each action of player 1. \square

To give some intuition as to what goes wrong when conditions (i) or (ii) are not met, we provide some examples of extensive-form structure fails these conditions.

Example 4.4: In Figure 4(a), Player 3 can observe Player 2's action, but not Player 1's. There are therefore multiple possibilities for off-path beliefs for Player 3 due to the violation of condition (ii) of Theorem 4.1. To see why, suppose that Player 2 always chooses action L on path. Then two distinct beliefs are possible from a slight perturbation: that Player 2 deviated when Player 1 chose $x_1^1 = L$, or when he chose $x_1^1 = R$. Thus we cannot simply pin down the beliefs in this case from the limit of on-path strategies. \square

Example 4.5: In Figure 4(b), condition (i) of Theorem 4.1 is violated. Player 2 cannot distinguish between $x_1^1 = L$ and $x_1^1 = M$, and so $q_{2,2}$ is generated by more than one choice of x_1^1 , but not all. This will lead to a discontinuous belief structure: if Player 1 were to always choose R on path, then if he were to deviate, Player 2 would be unable to tell whether $x_1^1 = L$ or $x_1^1 = M$. As in Figure 4(a), we will therefore be unable to pin down




Figure 4: Continuous beliefs not possible

beliefs from a sequence of strategy profiles in which (at least) one of these actions is on path. \square

An important feature of Theorem 4.1 is that it uniquely pins down the beliefs that must hold upon observing off-path actions. We give some intuition in the following example.

Example 4.6: Consider the scenario described in Figure 5(a). Player i has an action set $X_t^i = \{1, 2, 3\}$ and type θ_i distributed over $[0, 10]$. Player i uses a strategy that is monotonic in θ_i , and no action is off-path. This implies that in future periods, the beliefs of player j must restrict the prior to the interval of types that choose each action.

Note that the interval of types in Figure 5(a) that chooses $x_t^i = 2$ is very tightly clustered around $\theta_i = 5$. Consider a sequence of strategy profiles in which the set of types that choose $x_t^i = 2$ shrinks to one of measure 0 centered around $\theta_i = 5$. In the limit as the probability that $x_t^i = 2$ is chosen vanishes (Figure 5(b)), all other players must believe that, conditional on $x_t^i = 2$ being played, the distribution over θ_i must converge to the point mass on $\theta_i = 5$ in the appropriate topology. \square

A point of interest regarding the beliefs generated by ψ as found in Theorem 4.1 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: higher types are more likely to take higher actions, even when




Figure 5: Continuity of beliefs

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 6.

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 information set q_{it} , 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 3.1 to generate existence of equilibrium in the static game in which payoffs are given by (4).

Lemma 4.5: *Suppose that there exists continuous ψ , and that conditions (i)-(iv) of Theorem 3.1 are satisfied in the dynamic game. Then there exists a monotone pure-strategy equilibrium in the static game Γ^2 in which payoffs are given by (4).*

4.3 Constructing the PBE

We now use the equilibrium that we have constructed in Lemma 4.5 to construct an equilibrium in the original dynamic game. 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, and thus be consistent.

Lemma 4.6: *Define $\underline{\alpha}_i(q_{i\tau}, q_{it})$ and $\bar{\alpha}_i(q_{i\tau}, q_{it})$ as in the definition of $\hat{x}_{\tau,t}^i$. Then in any equilibrium found by Lemma 4.5, $\tilde{\theta}_i^t(q_{it}, \underline{\alpha}_i(q_{i\tau}, q_{it})) = \tilde{\theta}_i^\tau(q_{i\tau}, 0)$ and $\tilde{\theta}_i^t(q_{it}, \bar{\alpha}_i(q_{i\tau}, q_{it})) =$*

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

Another way of putting Lemma 4.6 is that $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ forms a consistent strategy on-path. To see this in the context of Example 4.1, the construction of $\tilde{\theta}_i^t$ via ψ ensures that, if strategies are as given by $\tilde{\mathbf{x}}_{t,t}^i$ and $\tilde{\mathbf{x}}_{t+1,t+1}^i$, then it must be that $\tilde{\theta}_i^{t+1}(q_{i,t+1}, 1) - \tilde{\theta}_i^{t+1}(q_{i,t+1}, 0) = \frac{2}{5}[\tilde{\theta}_i^t(q_{it}, 1) - \tilde{\theta}_i^t(q_{it}, 0)]$ since $\frac{2}{5}$ is the length of the interval of types α_i that choose $\tilde{\mathbf{x}}_{t,t}^i(q_{it}, \alpha_i) = 1$. Thus choosing $\hat{\mathbf{x}}_{t+1,t}^i(q_{i,t+1}, \alpha_i)$ will be consistent for $\alpha_i \in [0, 0.4]$ with $\tilde{\mathbf{x}}_{t+1,t+1}^i$.

Since $\tilde{\mathbf{x}}_{\tau,\tau}^i$ forms a best-reply, 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 4.7: *Suppose that all players i use strategies $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ in the equilibrium generated by Lemma 4.5. Then $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T$ is also a best-reply for each α_i .*

Note that if player i 's strategy is consistent, the portion of $\tilde{\mathbf{x}}_{\tau,t}^i$ that is chosen by α_i (indexed by q_{it}) that does not choose actions $\{\tilde{\mathbf{x}}_{t,t}^i, \dots, \tilde{\mathbf{x}}_{\tau-1,t}^i\}$ that generate $q_{i\tau}$ with positive probability is essentially irrelevant for the purposes of players $-i$, since type α_i will never reach such H^τ in equilibrium. We saw this in Figure 2(d), in which only $\alpha_i \in [0, 0.4]$ chose $x_t^i = 1$, and so only those types would reach the ensuing subgame; what types $\alpha_i \geq 0.4$ would do at that subgame makes no difference in equilibrium to the expected payoffs of the other players, as this never occurs on-path in equilibrium.

The payoff of player j thus only depends (given that $q_{i\tau}$ is reached) on $(\tilde{\mathbf{x}}_{\tau,\tau}^i, \tilde{\theta}_i^\tau)$. 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 $\hat{\mathbf{x}}_{\tau,t}^i$ from which $q_{i\tau}$ is reachable, and complete the strategy function with an arbitrary 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 $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{\tau=t}^T$ as generated from $\{\tilde{\mathbf{x}}_{\tau,\tau}^i\}_{\tau=t}^T$.

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

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

Otherwise, if θ_i is completely atomic, then $\tilde{\theta}_i(q_{i\tau}, 0) = \tilde{\theta}_i(q_{i\tau}, 1) = \tilde{\theta}_i(q_{it}, 0) = \tilde{\theta}_i(q_{it}, 1) =$

θ_i^* . Then define the mixed strategy over X_τ^i (regardless of whether it is from the perspective of q_{it} or $q_{i\tau}$) so that

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

where $q_{i,\tau+1}$ is reachable only from choosing x_τ^i at $q_{i\tau}$.

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

Theorem 4.2: *Suppose that ψ is continuous, and that conditions (i)-(iv) of Theorem 3.1 hold in the dynamic game. Then there exists a monotone PBE within subgames in mixed strategies.*

Note that the result here only guarantees the existence of a mixed-strategy equilibrium, as there is no guarantee that the function $\tilde{\theta}_t^i$ 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 q_{it} . 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 in Lemma 4.1 that any on-path informational events q_t also induce absolutely continuous distributions over types θ . 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 4.2 to incorporate pure strategies.

Corollary 4.3: *Under the assumptions of Theorem 4.2, there exists a monotone PBE within subgames in pure strategies for player i whenever the conditional distribution of θ_i is not completely atomic. Thus the PBE will be monotonic in pure strategies on-path.*

Remark: It may seem more intuitively simpler to attempt to use backward induction on the agent-normal form of the game to demonstrate existence of equilibrium in this model, by finding fixed points given beliefs at any subgame from period $t + 1$ onward, and using those to generate fixed points in the subgame from period t . However, the difficulty in this approach arises from the fact that in equilibrium in a dynamic incomplete-information game, the set of types that reaches a given subgame is endogenous. Thus in order to construct an equilibrium by backward induction, one must consider the continuation equilibria for any possible subset of types, and then determine which is part of an equilibrium for the entire

game. This requires the continuation equilibria to be “well-behaved” in the sense of allowing for the application of sufficient conditions such as those in Reny’s theorem. However, it turns out that when one takes one’s own continuation strategy as given, the set of best replies will no longer be join-closed, violating one of these conditions. This is illustrated through an example presented in Appendix B.

4.4 Extensions

4.4.1 Symmetric games

Our approach can be specialized 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 $\pi(i)$.¹⁵ 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 4.4: For players $i \in I$, Γ is *symmetric* if the following conditions hold for all $i, j \in I$:

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 $i \in I$ over their actions, types, and payoffs:

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

4. For any H^t, \hat{H}^t such that $H^t = (x_1, \dots, x_{t-1})$ and $\hat{H}^t = (\{x_1^{\pi(I)}, x_1^{-I}\}, \dots, \{x_{t-1}^{\pi(I)}, x_{t-1}^{-I}\})$, then if $j = \pi(i)$, then $Q_{jt}(\hat{H}^t) = Q_{it}(H^t)$.

The first three 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. Condition 4 is needed

¹⁵Though 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.

due to the dynamic nature of the game in order to guarantee that the informational events are symmetric.

Theorem 4.4: *If players $i \in I$ are symmetric, and the conditions of Theorem 4.2 are satisfied, then there exists a symmetric monotone PBE in Γ .*

4.4.2 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. To do this, we weight 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.

Assumption 4.1: *Continuity at infinity:* for any player i in period t , then for any $\varepsilon > 0$, there exists T_ε such that for all $t > T_\varepsilon$, for any history H^{t+1} , for any θ , and for any continuations C^t, \hat{C}^t , $|u_i(H^{t+1}, C^t, \theta) - u_i(H^{t+1}, \hat{C}^t, \theta)| < \varepsilon$

Assumption 4.2: 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 4.1 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 4.2. However, we cannot use their result directly, as they only derive their results for games with finitely many players and types.

Theorem 4.5: *Suppose that $T = \infty$, and that the conditions of Theorem 4.2, as well as Assumptions 4.1 and 4.2, 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.*

With our general existence result in place, we can now turn to the cases of more general action spaces, as well as conditions under which the existence results apply (i.e. best-replies

to monotone strategies will be monotone).

5 Continuum of actions

As seen in Section 3 in Example 3.2 and Figure 1, there are potential issues with a straight-forward extension to games with a continuum of possible actions due to potential “belief entanglement.” As these issues are specific to dynamic games, we dedicate a separate section to address them.

To attempt to avoid such issues, we define the continuum game by $\Gamma = (\Theta, N, \{X_t\}_{t=1}^T, \{Q_t\}_{t=2}^T, \mathcal{M}, \Psi, u)$, and endow it with the properties of the model in Section 2, except for the restrictions that X_t^i and Q_t are finite. In order to find a sequence that converges, we must define a sequence of discretizations of Γ by $\{\Gamma_m\}_{m=1}^\infty = \{(\Theta, N, \{X_{t,m}\}_{t=1}^T, \{Q_{t,m}\}_{t=2}^T, \mathcal{M}_m, \Psi_m, u)\}_{m=1}^\infty$. Let the with corresponding monotone equilibrium strategies $\{x_{1,m}(\cdot), \dots, x_{T,m}(\cdot)\}$. We can view $X_{t,m}^i$ as a partition of X_t^i , with each (non-degenerate) interval $I \subset X_t^i$ from the partition defining the elements of $X_{t,m}^i$ by (for example) defining, for all $x_{t,m}^i \in X_{t,m}^i$, $x_{t,m}^i \equiv \inf\{x_t^i \in I\}$ and such that there does not exist $\hat{x}_{t,m}^i \neq x_{t,m}^i$ such that $\hat{x}_{t,m}^i \in I$. This allows the definition of $\Psi_m(\sigma_m)$ at every $x_t^i \in X_t^i$ by assigning the same image when σ_m chooses any $x_t^i \in I$. We can similarly partition Q_t to approximate it by $Q_{t,m}$.

Assumption 5.1: For all m , $X_{t,m}$ is finite.

Assumption 5.2: The partitions of X_t are successively finer, so that $X_{t,m} \subset X_{t,m+1}$, and $\bigcup_{m=1}^\infty X_m$ is dense on X .

Assumption 5.3: Conditions (i) and (ii) of Theorem 4.1 are satisfied for both q_{it} and $q_{it,m}$ for all m ; moreover, if x_τ^j is revealed by q_{it} , then it is revealed by $q_{it,m}$.

Assumption 5.4: At any q_{it} at which x_τ^j is fully revealed, where $\tau < t$, then for almost all $\theta_j, \hat{\theta}_j \subset \Theta_j$,

$$\lim_{m \rightarrow \infty} \mathbf{x}_{\tau,m}^j(q_{j\tau,m}, \theta_j) = \lim_{m \rightarrow \infty} \mathbf{x}_{\tau,m}^j(q_{j\tau,m}, \hat{\theta}_j) \iff \exists M \forall m \geq M : \mathbf{x}_{\tau,m}^j(q_{j\tau,m}, \theta_j) = \mathbf{x}_{\tau,m}^j(q_{j\tau,m}, \hat{\theta}_j)$$

where $\{q_{j\tau,m}\}_{m=1}^\infty$ is a sequence of events such that there exists M such that for all $m > M$, θ_j and $\hat{\theta}_j$ are in the support of types of player j who reach $q_{j\tau,m}$ on-path.

Such a convergent sequence is not guaranteed to exist, as seen in Example 3.2. That being said, we can establish that if such a sequence does exist, then an equilibrium exists in game Γ .

Theorem 5.1: *Suppose there exists a sequence of discretizations $\{\Gamma_m\}$ such that Assumptions 5.1-5.4 hold. Then there exists a monotone PBE of Γ .*

We now explore some conditions that are sufficient to guarantee that the above construction is possible. Trivially, one can construct a sequence of finer and finer partitions of X and Q which satisfy Assumptions 5.1 and 5.2. Assumption 5.3 can be generated endogenously in any discretization that satisfies the conditions of Theorem 4.1.

The tricky assumption is Assumption 5.4, as this one can possibly run afoul of Helly's selection theorem, as seen in Example 3.2. Returning to that example, we see that Assumptions 5.1-5.3 are clearly satisfied. However, we saw that the point $x_1^1 = \frac{1}{2}$ induced beliefs that were different in the limit: for Γ_m , we had $\mu_{2,m}^2(\{\theta_1 \in [\frac{1}{3}, \frac{2}{3}] | x_1^1 = \frac{1}{2}\}) = 1$, but in the limit, we had $\mu_{2,m}^2(\{\theta_1 \in [\frac{1}{3}, \frac{2}{3}] | x_1^1 = \frac{1}{2}\}) = \frac{1}{3}$. Setting $x_1^1 = \frac{1}{2}$, we see that some types chose actions that were distinct in each discretization, yet were the same in the limit. We therefore see that Assumption 5.4 is violated.

Since Assumption 5.4 can only be satisfied endogenously, we must find conditions under which Assumption 5.4 holds. While we do not provide exhaustive primitive conditions under which it will be satisfied, we provide some conditions on the primitives under which either Assumption 5.4 holds, or it is unnecessary for the construction of an equilibrium in Γ . The idea, as is standard in much of the literature¹⁶ on existence of monotone equilibrium with a continuum of actions, is to show that if the limit strategies were not an equilibrium, then there would exist some strategy profile $\{\sigma_{t,m}^i\}_{i,t}$ for sufficiently high m along the sequence that would not be an equilibrium as well; but since equilibrium is guaranteed to exist in each discretization, this leads to a contradiction.

Proposition 5.2: *Assume that for any player i such that $X_t^i \neq \{\emptyset\}$, u_i is strictly increasing in x_τ^{-i} for all θ_i and $\tau > t$, or strictly decreasing in x_τ^{-i} for all θ_i and $\tau > t$. Then if equilibrium strategies are monotone across subgames in each Γ_m , there exists a monotone PBE of Γ .*

In Section 6, when we discuss single-crossing in dynamic games, we provide conditions under which there will exist equilibrium strategies that are monotone both within and across subgames.

Note that even when Assumption 5.4 fails, one can still generate an ϵ -PBE, i.e. where players choose an action that is within ϵ of the supremum of possible payoffs at any history

¹⁶See, for example, Athey (2001) in establishing existence of equilibrium in auctions with a continuum of possible bids.

given their beliefs. The idea is that, while belief entanglement exists in the limit, it does not for any game in the sequence $\{\Gamma_m\}$. Hence one can simply extend the beliefs and actions to the other elements of X_t^i , which by continuity of payoffs will make $x_t^i \in X_{t,m}^i \subset X_t^i$ ε -optimal for all players at all histories for sufficiently high m .

Theorem 5.3: *Suppose that Γ satisfies the conditions of Theorem 4.1, except that X_t^i is a continuum for at least one (i, t) . Then there exists a monotone ε -PBE of Γ .*

To summarize our results: while it is possible in many instances to reach a full monotone PBE in games with a continuum of actions, it will not always be the case that such an equilibrium exists. As such, stronger conditions are required. However, even in the absence of such conditions, an ε -equilibrium is guaranteed to exist, which may be sufficient for applications.

6 Conditions for Monotone Best Replies

We have demonstrated earlier that, under fairly general conditions 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 mixed strategy profiles by the other players in the current period, as well as monotone mixed strategies by all players in other periods. We now explore sufficient conditions to guarantee the existence of such monotone best-replies. While these conditions 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 7, we will present an application in which the conditions of this section are not satisfied, yet there exists monotone best-replies for each player.

We will need to define a sense of complementarity of the payoffs in actions and types. While in static games, a fairly broad class of complementarities is sufficient to generate existence, we have already seen in Example 3.3 that in general dynamic Bayesian games, even supermodularity is insufficient to guarantee monotonicity of best-replies. Thus we will need stronger conditions than those sufficient in static environments to enforce such monotonicity.

Definition 6.1: A function g satisfies Milgrom and Shannon's (1994) *single-crossing property* (SCP) in (y, s) if, for $y' \geq y$ and $s' \geq s$, $g(y', s) - g(y, s) \geq 0 \implies g(y', s') - g(y, s') \geq 0$.

Definition 6.2: 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 6.3: A function g satisfies *increasing differences* (ID) in (y, s) if for $s' \geq s$ and $y' \geq y$, $g(y', s') - g(y, s') \geq g(y', s) - g(y, s)$.

Definition 6.4: 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 ensure single-crossing, and thereby generate existence of monotone PBE for games in T periods. We introduce the following lemma which will be useful for several of our results. When combined in this way, these last two definitions encompass several better-known formulations of strategic complementarity, such as supermodularity and log-supermodularity.¹⁷

Lemma 6.1: (a) (Milgrom and Shannon (1994), Theorem 4) u_i satisfies SCP in (x_t^i, θ_i) 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 $\int u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) d\mu(H^t, x_t^{-i}, C^t, \theta_{-i})$ for some measure μ , if for any pair of vectors $(\hat{H}^t, \hat{x}_t^{-i}, \hat{C}^t, \hat{\theta}_{-i}), (H^t, x_t^{-i}, C^t, \theta_{-i})$, 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 $u_i(\hat{H}^t, \hat{x}_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i}) - u_i(\hat{H}^t, x_t^i, \hat{x}_t^{-i}, \hat{C}^t, \theta_i, \hat{\theta}_{-i})$ 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 $\int u_i(H^t, x_t^i, x_t^{-i}, C^t, \theta_i, \theta_{-i}) d\mu(H^t, x_t^{-i}, C^t, \theta_{-i})$ satisfies SCP in (x_t^i, θ_i) for some measure μ , if for $\hat{x}_t^i > x_t^i$ and $(\hat{H}^t, \hat{x}_t^{-i}, \hat{C}^t, \hat{\theta}_{-i}) > (H^t, x_t^{-i}, C^t, \theta_{-i})$, both

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

¹⁷ For supermodularity, see Athey (2001); for log-supermodularity, see Athey (2002). The reader is directed to Quah and Strulovici (2012) for further instances of functions that satisfy Definitions 6.4 and 6.5.

Note that by the remark following Theorem 2 in Quah and Strulovici (2012), (c) can be extended to environments in which $(H^t, x_t^{-i}, C^t, \theta_{-i})$ 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) in Lemma 6.1 hold, then u_i satisfies SCP and SRM in (x_t^i, θ_i) .

Recall that in Example 3.3, the failure of monotonicity of best-replies stemmed from a failure of beliefs to be monotone. However, this only manifested itself 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 6.5: 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 $C, C' \in C^\tau$, then $u_i(H^{\tau+1}, C, \theta) = u_i(H^{\tau+1}, C', \theta)$, 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. 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 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 depends on what happens before they choose an action.

Proposition 6.1: *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 6.1(c), θ is affiliated and $Q_{it}(H^t) = H^t$; and
- (iii) C^t is irrelevant for player i at all H^t and any x_t^i 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 6.1 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 6.1, 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 6.2: *Suppose that*

- (i) u_i satisfies SCP and SRM in (x_t^i, θ_i) ;
- (ii) θ is independently distributed; and
- (iii) For $\tau \geq t$, except following at most a unique $(x_t^i)^*$, C^τ is irrelevant for player i at all H^τ .

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 6.1, the distribution of types and actions for players $-i$ for all 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 choices by all players. This simplifies the comparison for the choice between x_t^i and \hat{x}_t^i , and allows the invocation of standard single-crossing arguments.

Though the conditions of Proposition 6.2 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 $x_t^i = 1$, respectively. 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 7.

It is possible to obtain even stronger results when $T = 2$: one no longer need 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.

Proposition 6.3: *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 6.3 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 higher actions in higher subgames. Putting this all together, we find that there exists best-replies to monotone strategies within and across subgames.

Remark: Proposition 6.3 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.¹⁸ 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 4.

With the results from Sections 4-6 in hand, we are now able to examine some applications to some economic questions.

7 Applications

7.1 Generalized games of strategic communication

In many economic environments, an agent wish 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 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. Non-babbling equilibria are of special interest, since they allow credible communication to take place.

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

¹⁸See Reny (2011), Lemmas A.13 and A.16.

ways:

1. There may be multiple 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;
3. We relax the strong assumptions on payoffs over messages, so that we do not exogenously require full separation or convex loss functions; and
4. We provide weaker single-crossing conditions than those previously found which are sufficient to generate a fully-separating equilibrium (i.e. reporting one's type truthfully need not minimize the cost of signalling).

We will thus provide sufficient conditions under which a non-babbling monotone equilibrium exists.

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) has 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) . In terms of information structure, we assume that x_1 is perfectly observable. These conditions are weaker than those of Van Zandt and Vives (2007), who make many additional assumptions,¹⁹ as well as those of Kartik et al. (2007) and Kartik (2009), who assumes not only supermodularity and utility maximization from truthful revelation (i.e, $\arg \max_{x_1^1 \in X_1^1} u_1^1(x_1^1, \theta_1) = \theta_1$), but also convexity of the loss function. Lastly, we assume that actions are perfectly observable.

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

¹⁹See Proposition 20 in their paper.

is of higher quality if it is indeed of higher quality. The investors may have their own private signals as well. 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 7.1: *There exists a monotone PBE in the generalized game of strategic communication.*

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 $j \geq 2$, 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_2^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. 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 signalling 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 to induce pooling.

Lemma 7.1: *In any monotone PBE with beliefs defined as in Theorem 4.1, 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 7.2: *There exists a non-babbling monotone PBE of the generalized strategic communication game if*

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

To illustrate Proposition 7.2, 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 7.2 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 7.2 to provide conditions under which, in fact, not only is there not complete pooling (i.e. babbling), but 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 7.3: *Suppose that $\sup_{x_1^1} U_1^1(x_1^1, \bar{\theta}_1) \neq U_1^1(\bar{x}_1^1, \bar{\theta}_1)$ regardless of the choice of monotone $\mathbf{x}_2^j(H^2, \theta_j)$ 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).

7.2 Stopping games

As noted earlier in Section 6, 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 . All actions from previous periods are observed by all players. The game ends when all players but (at most) one have exited, and the one who has not is declared the winner.

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 6.2 hold within each period t .

Proposition 7.4: *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 can apply the previous proposition to several auction environments. Previous existence results for dynamic auctions can be found in Milgrom and Weber (1982), Maskin (1992), Lizzeri and Persico (2000), Krishna (2003), and Birulin and Izmalkov (2011); these have mostly focused on efficiency rather than existence. Previous existence results for wars of attrition include Milgrom and Weber (1985), Fudenberg and Tirole (1986), Krishna and Morgan (1997), Bulow and Klemperer (1999), and Myatt (2005), though most of these focus on symmetric agents. Our results generalize all of these environments to show exis-

tence of monotone equilibrium in type: the higher one's type, the longer one remains in the game.

Proposition 7.5: *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 (with at least one strictly so), and*

$$t^* = \sup_t \{ \exists i, j : x_t^j = 1, x_t^i = 1 \}$$

$$t_i = \sup_t \{ x_t^i = 1 \}$$

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

The following games have a monotone PBE:

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

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

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

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

(iii) *Auctions with costly bidding with independent types,²⁰ 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.

The results for Proposition 7.5 apply to only finite T ; however, in the standard versions of stopping games, play can last for an arbitrary length of time. Furthermore, these games do not satisfy the continuity at infinity condition of Theorem 4.5, as a decision

to stay in the auctions until $t = \infty$ yields a payoff of $-\infty$ for some player. Fortunately, these results can be extended to $T = \infty$ using a similar approach.

Proposition 7.6: *In any of the games covered in Proposition 7.5, if $b_i \rightarrow 0$ or $c_i \rightarrow \infty$ as*

²⁰A 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”).

$t \rightarrow \infty$, there exists a monotone PBE when $T = \infty$ (symmetric if players are symmetric).

Proof: See Appendix. \square

We now turn to examine conditions under which the existence result of Proposition 7.5 can be extended to games of continuous time. We can view any strategy in a stopping game Γ as a function of which bidders dropped out (and when), as well as one's type, giving as output when one should drop out. Unfortunately, given that there will be an uncountable number of histories in games with $N > 2$, and it is not clear that the strategy function is monotonic in histories; indeed, for example, in the war of attrition, it may be that a later exit time by some player causes some other player j to exit so much later that it is worthwhile for player i to exit earlier rather than incur the cost of staying in the game. In the case of auctions with affiliated signals and interdependent values, the bidders would suffer the “winner’s curse” from winning, as they could only do so against lower types, yielding low payoffs while still having to pay a high bid. Hence it will not be possible to use the method of Proposition 7.6 involving Helly’s selection theorem to take the limit of a sequence of equilibria of discretizations of the game.²¹

Additionally, it could be that a positive measure of types θ_i drops out at a given time, in response to which (or in response to the lack of this decision) a positive measure of types θ_j drop out (almost) immediately in order to avoid the winner’s curse. However, if one were to take the limit using Helly’s selection theorem, all would be choosing to exit at the same time. This “race for the exits” which is present in the reactive element of the strategies would be missing in the limit, changing the incentives of the players.

In the case of $N = 2$, though, the game essentially has a single possible history, given that the choice of any player to stop ends the game, so the first possible issue is no longer present. Hence we can model the strategy of any player i with type θ_i as a stopping time decision $t_i(\theta_i)$. Given that strategies in any given period are monotone, it follows that $t_i(\theta_i)$ is monotone as well. It is therefore possible to avoid the issues related to reactions to exits, and so we can find a well-defined limit of the equilibrium stopping times. Moreover, if $v_i(\theta)$ is strictly increasing, it turns out that we find that the continuous-time limit of the war of attrition or the auction with costly bidding is an equilibrium, as it turns out that any possible issues involving reactions to the other player not dropping out would induce a

²¹Note that this issue does not arise in Athey (2001), since the games in her environment are static, and so there is only one history possible. This enabled her to ignore these types of reactions; in a dynamic environment, though, it will in general be impossible to ignore these. This is not necessarily problematic, since an immediate reaction to an action by another player must be taken into account in certain environments. Nevertheless, it shows that there is no neat way to model such an environment in continuous time.

profitable deviation at some point along the limiting sequence.

Formally, consider any sequence of discretizations $\{\Gamma_m\}$ given by a sequence of potential stopping times $\mathcal{T}_m = \{t_{m_1}, \dots, t_{m_K}\}$, where $\mathcal{T}_m \subset \mathcal{T}_{m+1}$ and $\bigcup_{m=1}^{\infty} \mathcal{T}_m$ is dense on $[0, \infty]$. We assume that $t_{m_1} = 0$. In any discretization, the payoff for stopping at any time $t \in (t_{m_k}, t_{m_{k+1}})$ is that of stopping at t_{m_k} , and generates the same history.

Proposition 7.7: *Consider any continuous-time version with $T \leq \infty$ of the war of attrition or auction with costly bidding, where b_i and c_i are continuous. If $N = 2$ and $v_i(\cdot)$ is strictly increasing in θ , then there exists an equilibrium in which (a) on path, a positive measure of types of player i can exit simultaneously only at $t = 0$, and (b) if play does not end with probability 1 at $t = 0$, there is a positive probability of exit at any interval $(t, t + \delta)$ for any $\delta > 0$ until the game ends with probability 1.*

As before, Proposition 7.7 extends to show existence of symmetric monotone PBE in continuous time when players are symmetric.

References

- Aradillas-López, A., Gandhi, A., and Quint, D. (2013). “Identification and inference in ascending auctions with correlated private values.” *Econometrica*, 81, 489-534.
- 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.
- Back, K., and Baruch, S. (2013): “Strategic liquidity provision in limit order markets.” *Econometrica*, 81, 363-392.
- 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.
- Echenique, F. (2004): "Extensive-form games and strategic complementarities." *Games and Economic Behavior*, 46, 348-364.
- Eilenberg, S., & 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.
- Gentry, M., and Li, T. (2014): "Identification in auctions with selective entry." *Econometrica*, 82, 315-344.
- 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.
- Hagenbach, J., Koessler, F., and Perez-Richet, E. (2014): "Certifiable pre-play communication: Full disclosure." *Econometrica*, 82, 1093-1132.
- 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. M., and Wilson, R. (1982): "Sequential equilibria." *Econometrica*, 50, 863-894.
- (1982): "Reputation and imperfect information." *Journal of Economic Theory*, 27,

253-279.

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, J., and Liu, Q. (2013): "Gambling reputation: Repeated bargaining with outside options." *Econometrica*, 81, 1601-1672.

Maskin, E. (1992): "Auctions and Privatization." In H. Siebert (ed): *Privitization*. Institut fur Weltwirtschaften der Universitat Kiel, Kiel, 115-136.

McAdams, D. (2003): "Isotone equilibrium in games of incomplete information." *Econometrica*, 71, 1191-1214.

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. (2014): "Sequential equilibria of multi-stage games with infinite sets of types and actions." Draft notes.

Okuno-Fujiwara, M., Postlewaite, A., and Suzumura, K. (1990): "Strategic information revelation." *Review of Economic Studies*, 57, 25-47.

Pavan, A., Segal, I., and Toikka, J. (2014): "Dynamic mechanism design: a Myersonian approach." *Econometrica*, 82, 601-654.

Quah, J. and Strulovici, B. (2012): "Aggregating the single crossing property." *Econometrica*, 80, 2333-2348.

Riley, J. G. (1979): "Evolutionary equilibrium strategies." *Journal of Theoretical Biology*, 76, 109-123.

- Reny, P. J. (2011): “On the existence of monotone pure strategy equilibria in Bayesian games.” *Econometrica*, 79, 499-553.
- 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: Proofs

Proof of Proposition 3.2: Suppose that there exists a monotone PBE. Clearly, in period 3, player 2 plays $x_3^2 = \theta_2$, so player 3 tries to match accordingly. So, if all types of player 2 pool in period 2, then in period 3, player 3 chooses $x_3^3 = 0$. However, if $\theta_2 = 0$ and $\theta_2 = 1$ separate, then in a monotone equilibrium, it must be that $\theta_2 = 0$ plays $x_2^2 = 1$ and $\theta_2 = 1$ plays $x_2^2 = 2$. Hence player 3 chooses $x_3^3 = 0$ if $x_2^2 = 1$ and $x_3^3 = 1$ if $x_2^2 = 2$.

Next, it is clear that $\theta_1 = 0$ always plays $x_1^1 = 1$, since then player 1 does not care about what happens in period 3. Likewise, $\theta_1 = 1$ always plays $x_1^1 = 2$, since this gives an additional payoff of 0.5 compared to $x_1^1 = 1$, while in period 3, the worst that can happen is a loss of payoff of 0.1 (which would occur if players 2 and 3 chose different actions). Therefore, in any monotone equilibrium, it must be that all types $\theta_1 > 1$ choose $x_1^1 = 2$.

In period 2, player 1 chooses

$$x_2^1 = \begin{cases} 0.5, & \theta_1 \in [0, 1) \\ 1.5, & \theta_1 \in [1, 2] \end{cases}$$

By the monotonicity of actions of player 1 in θ_1 , the support of the types θ_1 which choose $x_1^1 = 2$ must be $[\theta_1^*, 2]$, where $\theta_1^* \geq 0$. Hence all types θ_2 will choose to play $x_2^2 = 1$ upon observing $x_1^1 = 2$ if almost all types of player 1 play $x_1^1 = 2$. In this case, even if $\theta_1^* = 0$, the payoff to $\theta_2 = 0$ will be $-\frac{1}{2}(1.1)^2 - \frac{1}{2}(0.1)^2$ from choosing $x_2^2 = 1$, while from playing $x_2^2 = 2$, the payoff will be $-\frac{1}{2}(0.9)^2 - \frac{1}{2}(0.1)^2$; the incentive to play x_2^2 grows even larger if $\alpha > 0$. Hence $\theta_2 = 0$ will choose to play $x_2^2 = 2$. By monotonicity, it must be that $\theta_2 > 0$ chooses $x_2^2 = 2$.

On the other hand, by monotonicity, there cannot be $\theta_1 > 1$ playing $x_1^1 = 1$, so $\theta_1^* \leq 1$. Then the payoff for $\theta_2 = 0$ to play $x_2^2 = 1$ upon observing $x_1^1 = 1$ would be $-(0.1)^2$, while from playing $x_2^2 = 2$, it would be $-(0.9)^2$, so it will be optimal for $\theta_2 = 0$ to play $x_2^2 = 1$. On the other hand, if $\theta_2 = 1$, then the payoff from playing $x_2^2 = 1$ would be $-(0.1)^2 + 1$, while from playing $x_2^2 = 2$, it would be $-(0.9)^2 + 2$, so $\theta_2 = 1$ will play $x_2^2 = 2$.

Now consider the choice of $\theta_1 = 2$. Because of the induced pooling, the payoff from choosing $x_1^1 = 2$ will be $2(1.5) - (0.5)^2 + 0.1(1)(64) = 9.15$, while from choosing $x_1^1 = 1$, the payoff will be at least $(1)(1.5) - (0.5)^2 + (0.5)(0.1)(1)(64) + (0.5)(0.1)(2)(64) = 10.85$. Hence $\theta_1 = 2$ would want to deviate to $x_1^1 = 1$, violating monotonicity in period 1.

□

Proof of Lemma 4.1: 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). This generates a unique $q_{i,t-1}$. Since strategies are monotone, the support of θ_i conditional on choosing x_{t-1}^i at $q_{i,t-1}$ 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 I of players and is absolutely continuous for all other players, which we write as $-i \setminus I$. \square

Proof of Lemma 4.2: Note that $\hat{\mathbf{x}}_{\tau,t}^i$ (respectively, $\hat{\mathbf{x}}_{\tau,t,m}^i$) divides the strategy of player i at $q_{i\tau}$ 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 $(\underline{\alpha}_i(q_{i\tau}, q_{it}), \bar{\alpha}_i(q_{i\tau}, q_{it}))$. We inductively show that $\hat{\mathbf{x}}_{\tau,t,m}^i \rightarrow \hat{\mathbf{x}}_{\tau,t}^i$, $\underline{\alpha}_{i,m}(q_{i\tau}, q_{it}) \rightarrow \underline{\alpha}_i(q_{i\tau}, q_{it})$, and $\bar{\alpha}_{i,m}(q_{i\tau}, q_{it}) \rightarrow \bar{\alpha}_i(q_{i\tau}, q_{it})$. For $\tau = t$, this is trivial because $\tilde{\mathbf{x}}_{t,t}^i = \hat{\mathbf{x}}_{t,t}^i$, $\underline{\alpha}_{i,m}(q_{it}, q_{it}) = \underline{\alpha}_i(q_{it}, q_{it}) = 0$ and $\bar{\alpha}_{i,m}(q_{it}, q_{it}) = \bar{\alpha}_i(q_{it}, q_{it}) = 1$.

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

Now look at $\alpha_i \in (\underline{\alpha}_i(q_{i,\tau+1}, q_{it}), \bar{\alpha}_i(q_{i,\tau+1}, q_{it}))$. Suppose that $\hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i) = x_{\tau+1}^i$. For almost all such α_i , there exists $\varepsilon > 0$ such that $\hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i - \varepsilon) = \hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i + \varepsilon) = x_{\tau+1}^i$. Let $\hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i(q_{i,\tau+1}, q_{it})}{\bar{\alpha}_i(q_{i,\tau+1}, q_{it}) - \underline{\alpha}_i(q_{i,\tau+1}, q_{it})}$ and define $\hat{\alpha}_{i,m}$ analogously. Note that $\hat{\alpha}_{i,m} \rightarrow \hat{\alpha}_i$. For almost all $\hat{\alpha}_i \in [0, 1]$, it will be the case that if $\tilde{\mathbf{x}}_{\tau+1,\tau+1}^i(q_{i,\tau+1}, \hat{\alpha}_i) = x_{\tau+1}^i$, then $\tilde{\mathbf{x}}_{\tau+1,\tau+1}^i(q_{i,\tau+1}, \hat{\alpha}_i - \varepsilon) = \tilde{\mathbf{x}}_{\tau+1,\tau+1}^i(q_{i,\tau+1}, \hat{\alpha}_i + \varepsilon) = x_{\tau+1}^i$. 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 $m > M$, $\tilde{\mathbf{x}}_{\tau+1,\tau+1,m}^i(q_{i,\tau+1}, \hat{\alpha}_i - \frac{\varepsilon}{2}) = \tilde{\mathbf{x}}_{\tau+1,\tau+1,m}^i(q_{i,\tau+1}, \hat{\alpha}_i + \frac{\varepsilon}{2}) = x_{\tau+1}^i$. Moreover, for sufficiently high M , since $\hat{\alpha}_{i,m} \rightarrow \hat{\alpha}_i$, it follows that for $m >$

M , $\hat{\alpha}_{i,m} \in (\hat{\alpha}_i - \frac{\varepsilon}{2}, \hat{\alpha}_i + \frac{\varepsilon}{2})$. Since $\hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i) = \tilde{\mathbf{x}}_{\tau+1,\tau+1}^i(q_{i,\tau+1}, \hat{\alpha}_i)$, it follows that $\hat{\mathbf{x}}_{\tau+1,t,m}^i(q_{i,\tau+1}, \alpha_i) \rightarrow \hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i)$ for $\alpha_i \in (\underline{\alpha}_i(q_{i,\tau+1}, q_{it}), \bar{\alpha}_i(q_{i,\tau+1}, q_{it}))$.

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

Since $\hat{\mathbf{x}}_{\tau+1,t,m}^i(q_{i,\tau+1}, \alpha_i) \rightarrow \hat{\mathbf{x}}_{\tau+1,t}^i(q_{i,\tau+1}, \alpha_i)$ 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

Proof of Lemma 4.3: By Lemma 4.1, the conditional distributions given q_{it} and H^t over Θ_{-i} are those defined by the prior restricted to a Cartesian product of subintervals $[\theta_{j,k}^1, \theta_{j,k}^2]$. If $\mu_{t,k}^i \rightarrow \mu_t^i$ in the weak-* topology, then it must be that $\theta_{j,k}^1 \rightarrow \theta_j^1$ and $\theta_{j,k}^2 \rightarrow \theta_j^2$; hence $\tilde{\theta}_{j,k}^t \rightarrow \tilde{\theta}_j^t$. The proof for $g_{t,k}$ then follows from the fact that $f(\Theta_{-i}|\theta_i)$ is continuous in θ_i , and $\tilde{\theta}_{i,k}^t \rightarrow \tilde{\theta}_i^t$ everywhere. \square

Proof of Theorem 4.1: Suppose that (i) is false. Without loss of generality, set $\tau = t - 1$ (the same argument will hold for other values of τ). Then for a given H^{t-1} , there exists some strategy profile σ_{t-1} that does not generate q_{it} . Suppose that all players choose the same actions in periods $\tau < t - 1$ which generate H^{t-1} , regardless of type. We first show that it is not possible for there to exist some x_{t-1}^{-j} for which there exist $\{x_{t-1}^{j,k}\}_{k=1}^3$ such that q_{it} is generated by $\{H^{t-1}, x_{t-1}^{j,k}, x_{t-1}^{-j}\}_{k=1,2}$ but not by $\{H^{t-1}, x_{t-1}^{j,3}, x_{t-1}^{-j}\}$. Suppose that σ_{t-1} mandates that players select $(x_{t-1}^{j,3}, x_{t-1}^{-j})$ in period $t - 1$ with probability 1, regardless of type. Then if some measure ε of player j deviates to $x_{t-1}^{j,1}$, by Bayes' Theorem we must have that $\mu_t^i(\Theta_{-i}, H^{t-1}, x_{t-1}^{j,1}, x_{t-1}^{-j} | q_{it}, \theta_i) = 1$. By continuity, the same must be true at $\varepsilon = 0$. Similarly, from the possibility that some measure ε deviates to $x_{t-1}^{j,2}$, we must have $\mu_t^i(\Theta_{-i}, H^{t-1}, x_{t-1}^{j,2}, x_{t-1}^{-j} | q_{it}, \theta_i) = 1$. But these beliefs are contradictory, and so ψ cannot be continuous.

Next, suppose that there exist $x_{t-1}, \hat{x}_{t-1} \in X_{t-1}$ such that q_{it} is generated by $\{H^{t-1}, x_{t-1}\}$ and $\{H^{t-1}, \hat{x}_{t-1}\}$ but not $\{H^{t-1}, \hat{x}_{t-1}^j, x_{t-1}^{-j}\}$. Suppose that σ_{t-1} mandates that players choose $(\hat{x}_{t-1}^j, x_{t-1}^{-j})$ with probability 1. Then from having measure ε of player j choose x_{t-1}^j instead and taking the limit as $\varepsilon \rightarrow 0$ yields $\mu_t^i(\Theta_{-i}, H^{t-1}, x_{t-1} | q_{it}, \theta_i) = 1$. Similarly, having measure ε of each player k among $-j$ chose \hat{x}_{t-1}^k instead and taking the limit as $\varepsilon \rightarrow 0$ yields $\mu_t^i(\Theta_i, H^{t-1}, \hat{x}_{t-1}^j, X_{t-1}^{-j} | q_{it}, \theta_i) = 1$. These beliefs are also contradictory, and

so ψ cannot be continuous in this case as well.

Next, if (ii) is false, then there exist such players i, j, k . Suppose that according to σ , player k splits his choice between $x_{t_1}^k, \hat{x}_{t_1}^k \in X_{t_1}^k$ with positive probability, and that in period t_2 , player j always chooses $x_{t_2}^j$. Lastly, we assume that q_{it_3} is only generated $\hat{x}_{t_2}^j \in X_{t_2}^j$, which we can do since i observes j 's period t_2 action. Consider the following two deviations. First, if a measure $\varepsilon > 0$ of player j chooses $\hat{x}_{t_2}^j$ if and only if he observes $x_{t_1}^k$, then in period t_3 (taking the limit as $\varepsilon \rightarrow 0$), player i 's beliefs must satisfy $\mu_{t_3}^i(\Theta_{-i}, \{H^{t_3} : x_{t_1}^k \in H^{t_3}\} | q_{it_3}, \theta_i) = 1$. Similarly, from a measure of ε of player j choosing $\hat{x}_{t_2}^j$ if and only if he observes $\hat{x}_{t_1}^k$, then player i 's beliefs must satisfy $\mu_{t_3}^i(\Theta_{-i}, \{H^{t_3} : \hat{x}_{t_1}^k \in H^{t_3}\} | q_{it_3}, \theta_i) = 1$. Since these beliefs are contradictory, ψ cannot be continuous.

Now suppose that both (i) and (ii) are true. Recall that q_{jt} fully reveals $q_{j\tau}$ for all $\tau \leq t$. Let $\mathcal{H}^{q_{it}} \subset \mathcal{H}^t$ be the set of time- t histories that generate q_{it} . Thus for q_{it} that is on path and any open $A \subset \Theta$, by the construction of ψ and $\hat{\mathbf{x}}_{\tau,1}^i$ (suppressing arguments),

$$\mu_t^i(A, H^t | q_{it}, \tilde{\theta}_i^1(q_{i1}, \alpha_i)) = \frac{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\hat{\mathbf{x}}_{\tau,1}^i, \hat{\mathbf{x}}_{\tau,1}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i} | q_1, \alpha_i) d\alpha_{-i}}{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\hat{\mathbf{x}}_{\tau,1}^i, \hat{\mathbf{x}}_{\tau,1}^{-i}\}_{\tau=1}^{t-1}) \in (\Theta, \mathcal{H}^{q_{it}})\}} (\alpha_{-i}) g_1(\alpha_{-i} | q_1, \alpha_i) d\alpha_{-i}} \quad (5)$$

These integrals converge when $\sigma_m \rightarrow \sigma$ since $\tilde{\theta}^1$ 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 q_{it} is off-path, note that the denominator in (5) is 0. To circumvent this issue, we invoke Lemma 4.4 (included here again for exposition).

Lemma 4.4: *q_{it} is off-path if and only if q_{it} fully reveals some x_{τ}^j for some (j, τ) such that, according to the equilibrium strategy of j , j does not choose x_{τ}^j .*

Proof of Lemma 4.4: The “if” direction is trivial. For the “only if” direction, we show this inductively. In period 2, we know that if all such j choose x_1^j , then if q_{i2} can be reached by any H^2 , it must be generated by all j who fall under (i)(b). Now suppose that the lemma is true through period $t - 1$. Then in period t , given H^{t-1} such that q_{it} is reachable, the only way that q_{it} can be off-path is from some choice of x_{t-1} . This can happen from x_{t-1} not occurring on-path given H^{t-1} , or from x_{t-1} revealing that H^{t-1} has not occurred. The former case is identical to that of period 2. In the latter case, by (ii), we know that if H^{t-1} indeed has not occurred, then any player j in period $t - 1$ cannot observe this if i cannot observe this in period t but can observe x_{t-1}^j , i.e. for any H^{t-1}, \hat{H}^{t-1} such that q_{it} is generated for some x_{t-1} , then for any $q_{j,t-1}$, either $q_{j,t-1}$ is generated by both H^{t-1} and

\hat{H}^{t-1} , or q_{it} is generated by both $\{H^{t-1}, \hat{x}_{t-1}\}$ and $\{H^{t-1}, x_{t-1}\}$ for any \hat{x}_{t-1} that satisfies (i). Thus it cannot be that x_{t-1} induces i to believe that H^{t-1} has not occurred by the choice of action by any j in period $t-1$ as long as (i) is satisfied, as i 's observation of the choice of x_{t-1} implies that it could have occurred after either H^{t-1} or \hat{H}^{t-1} . \square

Continuing the proof of Theorem 4.1, suppose that q_{it} is off-path. Then H^t must be off path if H^t generates q_{it} . Let R be the set of pairs $\{(j, \tau)\}$ representing the set of players whose period τ actions are completely revealed by q_{it} if H^t is the true history. 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 $\tilde{x}_{\tau, \tau}^j$. Then for any $x_{t_j}^j \in X_{t_j}^j$ and almost all α_j such that $\tilde{x}_{t_j, t_j}^j(q_{jt_j}, \alpha_j) < x_{t_j}^j$, there must exist some M such that for all $m > M$, $\tilde{x}_{t_j, t_j, m}^j(q_{jt_j}, \alpha_j) < x_{t_j}^j$. A similar argument holds for $\tilde{x}_{t_j, t_j}^j(q_{jt_j}, \alpha_j) > x_{t_j}^j$.

Let $\underline{\theta}_{j, m} = \sup\{\theta_j : \tilde{x}_{t_j, t_j, m}^j(q_{jt_j}, \alpha_j) < x_{t_j}^j \text{ and } \tilde{\theta}_{j, m}^{t_j}(q_{jt_j}, \alpha_j) = \theta_j\}$ and $\bar{\theta}_{j, m} = \inf\{\theta_j : \tilde{x}_{t_j, t_j, m}^j(q_{jt_j}, \alpha_j) > x_{t_j}^j \text{ and } \tilde{\theta}_{j, m}^{t_j}(q_{jt_j}, \alpha_j) = \theta_j\}$. Then, as shown in Lemma 4.1, the support of types who choose action x_τ^j will be contained in $[\underline{\theta}_{j, m}, \bar{\theta}_{j, m}]$. By the argument from the previous paragraph, $\lim_{m \rightarrow \infty} \underline{\theta}_{j, m} = \lim_{m \rightarrow \infty} \bar{\theta}_{j, m}$; call this limit θ_j^* . Hence μ_t^i as generated by ψ must place probability 1 on θ_j^* conditional on observing q_{it} . For subsequent t' , because θ_j must be contained in the set of types who chose action x_τ^j for each σ_m , it follows that the beliefs over θ_j for the subsequent t' converge to the same θ_j .

Lastly, for $(j, \tau) \notin R$, we have already shown that conditional on $\{x_\tau^j\}_{(j, t) \in R} \subset H^t$, the beliefs over $(j, \tau) \notin R$ must follow by Bayes' rule on path. For H^t that is off-path, any deviation by $(j, \tau) \in R$ cannot affect the choice of strategies by $(j, \tau) \notin R$ 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 $H^t \supset \{x_\tau^j\}_{(j, t) \in R}$, μ_t^i is uniquely determined to set

$$\mu_t^i(A, H^t | q_{it}, \tilde{\theta}_i^1(q_{i1}, \alpha_i)) = \lim_{m \rightarrow \infty} \frac{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\tilde{x}_{\tau, 1, m}^i, \tilde{x}_{\tau, 1, m}^{-i}\}_{\tau=1}^{t-1}) \in (A, H^t)\}} (\alpha_{-i}) g_1(\alpha_{-i} | q_1, \alpha_i) d\alpha_{-i}}{\int 1_{\{(\tilde{\theta}_i^1, \tilde{\theta}_{-i}^1, \{\tilde{x}_{\tau, 1, m}^i, \tilde{x}_{\tau, 1, m}^{-i}\}_{\tau=1}^{t-1}) \in (\Theta, \mathcal{H}^{q_{it}})\}} (\alpha_{-i}) g_1(\alpha_{-i} | q_1, \alpha_i) d\alpha_{-i}}$$

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

Proof of Lemma 4.5: We must show that conditions (i)-(iv) also hold in the static interpretation of the game as well. 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 dynamic 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 dynamic game operates lexicographically over periods, so that if player i has two potential vectors of actions that he considers from $\prod_{t=1}^T (X_t^i)^{|Q_{it}|}$ 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 $q_{i\tau}$, and the meet picks from the lower. Thus, if given i 's opponents' actions, the two vectors of actions chosen would induce the sequences of actions $x^i, \hat{x}^i \in \prod_{t=1}^T X_t^i$, such that x^i and \hat{x}^i agree up to period τ , and $x_\tau^i > \hat{x}_\tau^i$, then $x^i \vee \hat{x}^i = x^i$.²² 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 q_{it} from choosing x_t^i can be rewritten as (suppressing arguments for $\tilde{\mathbf{x}}_{\tau,t}$, $\hat{\mathbf{x}}_{\tau,t}$, and $\tilde{\theta}^t$)

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

By Lemma 4.4 and Theorem 4.1, 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 will hold in the static interpretation with the partial order given above 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 3.1 are satisfied for the static interpretation as well, and so a monotone equilibrium will exist. \square

Proof of Lemma 4.6: By construction of ψ , if the distribution of θ_i is completely atomic at q_{it} , then

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

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 $[\tilde{\theta}_i^t(q_{it}, 0), \tilde{\theta}_i^t(q_{it}, 1)]$ must be equal to the set of types who choose

²²Unlike 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 $q_{it'}$, but is admissible within a best-reply since $q_{it'}$ is off-path due to some action by i in period τ .

$\{\hat{\mathbf{x}}_{t',1}^i\}_{t'=1}^{t-1} = \{x_{t',t}^i\}_{t'=1}^{t-1}$ at the respective $q_{it'} \subset q_{it}$ for which q_{it} is reachable from the perspective of period 1. The same holds for $[\tilde{\theta}_i^\tau(q_{i\tau}, 0), \tilde{\theta}_i^\tau(q_{i\tau}, 1)]$.

Next, note that in equilibrium, $\hat{\mathbf{x}}_{t,1}^i(q_{it}, \alpha_i) = \hat{\mathbf{x}}_{t,t}^i(q_{it}, \hat{\alpha}_i)$, where $\hat{\alpha}_i = \frac{\alpha_i - \underline{\alpha}_i(q_{it}, q_{i1})}{\bar{\alpha}_i(q_{it}, q_{i1}) - \underline{\alpha}_i(q_{it}, q_{i1})}$. 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 $\hat{\alpha}_i \in [\underline{\alpha}_i(q_{i,t+1}, q_{it}), \bar{\alpha}_i(q_{i,t+1}, q_{it})]$ who choose $\tilde{\mathbf{x}}_{t,t}^i(q_{it}, \hat{\alpha}_i) = x_t^i \in X_t^i$ which can reach some $q_{i,t+1}$ generates the same set of types $\{\theta_i : \tilde{\theta}_i^t(q_{it}, \hat{\alpha}_i) = \theta_i\}$ as $\{\theta_i : \tilde{\theta}_i^1(q_{i1}, \alpha_i) = \theta_i\}$, where $\alpha_i \in [\underline{\alpha}_i(q_{i,t+1}, q_{i1}), \bar{\alpha}_i(q_{i,t+1}, q_{i1})]$ such that α_i chooses (conditional on each relevant $q_{it'} \subset q_{i,t+1}$ being reached) $\{\tilde{\mathbf{x}}_{t',1}^i\}_{t'=1}^t = \{x_{t',t}^i\}_{t'=1}^{t-1}$ which can reach $q_{i,t+1}$.

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

$$\tilde{\theta}_i^t(q_{it}, \underline{\alpha}_i(q_{i\tau}, q_{it})) = \tilde{\theta}_i^1(q_{i1}, \underline{\alpha}_i(q_{i\tau}, q_{i1})) = \tilde{\theta}_i^\tau(q_{i\tau}, 0)$$

$$\tilde{\theta}_i^t(q_{it}, \bar{\alpha}_i(q_{i\tau}, q_{it})) = \tilde{\theta}_i^1(q_{i1}, \bar{\alpha}_i(q_{i\tau}, q_{i1})) = \tilde{\theta}_i^\tau(q_{i\tau}, 1)$$

□

Proof of Lemma 4.7: 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 $q_{i,\tau+1}$. 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 $\{\tilde{\mathbf{x}}_{t',t}^i\}_{t'=\tau+1}^T$ (suppressing arguments) is

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

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

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

We now show that the conditional payoff upon reaching $q_{i\tau}$ is maximized by replacing $\tilde{\mathbf{x}}_{\tau,t}^i$ with $\hat{\mathbf{x}}_{\tau,t}^i$ is not decreased. In the case where $\tilde{\mathbf{x}}_{\tau,t}^i = \hat{\mathbf{x}}_{\tau,t}^i$, this is true by definition since $\tilde{\mathbf{x}}_{\tau,t}^i$ was optimal. On the other hand, by Lemma 4.6, if $\hat{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \alpha_i) = \hat{\mathbf{x}}_{\tau,\tau}^i(q_{i\tau}, \hat{\alpha}_i)$ where $\tilde{\theta}_i^{\tau+1}(q_{i,\tau+1}, \hat{\alpha}_i) = \tilde{\theta}_i^\tau(q_{i\tau}, \alpha_i)$, then since $\hat{\mathbf{x}}_{\tau,\tau}^i(q_{i\tau}, \hat{\alpha}_i)$ is optimal for $\hat{\alpha}_i$ from the perspective

of period τ , $\hat{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \alpha_i)$ must be optimal as well from the perspective of period t because $\{\hat{\mathbf{x}}_{\tau,t}^{-i}\}_{-i,\tau,t}$ forms a consistent strategy profile. Thus in either case, $\hat{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \alpha_i)$ is optimal. \square

Proof of Theorem 4.2: By Lemma 4.5, there exists an equilibrium in the static game in which payoffs are given by (4). As argued above in Lemma 4.7, choosing $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ is also a best reply given equilibrium strategy profile $\{\hat{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ when payoffs are given by (4); 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 4.6, the set of types of player i that choose a given x_t^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 $q_{i\tau}$, we must ensure that there is then a best-reply $x_\tau^i \in BR_\tau^i(q_{i\tau}, \tilde{\theta}_i^i(q_{it}, \alpha_i))$ for $\alpha_i \geq \bar{\alpha}_i(q_{i\tau}, q_{it})$ that is at least as great as $\hat{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \bar{\alpha}_i(q_{i\tau}, q_{it}))$. Fortunately, by the fact that best-replies are increasing in the strong-set order, it must be that $\max\{x_\tau^i \in BR_\tau^i(q_{i\tau}, \tilde{\theta}_i^i(q_{it}, \alpha_i))\} \geq \hat{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \bar{\alpha}_i(q_{i\tau}, q_{it}))$. An analogous argument holds for $\alpha_i \leq \underline{\alpha}_i(q_{i\tau}, q_{it})$. Choosing such values for what type α_i would do at $q_{i\tau}$ from the perspective of q_{it} is therefore a monotone best-reply.

Suppose that we look at the interim payoffs in Γ^1 as given by (2), i.e. if the strategy profile is $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$, then

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

Note now that if we use the strategies described in the above by $(x_t^i)^*$, the strategies are consistent, and so these coincide with the payoffs given by (4) by Lemma 4.6. Therefore the strategy given by $\mathbf{x}_\tau^i(\cdot, \cdot)$ is optimal from the perspective of period t for type $\tilde{\theta}_i^i(q_{it}, \alpha_i)$ when the conditional distribution over θ_i is absolutely continuous at q_{it} ; similarly, when the distribution is completely atomic, the strategy given by $\rho_\tau^i(x_t^i | q_{i\tau}, \theta_i)$ gives a correct prediction of what will be (optimally) done at $q_{i\tau}$ from the perspective of q_{it} and θ_i^* . Thus these strategies form an equilibrium of the original dynamic game. \square

Proof of Theorem 4.4: As in the proof of Theorem 4.2, we reinterpret the dynamic game as a static one. By assumption, such a reinterpretation is possible; thus the only remain-

ing objective is to show that the symmetry of players is preserved in all subgames. We proceed inductively. Let I_t be the set of players who are symmetric at q_t . In period T when $q_{\pi(I_T),T} = q_{I_T,T}$, we immediately have symmetry as this is essentially a static environment, since current actions do not affect the future. Hence for $i, j \in I_T$, for any given $x_T^i \in X_T^i = X_T^j$, and $q_{iT} = q_{jT}$,

$$\begin{aligned} & \int u_i(H^T, x_T^i, \tilde{x}_{T,T}^{-i}; \tilde{\theta}_i^T, \tilde{\theta}_{-i}^T) g_T(\alpha_{-i} | q_{-iT}, \alpha_i) d\alpha_{-i} d\mu_T^i(\Theta_{-i}, H^T | q_{iT}, \tilde{\theta}_i^T(q_{iT}, \alpha_i)) \\ &= \int u_j(H^T, x_T^j, \tilde{x}_{T,T}^{-j}; \tilde{\theta}_j^T, \tilde{\theta}_{-j}^T) g_T(\alpha_{-j} | q_{-jT}, \alpha_j) d\alpha_{-j} d\mu_T^j(\Theta_{-j}, H^T | q_{jT}, \tilde{\theta}_j^T(q_{jT}, \alpha_j)) \end{aligned}$$

Now suppose that players $i \in I_t$ use symmetric monotone strategies in period t , and that we restrict our attention to subgames given by $Q_{\pi(I_t),t+1}$ that are symmetric in the sense of $\pi(I_{t+1})$, i.e. the strategies given by $C^t(q_{t+1}, \theta)$ are permuted by $\pi(I_t)$ if x_t and $\mu_{t+1}(\cdot | q_{t+1}, \theta)$ are permuted via $\pi(I_t)$. Then the distribution of continuations subgames starting from period t will be symmetric (in the sense of Condition (4)), implying that the incentives in period t given by $U_i^t(q_{it}, x_t^i, \theta_i)$ are symmetric (in the sense of Condition (3)), i.e. for $i, j \in \pi(I_t)$, choosing the vector of actions $\{\tilde{x}_{\tau,t}^i\}_{\tau=t}^T$ at $q_{it} = q_{jt}$ yields, when payoffs are given by equation (2),

$$\begin{aligned} & \int u_i(H^t, \{\tilde{x}_{\tau,t}^i, \tilde{x}_{\tau,t}^{-i}\}; \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i} | q_{-it}, \alpha_i) d\alpha_{-i} d\mu_t^i(\Theta_{-i}, H^t | q_{it}, \tilde{\theta}_i^t(q_{it}, \alpha_i)) \\ &= \int u_j(H^t, \{\tilde{x}_{\tau,t}^j, \tilde{x}_{\tau,t}^{-j}\}; \tilde{\theta}_j^t, \tilde{\theta}_{-j}^t) g_t(\alpha_{-j} | q_{-jt}, \alpha_j) d\alpha_{-j} d\mu_t^j(\Theta_{-j}, H^t | q_{jt}, \tilde{\theta}_j^t(q_{jt}, \alpha_j)) \end{aligned}$$

This implies symmetry of the subgame in period t where $q_{\pi(I_t),t} = q_{I_t,t}$. 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 4.5). To translate this into the dynamic game, we apply Lemmas 4.6 and 4.7 to show that we can generate consistent strategies $\{\tilde{x}_{\tau,t}^i\}$ such that the payoffs as given by equation (2) match those given by equation (4). Thus there will exist a monotone PBE which is symmetric, i.e. for $i, j \in I$, if in $q_{I,t}$, player i follows strategy $\mathbf{x}_t^i(q_{it}, \cdot)$ and player j follows strategy $\mathbf{x}_t^j(q_{jt}, \cdot)$, then at $q_{\pi(I),t}$, i follows strategy $\mathbf{x}_t^j(q_{jt}, \cdot)$ and j follows $\mathbf{x}_t^i(q_{it}, \cdot)$ (randomizing with the same probabilities if necessary). \square

Proof of Theorem 4.5: 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 q_{it} ; by Assumption 4.2, there are a countable number of such players in Γ_m . We define ψ as in the finite case; the restriction of ψ to T_m periods will be continuous if ψ is, and so we may use the same ψ for every Γ_m , as well as Γ itself. For each indexed player, the equilibrium function $\sigma_{t,m}^i$ (as defined in Section 4) 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 Lemma 4.4 and Theorem 4.1, and $\hat{x}_{\tau,t,m}^i \rightarrow \hat{x}_{\tau,t}^i$ for all $\tau \geq t$ by Lemma 4.3. We check that the limit strategies form an equilibrium in the static game in which payoffs are given by (4).

The last item that must be checked is that payoffs in this sequence converge to an equilibrium as $m \rightarrow \infty$. 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.²³ By continuity of payoffs, we have that for any $\varepsilon > 0$ and any t , there exists M such that for all $m > M$,

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

²³ See Fudenberg and Levine (1983), Lemma 4.1 for the proof.

$$< \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3} = \varepsilon$$

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 4.2 and 4.3 and Theorem 4.1. We then translate this into an equilibrium of the dynamic game in the manner analogous to Theorem 4.2 using Lemmas 4.6 and 4.7.

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

Proof of Theorem 5.1: Consider any sequence of equilibrium consistent monotone strategy profiles for period t , $\{\sigma_{t,m}(\cdot)\} = \{\tilde{\mathbf{x}}_{\tau,t}^i, \tilde{\theta}_t^i\}_{i,\tau,t,m}$ in the static game with payoffs given by (4). Given Assumptions 5.1-5.3, and the fact that $\bigcup_{m=1}^{\infty} \{(X_{\tau,m}, Q_{\tau,m})\}_{\tau=1}^T$ is countable, we can use Helly's selection theorem in conjunction with Tychonoff's theorem to generate a convergent subsequence to some monotone strategy profile $\sigma_t(\cdot)$ (without loss of generality, this can be the sequence $\{\sigma_{t,m}(\cdot)\}$ itself). This is generated inductively in τ , exploiting in the following manner the fact that in any period, there can only be a countable number of actions x_{τ}^i for which a positive measure of types θ_i (with respect to the prior F) choose the same action. Specifically, fix $\{\tilde{\mathbf{x}}_{t',t'}^i\}_{i,t'}$ for $t' < \tau$ as the limit strategy profile of all players restricted to those periods before τ , $\lim_{m \rightarrow \infty} \{\tilde{\mathbf{x}}_{t',t',m}^i\}_{i,t'}$. Then we look at sequences $\{q_{\tau,m}\}_{m=1}^{\infty}$ as follows. If there is a positive measure of types α_i from which q_{τ} is reachable by the limit strategy $\{\tilde{\mathbf{x}}_{t',t'}^i\}_{t'}$ (i.e. $\bar{\alpha}_i(q_{i\tau}, q_{i1}) - \underline{\alpha}_i(q_{i\tau}, q_{i1}) > 0$), then the sequence $\{q_{\tau,m}\}_{m=1}^{\infty}$ assigns the value to $q_{i\tau,m}$ as given by some $\hat{q}_{i\tau,m} \in Q_{i\tau,m}$ which are reachable for $\alpha_i \in (\underline{\alpha}_i(q_{i\tau}, q_{i1}), \bar{\alpha}_i(q_{i\tau}, q_{i1}))$; by Assumption 5.4 this sequence is well-defined as all such α_i choose the same actions $\{x_{t',m}^i\}_{t' < \tau}$ for sufficiently high m . Otherwise, for $q_{\tau} \in \bigcup_{m=1}^{\infty} Q_{\tau,m}$, we set $\tilde{\mathbf{x}}_{\tau,\tau,m}^i(q_{i\tau}, \alpha_i)$ to be equal to $\tilde{\mathbf{x}}_{\tau,\tau,m}^i(q_{i\tau,m}, \alpha_i)$, where $q_{i\tau,m} = \max\{\hat{q}_{i\tau,m} \in Q_{i\tau,m} : \hat{q}_{i\tau,m} \leq q_{i\tau}\}$. Thus, given $q_{\tau-1}$, the set of possible $\{q_{\tau}\}$ will be given by the information events generated by the Cartesian product of actions $\{x_{\tau}^i\}$ that are contained in the union of $\bigcup_{m=1}^{\infty} X_{\tau,m}^i$ and the set of $\{x_{\tau}^i\}$ that were played by a subset $\{\alpha_i\} \subset [0, 1]$ of positive measure as limits of actions $x_{\tau,m}^i \in X_{\tau,m}^i$.

We show inductively in τ that the set of $\{q_{t'}\}_{t' < \tau}$ as generated in the above paragraph is countable for each τ , and hence the number of such sequences $\{q_{t',m}\}_{t' < \tau}$ is countable as well. For $\tau = 1$, there is only one possible sequence (as no actions have yet been taken). Now suppose that the induction hypothesis is true up to period $\tau - 1$. Because for each $q_{\tau-1}$, there can only be a countable number of values of $x_{\tau-1}^i$ for which a positive measure of

types α_i can reach the same q_τ , this implies that only a countable number of informational events q_τ need be considered.

As in Section 4, we can define the beliefs by the corresponding mappings ψ and ψ_m in the limit game and in the discretizations, respectively. As seen in Example 3.2, there is a potential for belief entanglement in the limit, in that the supports of types that choose a given action x_τ^j in period τ may be different in the limit from that in any of the discretizations. However, Assumption 5.4 rules this out. We formalize this in the following two lemmas, broken down by whether x_τ^j is off-path or on-path in the limit game via $\tilde{x}_{\tau,\tau}^j(q_{j\tau}, \cdot)$.

Lemma 5.1: *Suppose that x_τ^j is fully revealed to player i in period $t > \tau$, but $(x_\tau^j)^*$ is off-path at $q_{j\tau}$ in the limit game. Then the beliefs of player i at q_{it} over (θ_j, H^t) such that H^t generates $q_{j\tau}$, upon observing $(x_\tau^j)^*$, converge in the weak-* topology.*

Proof: The argument is very similar to that of Theorem 4.1 in establishing beliefs that must hold off-path for continuous ψ . Consider the sequence $\{\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \cdot)\}$. Then for $(x_\tau^j)^*$ that is off-path, let $\alpha_j^* = \sup\{\alpha_j : \tilde{x}_{\tau,\tau}^j(q_{j\tau}, \alpha_j) < (x_\tau^j)^*\}$ and let $\theta_j^* = \tilde{\theta}_j^\tau(q_{j\tau}, \alpha_j^*)$. By Helly's selection theorem, for any $\alpha_j < \alpha_j^*$, there exists M such that for all $m > M$, it will be the case that $\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \alpha_j) < x_\tau^j$ and $\tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \alpha_j) \leq \theta_j^*$. Similarly, for all $\alpha_j > \alpha_j^*$, there exists M such that for all $m > M$, then $\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \alpha_j) > x_\tau^j$ and $\tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \alpha_j) \geq \theta_j^*$. Hence for $H^t \ni x_\tau^j$, by the construction of the beliefs via ψ and ψ_m , for any open $A \subset \Theta_j$ for which $\theta_j^* \in A$,

$$\frac{\mu_t^i(\theta_j^*, H^t | q_{it}, \theta_i)}{\mu_t^i(\Theta_j, H^t | q_{it}, \theta_i)} = \lim_{m \rightarrow \infty} \frac{\mu_{t,m}^i(A, H_m^t | q_{it,m}, \theta_i)}{\mu_{t,m}^i(\Theta_j, H_m^t | q_{it,m}, \theta_i)} = 1$$

where the above limit is taken in the weak-* topology. \square

Lemma 5.2: *Suppose that x_τ^j is fully revealed, but $(x_\tau^j)^*$ is on-path at some on-path $q_{j\tau}$ in the limit game. Then the beliefs of player i at q_{it} over (θ_j, H^t) such that H^t generates $q_{j\tau}$, upon observing $(x_\tau^j)^*$, converge in the weak-* topology, i.e. $\mu_{t,m}^i(\cdot, H_m^t | q_{it,m}, \theta_i) \rightarrow \mu_t^i(\cdot, H^t | q_{it}, \theta_i)$ where $H_m^t \rightarrow H^t$ and $q_{it,m} \rightarrow q_{it}$.*

Proof: The argument is very similar to that of Lemma 4.2. Suppose that the hypothesis holds, but that t is the first period in which beliefs diverge at some q_{it} . By Helly's selection theorem, $\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \cdot) \rightarrow \tilde{x}_{\tau,\tau}^j(q_{j\tau}, \cdot)$ pointwise almost everywhere for $\tau < t$. Because strategies are monotone, the support of player i 's beliefs over types θ_{-i} given any particular history must be a product of intervals by Lemma 4.1. Suppose that under the limit strategy, $\tilde{x}_{\tau,\tau}^j(q_{j\tau}, \alpha_j) = \tilde{x}_{\tau,\tau}^j(q_{j\tau}, \hat{\alpha}_j)$ but $\tilde{\theta}_j^\tau(q_{j\tau}, \alpha_j) < \tilde{\theta}_j^\tau(q_{j\tau}, \hat{\alpha}_j)$. Without loss of gener-

ality, since the support of beliefs over θ_j must be an interval, we can let α_j and $\hat{\alpha}_j$ be the minimal and maximal values (respectively) which satisfy this property. Then by Assumption 5.4, for all $\alpha_j^1 \notin [\alpha_j, \hat{\alpha}_j]$ and $\alpha_j^2 \in (\alpha_j, \hat{\alpha}_j)$, there exists M such that for all $m > M$, $\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \alpha_j^1) \neq \tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \alpha_j^2)$, and so $\tilde{x}_{\tau,\tau}^j(q_{j\tau}, \alpha_j^1) \neq \tilde{x}_{\tau,\tau}^j(q_{j\tau}, \alpha_j^2)$. Since, $\tilde{\theta}_j^\tau$ is increasing in α_j (strictly so because $\tilde{\theta}_j^\tau(q_{j\tau}, \alpha_j) < \tilde{\theta}_j^\tau(q_{j\tau}, \hat{\alpha}_j)$) and $\tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \cdot) \rightarrow \tilde{\theta}_j^\tau(q_{j\tau}, \cdot)$ by Lemma 4.3, we conclude that for H^t which generates q_{it} and H_m^t which generates $q_{it,m}$, where $H_m^t \rightarrow H^t$,

$$\begin{aligned} & \mu_t^i((\tilde{\theta}_j^\tau(q_{j\tau}, \alpha_j), \tilde{\theta}_j^\tau(q_{j\tau}, \hat{\alpha}_j)), \Theta_{-\{i,j\}}, H^t | q_{it}, \theta_i) \\ & \leq \lim_{m \rightarrow \infty} \mu_{t,m}^i((\tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \alpha_j), \tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \hat{\alpha}_j)), \Theta_{-\{i,j\}}, H_m^t | q_{it,m}, \theta_i) \end{aligned}$$

as it is clear that every type $\alpha_j' \notin [\alpha_j, \hat{\alpha}_j]$ will be excluded from the support of $\lim_{m \rightarrow \infty} \mu_{t,m}^i$. On the other hand, for each $\varepsilon > 0$, there must exist M such that for $m > M$, $\tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \alpha_j + \varepsilon) = \tilde{x}_{\tau,\tau,m}^j(q_{j\tau,m}, \hat{\alpha}_j - \varepsilon)$. Therefore, all types $\alpha_j' \in (\alpha_j, \hat{\alpha}_j)$ are included in the support of $\lim_{m \rightarrow \infty} \mu_{t,m}^i$, and so

$$\begin{aligned} & \mu_t^i((\tilde{\theta}_j^\tau(q_{j\tau}, \alpha_j), \tilde{\theta}_j^\tau(q_{j\tau}, \hat{\alpha}_j)), \Theta_{-\{i,j\}}, H^t | q_{it}, \theta_i) \\ & \geq \lim_{m \rightarrow \infty} \mu_{t,m}^i((\tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \alpha_j), \tilde{\theta}_{j,m}^\tau(q_{j\tau,m}, \hat{\alpha}_j)), \Theta_{-\{i,j\}}, H_m^t | q_{it,m}, \theta_i) \end{aligned}$$

Combining these two inequalities, we find that $\mu_{t,m}^i \rightarrow \mu_t^i$ in the weak-* topology, and so beliefs do not diverge at q_{it} . \square

Returning to the proof of Theorem 5.1 we extend the argument to the cases of q_τ not covered above. To generate the limit strategies for those q_τ that are covered, we have already defined a convergent subsequence of strategies (again, without loss of generality, the sequence itself, $\{\tilde{x}_{\tau,\tau,m}\}$). We consider two possible additional cases: that there is some type α_j that plays the given x_τ^j not included in the above set, or that no α_j plays that action. In the former case, note that for every element of the sequence Γ_m , player j is choosing some action $x_{\tau,m}^j \in X_{\tau,m}^j$. Thus we set the strategies $\tilde{x}_{t',t'}^j(q_{t'}, \cdot)$ for subgames $q_{t'}$ where $t' > \tau$ as the limit of the sequence of strategies $\{\tilde{x}_{t',t',m}^j(q_{t',m}, \cdot)\}$ chosen at the subgames $q_{t',m}$ that were reached when $x_{\tau,m}^j$ was chosen at $q_{\tau,m}$ in Γ_m ; such a limit is well-defined by Helly's selection theorem. Similarly, for the case where there is no type that chooses x_τ^j , we set $\tilde{x}_{t',t'}^j(q_{t'}, \cdot)$ to be the limit of the sequence of strategies $\{\tilde{x}_{t',t',m}^j(q_{t',m}, \cdot)\}$ that were reached by choosing $x_{\tau,m}^j = \sup\{x_{\tau,m}^j \in X_{\tau,m}^j : x_{\tau,m}^j < x_\tau^j\}$. Since in either case, these subgames $q_{t'}$ are off-path, the same logic as in Lemma 5.1 applies to the convergence of beliefs over the type of player j .

Finally, since beliefs do not diverge at any q_{it} , it follows that $\hat{\mathbf{x}}_{\tau,t,m}^i$ converges to $\hat{\mathbf{x}}_{\tau,t}^i$ by Lemma 4.3, as one recalls that $\hat{\mathbf{x}}_{\tau,t}^i$ is constructed in the same manner as beliefs via ψ . Therefore, for the sequence of strategies $\{\tilde{\mathbf{x}}_{\tau,t,m}^i\}_{m=1}^\infty$, by the continuity of u_i it will be the case that

$$\begin{aligned} \lim_{m \rightarrow \infty} \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t,m}^i, \hat{\mathbf{x}}_{\tau,t,m}^{-i}\}_{\tau=t}^T, \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i}|q_{it,m}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t|q_{it,m}, \tilde{\theta}_{i,m}^t) \\ = \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i, \hat{\mathbf{x}}_{\tau,t}^{-i}\}_{\tau=t}^T, \tilde{\theta}_i^t, \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|q_{it}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t|q_{it}, \tilde{\theta}_i^t) \end{aligned}$$

Hence the limit of the choices of actions in period t in games $\{\Gamma_m\}$ according to strategies $\{\tilde{\mathbf{x}}_{\tau,t,m}^i\}$ remains optimal in the limit game Γ . We then reconstruct the equilibrium in the dynamic game in the same manner as in Section 4 since $\{\tilde{\mathbf{x}}_{\tau,t}^i\}_{i,\tau,t}$ must be consistent, as each element in the sequence was consistent, and $\tilde{\theta}_{i,m}^t(q_{it,m}, \cdot) \rightarrow \tilde{\theta}_i^t(q_{it}, \cdot)$ for all i, t since beliefs converged by Lemmas 5.1 and 5.2. \square

Proof of Proposition 5.2: We show for the case that u_i is strictly increasing in x_τ^{-i} for $\tau > t$. Suppose that beliefs do not converge for x_t^i on-path. Let $\alpha_i, \hat{\alpha}_i$ be two such types such that $\tilde{\mathbf{x}}_{t,t}^i(q_{it}, \alpha_i) = \tilde{\mathbf{x}}_{t,t}^i(q_{it}, \hat{\alpha}_i)$ and $\tilde{\theta}_i^t(q_{it}, \hat{\alpha}_i) > \tilde{\theta}_i^t(q_{it}, \alpha_i)$ but such that for all M , there exists m such that $\tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \alpha_i) \neq \tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \hat{\alpha}_i)$. Without loss of generality, we can assume that this is true for all such m by taking the appropriate subsequence. Suppose further without loss of generality that $\hat{\alpha}_j$ is the supremum of all such types, and α_j is the infimum of all such types satisfying the above conditions.

Since $\tilde{\mathbf{x}}_{t,t,m}^i(\cdot)$ is monotone, if beliefs do not converge, by Lemma 5.2 it must be that $\tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \alpha_i) < \tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \hat{\alpha}_i)$. For convenience, we write

$$\{H_m^{t-1}, \tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \hat{\alpha}_i), \tilde{\mathbf{x}}_{t,t,m}^{-i}(q_{-it,m}, \alpha_{-i})\} = \hat{H}_m^{t+1}$$

$$\{H_m^{t-1}, \tilde{\mathbf{x}}_{t,t,m}^i(q_{it,m}, \alpha_i), \tilde{\mathbf{x}}_{t,t,m}^{-i}(q_{-it,m}, \alpha_{-i})\} = H_m^{t+1}$$

By monotonicity across subgames, it must consequently be that for the corresponding $\hat{Q}_{t,m} > Q_{t,m}$,

$$\lim_{m \rightarrow \infty} C_m^t(\hat{Q}_{t+1,m}, \alpha) \geq \lim_{m \rightarrow \infty} C_m^t(Q_{t+1,m}, \alpha) \quad (6)$$

where $C_m^t(\cdot)$ describes the continuation play for all players from date t (exclusive) via strategies $\{\tilde{\mathbf{x}}_{\tau,t}^j\}$. Suppose that players $-i$ play the limits of their strategies $\tilde{\mathbf{x}}_{\tau,t}^{-i}(q_{-i\tau}, \alpha_{-i}) = \lim_{m \rightarrow \infty} \tilde{\mathbf{x}}_{\tau,t,m}^{-i}(q_{-i\tau,m}, \alpha_{-i})$ in Γ . Then if inequality (6) is strict for C_m^t for a positive measure

of α_{-i} given other players' strategies, it must be that

$$\begin{aligned}
& \lim_{m \rightarrow \infty} \int u_i(\hat{H}_m^{t+1}, C_m^t(\hat{Q}_m^{t+1}, \alpha), \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | q_{it,m}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | q_{it,m}, \tilde{\theta}_i(q_{it}, \alpha_i)) \\
& > \lim_{m \rightarrow \infty} \int u_i(\hat{H}_m^t, C_m^t(Q_m^{t+1}, \alpha), \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | q_{it,m}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | q_{it,m}, \tilde{\theta}_i(q_{it}, \alpha_i)) \\
& = \lim_{m \rightarrow \infty} \int u_i(H_m^t, C_m^t(Q_m^{t+1}, \alpha), \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | Q_{it,m}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | Q_{it,m}, \tilde{\theta}_i(Q_{it}, \alpha_i))
\end{aligned}$$

where the first inequality follows from the fact that payoffs are strictly increasing in x_τ^{-i} , and the second equality follows from the continuity of u_i . Hence type α_i would want to deviate at some $m > \infty$, contradicting the claim that $(x_m(\cdot), \mu_m)$ was an equilibrium strategy profile. Hence it must be that for almost all θ_i , $\lim_{m \rightarrow \infty} C_m^t(\hat{Q}_{t+1,m}, \alpha) = \lim_{m \rightarrow \infty} C_m^t(Q_{t+1,m}, \alpha)$.

This establishes convergence of period $\tau > t$ strategies on path. Meanwhile, since beliefs converge off-path by Lemma 5.1, one can set up a limit of the strategies $\tilde{x}_{\tau,t}^i(q_{i\tau}, \alpha_i) = \lim_{m \rightarrow \infty} \tilde{x}_{\tau,t,m}^i(q_{i\tau}, \alpha_i)$ by Helly's selection theorem. By continuity of u_i , and the convergence of beliefs, this will imply that if $\sigma_{\tau,m}$ was an equilibrium strategy profile for every m in the subgame following q_τ , then $\sigma_\tau = \lim_{m \rightarrow \infty} \sigma_{\tau,m}$ is also an equilibrium strategy profile in the subgame following q_τ . Since τ was arbitrary, $\sigma = \lim_{m \rightarrow \infty} \sigma_m$ will be an equilibrium strategy profile of Γ .

The proof for when u_i is strictly decreasing in x_τ^{-i} is analogous, and so is omitted. \square

Proof of Theorem 5.3: Consider the sequence of discretizations $\{\Gamma_m\}$ which satisfy Assumptions 5.1-5.3. These generate an induced equilibrium strategy profile σ_m . Suppose that $X_{\tau,m}^j$ has $K_{j,m}$ elements, given by $\{a_{k,m}^j\}_{k=1}^{K_{j,m}}$ where $a_{k+1,m}^j > a_{k,m}^j$ and is observable at $q_{it,m}$. Since an equilibrium exists regardless of the discretization, we can assume that $a_{1,m}^j = \min\{x_{\tau,m}^j \in X_{\tau,m}^j\}$ and $a_{K_{j,m},m}^j = \max\{x_{\tau,m}^j \in X_{\tau,m}^j\}$. Then we extend $\mu_{t,m}^i(\cdot | \cdot)$ to any observation $x_\tau^j \in X_\tau^j$ (and hence to q_{it} defined by replacing $x_{\tau,m}^j$ with x_τ^j) by stipulating that, for any $x_\tau^j \in (a_{k,m}^j, a_{k+1,m}^j)$, $\mu_{t,m}^i(\cdot, \{H^t : x_\tau^j \in H^t\} | q_{it}, \theta_i) = \mu_{t,m}^i(\cdot, \{H^t, a_{k,m}^j, x_{\tau-1}^{-i}\} | q_{it,m}, \theta_i)$, i.e. the beliefs for any action in that interval will be the same. We then stipulate that for any subgame following x_τ^j , all players play the same strategies $\sigma_{t,m}$ as those following $a_{k,m}^j$. Thus, we have essentially treated any action in the interval $(a_{k,m}^j, a_{k+1,m}^j)$ as being equivalent to taking action $a_{k,m}^j$, and so we can assume without loss of generality that all players j choose some action in $X_{\tau,m}^j$ in period τ . Thus the distribution over continuations from any period t , given by $C_{i,m}^t$, will be the same.

By Helly's selection theorem, there exists a subsequence of $\{\tilde{x}_{\tau,t,m}\}_{m=1}^\infty$ (without loss of

generality, the sequence itself) which converges to a well-defined limit $\tilde{x}_{\tau,t}$ for all $t \geq \tau$ for any history generated by values of $x_{\tau,m} \in \bigcup_{m=1}^{\infty} X_{\tau,m}$. By Assumption 5.2, for all $\delta > 0$, there exists M such that for all $m > M$, $\max_{j,k} \{a_{k+1,m}^j - a_{k,m}^j\} < \delta$. By the continuity of u_i in x , for any $\varepsilon > 0$, there exists $\delta > 0$ such that for any x and \hat{x} , if $\|x - \hat{x}\| < 2NT\delta$, then $|u_i(x, \theta) - u_i(\hat{x}, \theta)| < \varepsilon$. By the optimality of $\tilde{x}_{\tau,t,m}$ and the compactness of X , then for sufficiently high M , for all θ_i , it follows that (checking for one-stage deviations)

$$\begin{aligned}
& (x_{t,m}^i)^* \in \arg \max_{x_{t,m}^i \in X_{t,m}^i} U_{t,m}^i(q_{it,m}, x_{t,m}^i, \theta_i) \\
& \equiv \int u_i(H^t, x_{t,m}^i, \tilde{\mathbf{x}}_{t,t,m}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t,m}\}_{\tau=t+1}^T, \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | q_{it,m}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | q_{it,m}, \tilde{\theta}_{i,m}^t(q_{it,m}, \alpha_i)) \\
& \implies \int u_i(H^t, (x_{t,m}^i)^*, \tilde{\mathbf{x}}_{t,t,m}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t,m}\}_{\tau=t+1}^T, \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | q_{it}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | q_{it}, \tilde{\theta}_{i,m}^t(q_{it}, \alpha_i)) \\
& > \int u_i(H^t, x_{t,m}^i, \tilde{\mathbf{x}}_{t,t,m}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t,m}\}_{\tau=t+1}^T, \tilde{\theta}_{i,m}^t, \tilde{\theta}_{-i,m}^t) g_{t,m}(\alpha_{-i} | q_{it}, \alpha_i) d\alpha_{-i} d\mu_{t,m}^i(\Theta_{-i}, H^t | q_{it}, \tilde{\theta}_{i,m}^t(q_{it}, \alpha_i)) - \varepsilon
\end{aligned}$$

for all $x_t^i \in X_t^i$. Hence the strategy profile σ_m is an ε -PBE. \square

Proof of Proposition 6.1: At subgame q_{it} , 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 q_{it} is absolutely continuous, and that²⁴

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

In the case that θ is affiliated, μ_t^i will be increasing in MLR in θ_i given H^t , since H^t is perfectly observed and 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

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

The best-reply in period t for player i will therefore be increasing in θ_i in the SSO by Lemma 6.1(a). The proof for the case where some type θ_j might be completely atomic is analo-

²⁴Because C^t is irrelevant, we can substitute this into u_i without affecting the payoffs.

gous²⁵, and therefore omitted. \square

Proof of Proposition 6.2: At subgame q_{it} , suppose that $\hat{x}_t^i \geq x_t^i$ and $\tilde{\theta}_i^t(q_{it}, \hat{\alpha}_i) \equiv \hat{\theta}_i > \theta_i \equiv \tilde{\theta}_i^t(q_{it}, \alpha_i)$.²⁶ Suppose that²⁷

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

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

Suppose that $\hat{x}_t^i \neq (x_t^i)^*$. By revealed preference, type α_i prefers to follow his continuation strategy (given by $\tilde{\mathbf{x}}_{\tau,t}^i(q_{i\tau}, \alpha_i)$) after choosing x_t^i instead of that of $\hat{\alpha}_i$. Moreover, by future irrelevance, α_i would have the same payoff if everyone continued by playing $\{\tilde{\mathbf{x}}_{\tau,t}(q_\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T$ after choosing \hat{x}_t^i . Therefore,

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

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

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

Lastly, due to future irrelevance after \hat{x}_t^i , we replace Q_τ with \hat{Q}_τ , so that

$$\int u_i(H^t, \hat{x}_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(\hat{Q}_\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(q_{it}, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|q_{it}) d\alpha_{-i} d\mu_t^i(\Theta_{-i}, H^t|q_{it})$$

²⁵The only difference is that there may be some mixing by player j . Since the conditions of the proposition allow for aggregation of single-crossing since they satisfy the conditions of Lemma 6.1(b), this will not make a difference.

²⁶The proof for θ_i outside the support of $g_t(\cdot|q_t)$ is identical.

²⁷We suppress types in g_t where possible due to the independence of the distribution of θ .

$$- \int u_i(H^t, x_t^i, \tilde{\mathbf{x}}_{t,t}^{-i}, \{\tilde{\mathbf{x}}_{\tau,t}(Q_\tau, \hat{\alpha}_i, \alpha_{-i})\}_{\tau=t+1}^T; \tilde{\theta}_i^t(q_{it}, \hat{\alpha}_i), \tilde{\theta}_{-i}^t) g_t(\alpha_{-i}|q_{it}) d\alpha_{-i} d\mu_t^i(\Theta_{-i}, H^t|q_{it}) \geq 0$$

The case where $x_t^i \neq (x_t^i)^*$ is analogous, where we first note that one can replace $\{\tilde{\mathbf{x}}_{\tau,t}(Q_\tau, \alpha)\}_{\tau=t+1}^T$ with $\{\tilde{\mathbf{x}}_{\tau,t}(\hat{Q}_\tau, \alpha)\}_{\tau=t+1}^T$ 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 6.3: 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 Q_{i2} , so we have (slightly abusing notation)

$$U_i^1(x_1^i, \theta_i) = \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(Q_{i2}, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}, \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i}$$

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 generality for the purposes of his optimization as of period 1 that $\mathbf{x}_2^i(Q_{i2}, \theta_i)$ is a singleton. By monotonicity within and across subgames in period 2 and affiliation of (Q_2, x_1, θ) , it must be that $\mathbf{x}_2^i(Q_{i2}, \theta_i)$ and $\tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i})$ are increasing in all arguments. Since player i observes the same information about x_1^{-i} regardless of his own action, we can set \hat{Q}_2 to be the information that follows from $(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i})$, and Q_2 to be the information that follows from $(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i})$. Suppose that

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{Q}_{i2}, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{Q}_{-i2}, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(Q_{i2}, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \geq 0 \end{aligned}$$

For any α_{-i} , since each player's strategy is chosen independently from the others, $\hat{Q}_{-i2} > Q_{-i2}$ for all α_{-i} . By monotonicity within and across subgames, it follows that $\tilde{\mathbf{x}}_{2,1}^{-i}(\hat{Q}_{-i2}, \hat{\alpha}_{-i}) \geq \tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i})$ for $\hat{\alpha}_{-i} \geq \alpha_{-i}$.

A possible complication is that we do not know whether $\mathbf{x}_2^i(\hat{Q}_{i2}, \theta_i) \geq \mathbf{x}_2^i(Q_{i2}, \hat{\theta}_i)$ or vice versa. To address this, let

$$\check{\mathbf{x}}_2^i(\hat{Q}_{i2}, \theta_i; Q_{i2}, \hat{\theta}_i) = \mathbf{x}_2^i(\hat{Q}_{i2}, \theta_i) \vee \mathbf{x}_2^i(Q_{i2}, \hat{\theta}_i)$$

$$\hat{\mathbf{x}}_2^i(\hat{Q}_{i2}, \theta_i; Q_{i2}, \hat{\theta}_i) = \mathbf{x}_2^i(\hat{Q}_{i2}, \theta_i) \wedge \mathbf{x}_2^i(Q_{i2}, \hat{\theta}_i)$$

By revealed preference, since $\mathbf{x}_2^i(Q_{i2}, \theta_i)$ is optimal for θ_i upon reaching Q_{i2} in period 2,

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{Q}_{i2}, \theta_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{Q}_{-i2}, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \hat{\mathbf{x}}_2^i(\hat{Q}_{i2}, \theta_i; Q_{i2}, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i}), \theta_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \geq 0 \end{aligned}$$

Since one can aggregate supermodularity/ID under integration, we have

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \hat{\mathbf{x}}_2^i(\hat{Q}_{i2}, \theta_i; Q_{i2}, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{Q}_{-i2}, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(Q_{i2}, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \geq 0 \end{aligned}$$

By revealed preference again, since $\mathbf{x}_2^i(\hat{Q}_{i2}, \hat{\theta}_i)$ is optimal for $\hat{\theta}_i$ upon reaching \hat{Q}_{i2} ,

$$\begin{aligned} & \int u_i(\hat{x}_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(\hat{Q}_{i2}, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(\hat{Q}_{-i2}, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \\ & - \int u_i(x_1^i, \tilde{\mathbf{x}}_{1,1}^{-i}, \mathbf{x}_2^i(Q_{i2}, \hat{\theta}_i), \tilde{\mathbf{x}}_{2,1}^{-i}(Q_{-i2}, \alpha_{-i}), \hat{\theta}_i, \tilde{\theta}_{-i}^1) g_1(\alpha_{-i}|q_{i1}) d\alpha_{-i} \geq 0 \end{aligned}$$

Putting all of this together, we find that

$$U_i(\hat{x}_1^i, \theta_i) - U_i(x_1^i, \theta_i) \geq 0 \implies U_i(\hat{x}_1^i, \hat{\theta}_i) - U_i^1(x_1^i, \theta_i) \geq 0$$

Hence the optimal action will be increasing in the strong set order in period 1, as shown in Lemma 6.1(a).

To show that best replies are monotone within and across subgames in period 2, suppose that $\hat{x}_2^i \geq x_2^i$, and that

$$\begin{aligned} & \int u_i(x_1, \hat{x}_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \theta_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i}|q_{i2}) d\alpha_{-i} d\mu_2^i(\Theta_{-i}, x_1|q_{i2}) \\ & - \int u_i(x_1, x_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \theta_i, \tilde{\theta}_{-i}^2) g_t(\alpha_{-i}|q_{i2}) d\alpha_{-i} d\mu_2^i(\Theta_{-i}, x_1|q_{i2}) \geq 0 \end{aligned}$$

We know that beliefs are increasing in MLR in q_{i2} because period-1 actions are increasing in type.²⁸ Since $\tilde{\mathbf{x}}_{2,2}^{-i}$ and $\tilde{\theta}_{-i}$ are increasing in q_{-i2} and α (by monotonicity within and across subgames), the induced distribution of (x_2^{-i}, θ_{-i}) conditional upon observing \hat{q}_{i2} will first-order stochastically dominate that from upon observing q_{i2} . Hence we find that

²⁸See Milgrom (1981), Proposition 4, for the details.

(by aggregating the supermodularity and ID conditions under integration, which is possible between distributions ordered under first-order stochastic dominance) that

$$\begin{aligned} & \int u_i(x_1, \hat{x}_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \hat{\theta}_i, \tilde{\theta}_{-i}^2) g_i(\alpha_{-i} | \hat{q}_{i2}) d\alpha_{-i} d\mu_2^i(\Theta_{-i}, x_1 | \hat{q}_{i2}) \\ & - \int u_i(x_1, x_2^i, \tilde{\mathbf{x}}_{2,2}^{-i}, \hat{\theta}_i, \tilde{\theta}_{-i}^2) g_i(\alpha_{-i} | \hat{q}_{i2}) d\alpha_{-i} d\mu_2^i(\Theta_{-i}, x_1 | \hat{q}_{i2}) \geq 0 \end{aligned}$$

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 6.1(a). \square

Proof of Proposition 7.1: The payoff functions satisfy the conditions of Proposition 6.3, so there exists monotone best-replies for all players to monotone strategies of the other players. Since period 1 actions are perfectly observable, there exists monotone and continuous ψ in this game, so there will exist a monotone PBE by Theorem 4.2 in the game with a finite number of actions. To extend to a continuum of actions, note that for player 1, u_1 is strictly increasing in x_2 . Hence we can invoke Proposition 5.2 to generate existence of a monotone equilibrium even in the case of the existence of a continuum of actions. \square

Proof of Lemma 7.1: Suppose that there is pooling at some $x_1^1 \neq \bar{x}_1^1$. Let $I(q_{j2})$ be the conditional support over θ_1 given observing q_{j2} . By Lemma 4.1, we can set $I(x_1^1) \equiv [\theta_1^1, \theta_1^2]$. Then for any $\hat{x}_1^1 > x_1^1$, $I(\hat{x}_1^1) \equiv [\theta_1^3, \theta_1^4]$, where $\theta_1^3 \geq \theta_1^2$. Moreover, if $\hat{x}_1 > x_1$, $\mathbf{x}_2^j(\hat{x}_1, \theta_j) - \mathbf{x}_2^j(x_1, \theta_j) \geq 0$ by monotonicity across subgames. By differentiability in x_2^j and strictly increasing differences in (x_2^j, θ_1) , this inequality must be strict and bounded away from 0. By continuity in x_1^1 and θ_1 , and the fact that u_1 is strictly increasing in x_j , type θ_1^2 has a profitable deviation, as

$$\lim_{\hat{x}_1^1 \rightarrow (x_1^1)^+} u_1^1(\hat{x}_1^1, \theta_1) + \int u_1^2(\mathbf{x}_2(H^2, \theta_{-1}), \theta) f_{-1}(\theta_{-1}) d\theta_{-1} > u_1^1(x_1^1, \theta_1) + \int u_1^2(\mathbf{x}_2(H^2, \theta_{-1}), \theta) f_{-1}(\theta_{-1}) d\theta_{-1}$$

The only way to avoid such a profitable deviation is if there does not exist $\hat{x}_1^1 > x_1^1$, i.e. $x_1^1 = \bar{x}_1^1$. \square

Proof of Proposition 7.3: By Lemma 7.1, there can only be pooling at \bar{x}_1^1 . By Proposition 7.1, there exists a monotone PBE, which implies that in such an equilibrium, if there is pooling, type $\bar{\theta}_1$ must be pooling as well at \bar{x}_1^1 . Since we know that $\bar{\theta}_1$ will not choose \bar{x}_1^1 , there cannot be a monotone PBE with pooling at \bar{x}_1^1 . Hence there cannot be any pooling anywhere in this equilibrium, and so it is completely separating. \square

Proof of Proposition 7.5: Cases (ii) and (iii) follow immediately from Proposition 7.4.

For Case (i), though, we cannot immediately apply any single crossing conditions from previous results since types here are affiliated. So, in order to show that the best response of each player i in each period t is increasing in the strong set order, we must then show that there is single crossing of the conditional expected utility in θ . Continuity of ψ_t then follows from Lemma 5.2.

Suppose that $x_t^i(H^t, \theta_i) = 1$ is optimal for θ_i . Then suppose that type $\tilde{\theta}_i^t(H^t, \hat{\alpha}_i) = \hat{\theta}_i > \theta_i$ follows the same strategy as θ_i for all $\tau \geq t$; we indicate this strategy by $\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i)$ where $\tilde{\theta}_i^t(H^t, \alpha_i) = \theta_i$. If, under this strategy, type θ_i beats type vector θ_{-i} at some H^t , the payoff for $\hat{\theta}_i$ for beating θ_{-i} is strictly higher; moreover, $v(\hat{\theta}_i, \theta_{-i})$ is weakly increasing in θ_{-i} . Since types are affiliated, this implies that

$$\begin{aligned} & \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i), \hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_{-i})\}, \hat{\theta}_i, \tilde{\theta}_{-i}(H^t, \alpha_{-i}) g_t(\alpha_{-i} | H^t, \hat{\alpha}_i) d\alpha_{-i} \\ & \geq \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i), \hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_{-i})\}, \hat{\theta}_i, \tilde{\theta}_{-i}(H^t, \alpha_{-i}) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \\ & \geq \int u_i(H^t, \{\tilde{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_i), \hat{\mathbf{x}}_{\tau,t}^i(H^\tau, \alpha_{-i})\}, \theta_i, \tilde{\theta}_{-i}(H^t, \alpha_{-i}) g_t(\alpha_{-i} | H^t, \alpha_i) d\alpha_{-i} \end{aligned}$$

The case of symmetric players follows from Theorem 4.4. \square

Proof of Proposition 7.6: For the case of the English auction with affiliated types, it must be that all players exit at $t \leq \max_i v_i(\bar{\theta})$, as winning at a later time gives a negative payoff; hence the game reduces to that of $T < \infty$. However, for the other cases, we cannot simply invoke Theorem 4.5 because the games are not continuous at infinity; for instance, in the war of attrition, if one exits at $T = \infty$, then one receives a payoff of $-\infty$; however, if one stops at any finite t , then one's payoff is finite.

To get around this issue, we note that equilibrium payoffs for each player at any history H^t are bounded between $-c_i(t_i)$ and $v_i(\bar{\theta})$, and (weakly) increasing in θ_i . Hence we can find a convergent subsequence of truncated games $\{\Gamma_m\}_{m=1}^\infty$ indexed by final times for any on-path history $\{T_m\}$ as in Theorem 4.5, and then take further subsequences (if necessary) in which not only the strategy function $x_t^i(H^t, \cdot)$ converge pointwise, but the interim payoff functions $U_{i,m}^t(H^t, 1, \cdot)$ and $U_{i,m}^t(H^t, 0, \cdot)$ converge by Helly's selection theorem.

Now consider the limit strategy functions $x_t^i(H^t, \theta_i)$, and consider a one-stage deviation at history H^t . By continuity of beliefs in strategies from Theorem 4.1, and convergence of strategies to $\tilde{\mathbf{x}}_{t,t}^i(H^t, \alpha_{-i})$ it must be that the distribution of outcomes in period t conditional on player i staying in converges in the weak-* topology, so that payoffs converge as well.

That is, if player i stays in, then by the one-stage deviation principle,

$$U_i^t(H^t, 1, \theta_i) = \int U_i^{t+1}(\{H^t, 1, \tilde{\mathbf{x}}_{t,t}^{-i}(H^t, \alpha_{-i})\}, x_{t+1}^i(\theta_i), \theta_i) g_t(\alpha_{-i}|H^t) d\alpha_{-i} d\mu(\Theta_{-i}, H^t|H^t)$$

Since the payoffs $U_{i,m}^{t+1}$ converge to U_i^{t+1} and beliefs converge as well (making the likelihood of any given history H^{t+1} converge), it must be that if such a deviation to $x_t^i = 1$ were strictly profitable, then it would be profitable in game Γ_m for sufficiently high m as well. But since the strategy functions converge pointwise almost-everywhere, there exists M such that for any $m > M$ and almost all θ_i , $\mathbf{x}_{t,m}^i(H^t, \theta_i) = 0$; hence this would contradict that this is an equilibrium strategy. A similar argument establishes that it cannot be that type θ_i would want to deviate to $x_t^i = 0$ at H^t . \square

Proof of Proposition 7.7: That a limit of a subsequence of $\{t_{i,m}(\cdot)\}_{m=1}^\infty$ exists follows immediately from Helly's selection theorem, as the stopping time must be increasing in type. We must now check whether this function is an equilibrium of the continuous-time game. The only possible issue that can arise occurs when a positive measure of types θ_i stop at the same time, forming an accumulation point. Suppose that, given that player $-i$ plays $t_{-i}(\theta_{-i})$, it is strictly better for some type θ_i to stop at $t - \delta$ instead of $t_i(\theta_i) = t$. If, indeed, there is no positive measure of types stopping at any given time in the limit, it must be that the distribution of types θ_{-i} is the limit in the weak-* topology of the distribution of types θ_{-i} such that $t_{-i,m}(\theta_{-i}) \in [t - \delta, t]$. So, $t > t_{-i}(\theta_{-i})$ if and only if there exists M such that for all $m > M$, $t_{i,m}(\theta_i) > t_{-i,m}(\theta_{-i})$. This implies that the payoffs for choosing any given t for any given θ_i converge. Because no accumulation points exist, payoffs are continuous in the choice of t in the limit. Let $\tilde{U}_{i,m}(\tau, \theta_i)$ be the expected payoff from choosing some time τ to stop in Γ_m ; we correspondingly define $\tilde{U}_i(\tau, \theta_i)$ for Γ . If $\tilde{U}_i(\tau - \delta, \theta_i) - \pi_i(\tilde{U}, \theta_i) > 0$, then for any $\varepsilon > 0$, there exists M such that for $m > M$,

$$\begin{aligned} \tilde{U}_i(t - \delta, \theta_i) - \tilde{U}_i(t, \theta_i) &\leq \tilde{U}_{i,m}(t - \delta, \theta_i) - \tilde{U}_{i,m}(t_{i,m}(\theta_i), \theta_i) - \|\tilde{U}_i(t, \theta_i) - \tilde{U}_{i,m}(t, \theta_i)\| \\ &\quad - \|\tilde{U}_i(t - \delta, \theta_i) - \tilde{U}_{i,m}(t - \delta, \theta_i)\| - \|\tilde{U}_{i,m}(t_{i,m}(\theta_i), \theta_i) - \tilde{U}_i(t, \theta_i)\| \\ &< \tilde{U}_{i,m}(t - \delta, \theta_i) - \tilde{U}_{i,m}(t_{i,m}(\theta_i), \theta_i) + \varepsilon \end{aligned}$$

Since this holds for any $\varepsilon > 0$, this would imply that $t_{i,m}$ is not optimal for sufficiently large m , a contradiction.

We now show that accumulation points cannot occur in the limit except at $t = 0$. Suppose

that, without loss of generality, t_1^* is an accumulation point for player 1. We now divide this situation among several possibilities.

It cannot be that both players have an accumulation point at t_1^* . If so, then for any game Γ_m , one player i such that $t_i(\theta_i) = t_1^*$ loses with positive probability to player $-i$ such that $t_{-i}(\theta_{-i}) = t_1^*$. But then for any $\delta > 0$, there exists M such that for $m > M$, type θ_i can beat $t_{-i,m}(\theta_{-i})$ with strictly higher probability by dropping out at $t_1^* + \delta$, thereby discretely increasing her payoff for an arbitrarily small increase in costs from staying in (since $v_i(\theta)$ is strictly increasing), contradicting the optimality of $t_{i,m}(\theta_i)$.

Now suppose that $t_1^* > 0$. For some $\delta > 0$, there must exist interval $(t_1^* - \delta, t_1^*]$ such that there is no θ_2 such that $t_2(\theta_2) \in (t_1^* - \delta, t_1^*)$. Suppose to the contrary that it is optimal for some θ_2 to stop at such a $t' \in (t_1^* - \delta, t_1^*)$. Then in the limit as $\delta \rightarrow 0$, there would be a strictly profitable deviation to some $t_1^* + \eta$ for θ_2 in the limit game. But that would also imply that such a deviation be strictly profitable for sufficiently high m as well.

However, if such an interval $(t_1^* - \delta, t_1^*]$ exists, $t_1(\theta_1) = t_1^*$ cannot be optimal for any θ_1 . For any $\varepsilon > 0$, there would then exist M such that if $m > M$, type θ_1 's prescribed equilibrium strategy is $t_{1,m}(\theta_1) \geq t_1^* - \frac{\delta}{3}$; however, she could decrease her action to $t_1^* - \frac{2\delta}{3}$, which would decrease the probability of winning by no more than ε , but decrease continuation costs discretely (recall that ε is arbitrary while δ is fixed). Hence it would not be optimal to choose $t_{1,m}(\theta_1)$.

To summarize, we have found that any limit of equilibrium strategies can only have an accumulation point at $t = 0$, and only one player can exit then with positive probability; moreover, if there are any gaps in the exits over the interval $(t, t + \delta)$ of either player before the end of the game, then it is preferable for some player such that $t_i(\theta_i) \geq t + \delta$ to decrease her exit time to $t' \in (t, t + \delta)$ in game Γ_m for all sufficiently high m , contradicting the possibility of such a gap. \square

Appendix B: Problems with a Backward Induction Approach

We present an example to illustrate possible issues with an attempt to prove existence of equilibrium via backward induction. Throughout, we use the notation that was developed in Section 4.

Suppose that there are two periods. We focus on one particular player i with $\theta_i \sim U[0, 1]$ to show that the set of best replies will not be join-closed. i 's action set in period 1 is $X_1^i = \{1, 2, 3\}$, while in period 2 it is $X_2^i = \{1, 2\}$. In period 2, the optimal choice of x_2^i (regardless of what happens in period 1) is 2 if $\theta_i \in [\frac{1}{2}, 1]$, and 1 otherwise, with $\theta_i = \frac{1}{2}$ indifferent. Player i 's period 1 action is observable to all other players.

There are two possible approaches to using backward induction: either by normalizing the set of types in subgames (as in our translation into α_i in the construction that we have used), or by using the original set of types. In the latter case, there will be a discontinuity of payoffs for other player j in the case where $\theta_i > \frac{1}{2}$ choose $x_1^i = 3$ and $\theta_i < \frac{1}{2}$ choose $x_1^i = 1$, upon j 's observing $x_1^i = 2$. If we perturb the strategy of i so that some $\theta_i > \frac{1}{2}$ chooses $x_1^i = 2$, then j will have to believe that player i chooses $x_2^i = 2$ with probability 1 in period 2. On the other hand, if we perturb the strategy so that $\theta_i < \frac{1}{2}$ chooses $x_1^i = 2$, then j must believe that $x_2^i = 1$ in period 2. This violates continuity of payoffs for j in the strategy chosen by i , a necessary condition for Reny's theorem. Hence any attempt to make the equilibria in the subgame well-behaved will need to use a normalization of types to preserve continuity.

However, when one normalizes the set of types and then takes as given the best-replies in period 2, the set of best-replies will no longer be join-closed. Suppose that all types of θ_i are indifferent between all actions in X_1^i , and $\theta_i > \frac{1}{2}$ chooses $x_1^i = 3$. Then in period 2, all types α_i choose $x_2^i = 2$ after x_1^i being chosen in period 1. However, if all types θ_i choose $x_1^i = 3$, then in period 2, half of the types α_i choose $x_2^i = 1$, while the other half chooses $x_2^i = 2$. However, the join of these two potential strategies (in which all θ_i choose $x_1^i = 3$, and all α_i choose $x_2^i = 2$) is not a best reply: this would entail $\theta_i < \frac{1}{2}$ choosing $x_2^i = 2$, which is not optimal for those types. Thus the backward induction attempt fails to generate well-behaved best-replies in the subgame, and so will not allow for easy generation of equilibrium. \square