

Protocol Invariance and the Timing of Decisions in Dynamic Games*

Ulrich Doraszelski[†] Juan F. Escobar[‡]

September 3, 2018

Abstract

We characterize a class of dynamic stochastic games that we call separable dynamic games with noisy transitions and establish that these widely used models are protocol invariant provided that periods are sufficiently short. Protocol invariance means that the set of Markov perfect equilibria is nearly the same irrespective of the order in which players are assumed to move within a period. Protocol invariance can facilitate applied work and renders the implications and predictions of a model more robust. Our class of dynamic stochastic games includes investment games, R&D races, models of industry dynamics, dynamic public contribution games, asynchronously repeated games, and many other models from the extant literature.

*We thank Yuichiro Kamada, George Mailath, Juuso Toikka, the editor Simon Board, and two anonymous referees for very useful comments and suggestions. Doraszelski acknowledges financial support from the Dean's Research Fund at the Wharton School. Escobar acknowledges financial support from the Institute for Research in Market Imperfections and Public Policy, MIPP, ICM IS130002 and from Conicyt (Fondecyt No. 1180723).

[†]Wharton School, University of Pennsylvania, Steinberg Hall-Dietrich Hall, 3620 Locust Walk, Philadelphia, PA 19104, USA. E-mail: doraszelski@wharton.upenn.edu.

[‡]Center for Applied Economics, Department of Industrial Engineering, University of Chile, Beauchef 851, Santiago, Chile. E-mail: jescobar@dii.uchile.cl.

1 Introduction

The timing of decisions is an essential ingredient into modelling many strategic situations. Asynchronous decisions are a type of commitment, and being able to move first and thus set the stage for rivals can confer a considerable advantage on a player. Synchronous decisions, in contrast, take away the ability to commit as players are neither leaders nor followers. From the basically static models in Cournot (1838) and von Stackelberg (1934) to the genuinely dynamic models in Cyert and DeGroot (1970) and Maskin and Tirole (1987, 1988a, 1988b) and the anti-folk theorems in Rubinstein and Wolinsky (1995) and Lagunoff and Matsui (1997, 2001), a long and distinguished literature has pointed out cases where the protocol of moves matters crucially for equilibrium behavior.

Our paper provides a counterpoint to this literature. We show that a fairly general and widely used class of dynamic models is protocol invariant provided that periods are sufficiently short and moves are therefore sufficiently frequent. Protocol invariance means that the set of equilibria of a model is nearly the same irrespective of the order in which players are assumed to move within a period, including—and extending beyond—simultaneous, alternating, and sequential moves. Protocol invariance can facilitate applied work and renders the implications and predictions of a model more robust.

We focus on infinite-horizon dynamic stochastic games with finite states and actions and their stationary Markov perfect equilibria (henceforth Markov perfect equilibria for short). A dynamic stochastic game is a dynamic system that can be in different states in different periods according to a discrete-time Markov process that players can influence through their actions. Dating back to Shapley (1953), dynamic stochastic games have a long tradition in economics and are central to the analysis of strategic interactions among forward-looking players in dynamic environments. The main contribution of this paper is to characterize a class of dynamic stochastic games that are protocol invariant provided that periods are sufficiently short. We call this class *separable dynamic games with noisy transitions*.

We apply dynamic stochastic games to situations in which a player primarily influences his rivals' payoffs by taking action to change the state. While per-period payoffs and state-to-state transitions depend arbitrarily on the state in a separable dynamic game, they are assumed to depend on players' actions in an additive manner: to a first-order approximation, per-period payoffs and state-to-state transitions are built from parts that depend on the actions taken by individual players. To the extent that there are complementarities between players' actions and other non-separabilities in

per-period payoffs and state-to-state transitions, they must vanish as periods become short.

Noisy transitions preclude that there is an action that a player can take to guarantee a change in the state. We model the evolution of the state by a discrete-time approximation to a continuous-time Markov process in which the time spent in a state has an exponential distribution with a finite hazard rate. The finite hazard rate implies that transitions are noisy. This assumption reflects the view that models are only an approximation to reality, and so there always is some residual uncertainty associated with taking an action.

While the assumptions of separability and noisy transitions are restrictive and doubtlessly rule out some interesting applications, many dynamic models in the literature are amenable to these assumptions. Examples include investment games (Spence 1979, Fudenberg and Tirole 1983, Hanig 1986, Reynolds 1987, Reynolds 1991, Dockner 1992), R&D races (Reinganum 1982, Lippman and McCardle 1987), models of industry dynamics (Ericson and Pakes 1995), dynamic public contribution games (Marx and Matthews 2000, Compte and Jehiel 2004, Georgiadis 2015), and the recent continuous-time stochastic games with moves at random times (Arcidiacono, Bayer, Blevins, and Ellickson 2016, Ambrus and Lu 2015, Calcagno, Kamada, Lovo, and Sugaya 2014, Kamada and Kandori 2017). We also show that some dynamic models that are not obviously separable, including the asynchronously repeated Bertrand, Cournot, and coordination games in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997), can be re-cast to satisfy our assumptions.

Our main result is that separable dynamic games with noisy transitions are protocol invariant provided that periods are sufficiently short. To provide intuition, consider a prototypical investment game between two firms. A firm can undertake a risky investment project to increase its capital stock. A firm's per-period payoff increases in its own capital stock and decreases in its rival's capital stock. The separability assumption is satisfied, as whether its rival invests affects directly neither the firm's per-period payoff nor the probability that the firm succeeds in increasing its capital stock. Moreover, transitions from one state to another are noisy due to the risky nature of the investment project.

Now contrast two protocols of moves. When firms move alternately, a forward-looking firm deciding whether to invest understands that its rival's capital stock remains constant for (at least) the period. In contrast, when firms move simultaneously, the firm has to take into account the probability that its rival's capital stock increases over the course of the period. This probability, however, becomes negligible as periods become short because transitions are noisy. The protocol of moves is therefore almost immaterial to the firm's decision.

As intuitive as it may be that our assumptions imply protocol invariance, this intuition neither immediately translates into a proof nor is it always salient in the literature. Consider the large literature on technology adoption and the related debate about the persistence of monopoly (see Reinganum (1989) and the references therein). In a prototypical model with two firms, the state-of-the-art technology evolves over time and a firm decides whether to continue operating its current technology or pay a cost to undertake an upgrade to the state-of-the-art technology. The separability assumption is satisfied as a firm’s per-period payoff depends on its rival’s current technology but not directly on its rival’s adoption decision. Riordan and Salant (1994) approximate a continuous-time game by a discrete-time game with alternating moves and establish that a pattern of increasing dominance arises and all adoptions are by one of the firms. Giovannetti (2001) shows in a discrete-time game with simultaneous moves that firms take turns to adopt in a pattern of perpetual leapfrogging. Most recently, Iskhakov, Rust, and Schjerning (2017) show that a game with alternating moves has a unique Markov perfect equilibrium whereas the game with simultaneous moves has a vast number of equilibria, with some displaying increasing dominance and others leapfrogging.¹ Our main result clarifies that these stark differences between simultaneous and alternating moves hinge on the outcome of the adoption decision being certain. Provided that periods are sufficiently short, these differences vanish under the empirically plausible assumption that the process of updating to the state-of-the-art technology involves at least some uncertainty.

The proof of our main result proceeds from the equilibrium conditions for a separable dynamic game with noisy transitions. We fix an arbitrary protocol of moves and take the limit as periods become short and we therefore pass from discrete to continuous time. We observe that the limit conditions are independent of the protocol of moves used to pass to the limit. Protocol invariance follows immediately if the limit conditions admit a unique solution.

However, due to the richness of the class of separable dynamic games with noisy transitions, the limit conditions may admit multiple solutions. In this case, taking the limit of a sequence of equilibria under one protocol of moves may potentially lead to a different solution of the limit conditions than taking the limit under another protocol of moves, thereby causing protocol invariance to fail. To show that this does not happen, we introduce differential topology tools to study the limit conditions. We prove that generically all solutions of the limit conditions can be approximated by the Markov perfect equilibria of a separable dynamic game with noisy transitions and an

¹In the game with alternating moves uniqueness obtains under certain parameter restrictions; absent these restrictions, the “set of payoffs shrinks dramatically” in the game with alternating moves as compared to the game with simultaneous moves (Iskhakov, Rust, and Schjerning 2017, p. 47).

arbitrary protocol of moves provided that periods are sufficiently short. Our proof draws on ideas in Harsanyi (1973a, 1973b), who establishes robustness to perturbations of payoffs and information in normal-form games. While differential topology tools have been used in economic theory before (see, e.g., Chapter 17 of Mas-Colell, Whinston, and Green 1995), to the best of our knowledge our paper is the first attempt to use them to establish robustness to the timing of decisions.

While we mostly treat the limit conditions as a technical device, they are of interest by themselves. The limit conditions can be interpreted as the equilibrium conditions for a continuous-time stochastic game. We also provide an equivalence result showing that the limit conditions are identical to the equilibrium conditions for a dynamic stochastic game in which in any period one player is randomly selected to make a decision. In this game with random moves, the fact that a player can revise his decision only at random times confers a kind of commitment power on the player similar to that in the games with alternating moves in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997). Our equivalence result therefore underscores that the class of separable dynamic games with noisy transitions admits quite rich strategic interactions between players.

Separability, noisy transitions, and Markov perfection appear to be the key properties that underpin protocol invariance. We show that our assumptions are tight in the sense that counterexamples to protocol invariance can be constructed if any one of them is relaxed. In particular, we show that protocol invariance does not extend beyond Markov perfect equilibria to other equilibrium concepts.

Our main result facilitates and informs applied work in a number of ways. First and perhaps most important, determining the protocol of moves that is most realistic and appropriate for the application at hand may be amongst the most difficult choices a modeler has to make. In empirical work, in particular, the timing of decisions and the ability to commit is typically not observable to the researcher. Hence, we may be suspicious of any implication or prediction from a model that is driven by the protocol of moves that the modeler has chosen to impose, a point that has been made forcefully by Rosenthal (1991) and van Damme and Hurkens (1996) for normal-form games and by Kalai (2004) for large Bayesian games. Protocol invariance alleviates this concern and the burden of determining the protocol of moves for the class of separable dynamic games with noisy transitions by ensuring that equilibrium behavior is independent of the timing of decisions provided that periods are sufficiently short. Second, because the timing of decisions and the ability to commit is typically not observable, empirical work sometimes averages over different protocols of moves (Einav 2010). This average depends on the assumed probability distribution

over protocols of moves and may be difficult to interpret if it does not correspond to an equilibrium of any game. Protocol invariance renders averaging unnecessary. Third, our main result cautions against the presumption that imposing asynchronous instead of synchronous decisions on a dynamic stochastic game reduces the number of equilibria. Fourth, dynamic stochastic games are often not very tractable analytically and thus call for the use of numerical methods. Doraszelski and Judd (2018) show that the computational burden can vary by orders of magnitude with the protocol of moves. For the class of separable dynamic games with noisy transitions, protocol invariance justifies imposing the protocol of moves that is most convenient from a computational perspective.

We apply and extend our main result in three ways. First, we extend protocol invariance to the limiting case of deterministic transitions and provide a novel dynamic programming characterization of separable dynamic games with noisy transitions as moves become arbitrarily frequent and hazard rates arbitrarily large. As a particular example we show that a class of coordination games has a unique equilibrium that attains the efficient outcome regardless of the protocol of moves. We emphasize that this extension of our protocol-invariance theorem relies on a particular way of taking the joint limit. We exposit a discontinuity in the set of Markov perfect equilibria as moves become arbitrarily frequent and hazard rates arbitrarily large and argue that many examples from the extant literature where equilibrium behavior hinges on the protocol of moves can be seen as a manifestation of this discontinuity.

Second, we contribute a rationale for focusing on Markov perfect equilibria to the literature that provides foundations for these equilibria in a variety of dynamic models (Maskin and Tirole 2001, Bhaskar and Vega-Redondo 2002, Sannikov and Skrzypacz 2007, Faingold and Sannikov 2011, Bhaskar, Mailath, and Morris 2013, Bohren 2014). Provided that periods are sufficiently short, we show that if a strict finite-memory equilibrium payoff profile in a separable dynamic game with noisy transitions and simultaneous moves is protocol invariant, then it is arbitrarily close to a Markov perfect equilibrium payoff profile. Thus, the Markovian restriction on equilibrium strategies is not only sufficient but also necessary for protocol invariance. Markov perfect equilibria are therefore the only equilibria that are robust to changes in the protocol of moves.

Third, we contribute to the literature on computing Markov perfect equilibria (Pakes and McGuire 1994, Pakes and McGuire 2001, Doraszelski and Judd 2012, Weintraub, Benkard, and Van Roy 2008, Ifrach and Weintraub 2017, Doraszelski and Judd 2018). Doraszelski and Judd (2012) show that the limit conditions that arise as we pass from discrete to continuous time are particularly easy to solve numerically, often reducing the computational burden by orders of magni-

tude. Our main result provides a justification for solving the limit conditions by establishing that these solutions almost coincide with the Markov perfect equilibria of separable dynamic games with noisy transitions and arbitrary protocols of moves provided that periods are sufficiently short. Moreover, we show that this one-to-one correspondence between the solutions to the limit conditions and the Markov perfect equilibria of dynamic stochastic games obtains more broadly. In particular, we can dispense with the separability assumption if we restrict attention to games with simultaneous moves.

Our paper is related to two strands of literature. First, our notion of protocol invariance builds on and extends the notion of a commitment robust equilibrium in Rosenthal (1991) and van Damme and Hurkens (1996) from two-player normal-form games to N -player dynamic stochastic games. Rosenthal (1991) defines a Nash equilibrium of a two-player normal-form game to be commitment robust if it is also a subgame perfect equilibrium outcome of each of the two extensive-form games in which one of the players moves first, and provides a series of illustrative examples.² Neither Rosenthal (1991) nor van Damme and Hurkens (1996) provide conditions for a commitment robust equilibrium to arise in a normal-form game. In contrast to the notion of a commitment robust equilibrium, our notion of protocol invariance pertains to the entire set of equilibria of a fairly general class of dynamic models.

We share the focus on the entire set of equilibria with Kalai (2004), who shows that with many semi-anonymous players, the equilibria of Bayesian games with simultaneous moves are extensively robust and survive even if the simultaneous-moves assumption is relaxed to allow for a variety of extensive modifications. In contrast to our paper, Kalai (2004) cannot exclude the possibility that an equilibrium of an extensive modification is not an equilibrium of the game with simultaneous moves (see his Remark 6). Our notion of protocol invariance is more demanding in this respect. The driving force behind Kalai's (2004) result is the vanishing impact that a player's action has on other players' payoffs as the number of players grows large. In our setting, the impact that a player's action has on other players' payoffs (other than through a change in the state) vanishes as periods become short. From a more technical perspective, Kalai (2004) allows for ϵ -equilibria, while we impose exact equilibrium and prove results for generic payoffs.

Second, previous attempts to exposit dynamic models where the protocol of moves does not matter for equilibrium behavior are few and far between and confined to very specific models.

²van Damme and Hurkens (1996) provide a slightly different notion of a commitment robust equilibrium than Rosenthal (1991) and link it to the notions of persistent equilibria (Kalai and Samet 1984) and curb* equilibria (Basu and Weibull 1991).

Abreu and Gul (2000) study bilateral bargaining and show that independent of the bargaining protocol the same limit is reached as the time between offers becomes short. Caruana and Einav (2008) study a model in which players repeatedly announce an action but only the final announced action is relevant for payoffs. While players can revise their announcements, they pay a cost each time they do so; in this way, announcements play the role of an imperfect commitment device. Caruana and Einav (2008) show that the order in which players make announcements does not matter as long as the time between announcements is sufficiently short. In contrast to Abreu and Gul (2000) and Caruana and Einav (2008), we do not presuppose that the limit conditions admit a unique solution. Because our class of dynamic stochastic games is much less tightly specified, we require differential topology tools to analyse the limit conditions.

The remainder of this paper is organized as follows. Section 2 introduces separable dynamic games with noisy transitions. Section 3 develops our main result. Section 4 discusses a number of applications and extensions of our main result and Section 5 concludes. Unless noted otherwise, proofs are in Appendix A. Appendix B provides examples and extensions of our main result.

2 Separable Dynamic Games with Noisy Transitions

We focus on dynamic stochastic games with finite sets of players, states, and actions. Time $t = 0, \Delta, 2\Delta, \dots$ is discrete and measured in units of $\Delta > 0$. We refer to Δ as the length of a period; as $\Delta \rightarrow 0$, moves become frequent. The time horizon is infinite.³ We focus on dynamic stochastic games with finite sets of players, states, and actions. Let $\{1, 2, \dots, N\}$ denote the set of players, Ω the set of states, and $A_i(\omega)$ the set of actions of player i in state ω . Each player strives to maximize the expected net present value of his stream of payoffs and discounts future payoffs using a discount rate $\rho > 0$. Monitoring is perfect.

The protocol of moves determines which players can take an action at time t and which players cannot. We allow for a general protocol of moves that encompasses—and goes beyond—simultaneous, alternating, and sequential moves. To this end, we allow the set of players who have the move to change from one period to the next. The set of players $J^t \subseteq \{1, 2, \dots, N\}$ who have the move at time t thus becomes part of the state of the system, and we refer to it as the

³ A finite set of states rules out finite-horizon games in which time is a state variable that, as moves become frequent, belongs to an arbitrarily large set. Our main result extends to games with infinite state spaces if the limit conditions admit a unique solution as in Caruana and Einav (2008) and Abreu and Gul (2000). In Section 4.3, we discuss an extension to a continuum of actions.

“protocol” state to distinguish it from the familiar “physical” state $\omega^t \in \Omega$. In contrast to the physical state, we assume that the protocol state evolves independently of players’ actions.

The game proceeds as follows. It starts at time $t = 0$ from an initial state $(\omega^{t=0}, J^{t=0})$. After observing $(\omega^{t=0}, J^{t=0})$, the players $j \in J^{t=0}$ who have the move choose their actions $a_{j^{t=0}}^{t=0} = (a_j^{t=0})_{j \in J^{t=0}}$ simultaneously and independently from each other. Now two things happen, depending on the state $(\omega^{t=0}, J^{t=0})$ and the actions $a_{j^{t=0}}^{t=0}$. First, player i receives a payoff $u_i^\Delta(\omega^{t=0}, J^{t=0}, a_{j^{t=0}}^{t=0})$. Second, the system transits from state $(\omega^{t=0}, J^{t=0})$ to state $(\omega^{t=\Delta}, J^{t=\Delta})$. Independent of each another, the transition from $\omega^{t=0}$ to $\omega^{t=\Delta}$ happens with probability $\Pr^\Delta(\omega^{t=\Delta} | \omega^{t=0}, J^{t=0}, a_{j^{t=0}}^{t=0})$ and that from $J^{t=0}$ to $J^{t=\Delta}$ with probability $\Pr(J^{t=\Delta} | J^{t=0})$. While player i receives a payoff irrespective of whether he has the move ($i \in J^{t=0}$) or not ($i \notin J^{t=0}$), the exact amount depends on the actions $a_{j^{t=0}}^{t=0}$ of the players who have the move, as do the state-to-state transitions. In the next round at time $t = \Delta$, after observing $(\omega^{t=\Delta}, J^{t=\Delta})$, the players $j \in J^{t=\Delta}$ who have the move choose their actions $a_{j^{t=\Delta}}^{t=\Delta}$. Then player i receives a payoff $u_i^\Delta(\omega^{t=\Delta}, J^{t=\Delta}, a_{j^{t=\Delta}}^{t=\Delta})$ and the state changes again from $(\omega^{t=\Delta}, J^{t=\Delta})$ to $(\omega^{t=2\Delta}, J^{t=2\Delta})$. The game goes on in this way ad infinitum.

To allow for a general protocol of moves, we partition the set of players $\{1, 2, \dots, N\}$ and assume that the set of players J^t who have the move at time t evolves according to a Markov process that is defined over this partition as follows:⁴

Assumption 1 (Protocol of Moves) *Let \mathcal{J} be a partition of $\{1, 2, \dots, N\}$ and $\mathcal{P} = (\Pr(J' | J))_{J, J' \in \mathcal{J}}$ a $|\mathcal{J}| \times |\mathcal{J}|$ transition matrix. \mathcal{P} is irreducible and its unique stationary distribution is uniform on \mathcal{J} .*

In stating Assumption 1 and throughout the remainder of the paper we omit the time superscript whenever possible and use a prime to distinguish future from current values.

We denote the protocol of moves as $\langle \mathcal{J}, \mathcal{P} \rangle$ in what follows. Because \mathcal{J} is a partition of $\{1, 2, \dots, N\}$, Assumption 1 ensures that player i always has the move in conjunction with the same rivals. By requiring the transition matrix \mathcal{P} to have a unique stationary distribution that is uniform on \mathcal{J} , Assumption 1 further ensures that all players have the move with the same frequency over a sufficiently large number of periods.

Assumption 1 accommodates synchronous and asynchronous decisions and thus encompasses most dynamic stochastic games in the literature, including games with simultaneous moves (Shapley 1953, Ericson and Pakes 1995), games with alternating moves (Maskin and Tirole 1987, Maskin and

⁴Our analysis easily extends to allow for $J^t = \emptyset$.

Tirole 1988b, Maskin and Tirole 1988a, Lagunoff and Matsui 1997), and games with random moves (Iskhakov, Rust, and Schjerning 2017, Doraszelski and Judd 2018). In games with simultaneous moves, the partition is $\mathcal{J} = \{\{1, \dots, N\}\}$ with the trivial 1×1 transition matrix \mathcal{P} ; in games with alternating moves the partition is $\mathcal{J} = \{\{1\}, \dots, \{N\}\}$ with the $N \times N$ transition matrix \mathcal{P} with entries $\Pr(\{ \bmod_N(i+1) \} | \{i\}) = 1$.⁵ In games with asynchronous moves, $\mathcal{J} = \{\{1\}, \dots, \{N\}\}$ and the identity of the player who has the move in a given period may follow a deterministic sequence as in games with alternating moves or it may be stochastic. Games with random moves are another special case of games with asynchronous moves. In these games, the probability that a player has the move in a given period is uniform across players and periods. Finally, Assumption 1 accommodates more than one—but less than all—players having the move in a given period and thus settings where decisions are partially synchronous.

Assumption 1 can be relaxed in several ways. In Appendix B.1, we show that our protocol-invariance theorem remains valid if the evolution of the protocol state J depends on the physical state ω and that the uniform stationary distribution in Assumption 1 can be replaced by a non-uniform stationary distribution. In Appendix B.2, we provide a partial extension of our protocol-invariance theorem that does not require \mathcal{J} to be a partition of the set of players. This introduces complexities that are not easily dealt with using our differential topology tools and forces us to restrict attention to pure strategies.

We model the evolution of the physical state by a discrete-time approximation to a continuous-time Markov process in order to impose that transitions are noisy:⁶

Assumption 2 (Noisy Transitions) *The transition probability $\Pr^\Delta(\omega' | \omega, J, a_J)$ is differentiable in Δ and can be written as*

$$\Pr^\Delta(\omega' | \omega, J, a_J) = \begin{cases} 1 - q_J(\omega, a_J)\Delta + \mathcal{O}(\Delta^2) & \text{if } \omega' = \omega, \\ q_J(\omega, a_J)p_J(\omega' | \omega, a_J)\Delta + \mathcal{O}(\Delta^2) & \text{if } \omega' \neq \omega, \end{cases}$$

where $q_J: \{(\omega, (a_j)_{j \in J}) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{R}^+ \cup \{0\}$, $p_J: \{(\omega, (a_j)_{j \in J}) \mid a_j \in A_j(\omega)\} \rightarrow \mathbf{P}(\Omega)$, and $\mathbf{P}(\Omega)$ is the set of probability distributions over Ω . We normalize $p_J(\omega | \omega, a_J) = 0$.

Without loss of generality, we decompose the transition probability $\Pr^\Delta(\omega' | \omega, J, a_J)$ into a probability that the state changes in a given period—or that a jump occurs in the lingo of stochastic

⁵The notation $\bmod_N(x)$ refers to the modulo N congruence.

⁶An alternative modelling approach is to start from a discrete-time Markov process and then “rescale” it to shorter periods. This presumes that the discrete-time Markov process is embeddable (Elfving 1937).

processes—and a probability distribution over successor states conditional on the state changing. The probability that the state changes is $q_J(\omega, a_J)\Delta$ in proportion to the length of a period Δ and, conditional on the state changing, the probability that it changes from ω to ω' is $p_J(\omega'|\omega, a_J)$. Normalizing $p_J(\omega|\omega, a_J) = 0$ amounts to ignoring a jump from a state to itself and adjusting the hazard rate $q_J(\omega, a_J)$ of a jump occurring accordingly. Importantly, Assumption 2 restricts the model to have a finite hazard rate in a given state so that there is no action that a player can take to guarantee a change in the state.

We finally assume that per-period payoffs and state-to-state transitions have an additively separable structure:

Assumption 3 (Separability) *The per-period payoff $u_i^\Delta(\omega, J, a_J)$ is differentiable in Δ and can be written as*

$$u_i^\Delta(\omega, J, a_J) = |\mathcal{J}| \sum_{j \in J} u_{i,j}(\omega, a_j) \Delta + \mathcal{O}(\Delta^2),$$

where $u_{i,j}: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{R}$. The hazard rate $q_J(\omega, a_J)$ and transition probability $p_J(\omega' \mid \omega, a_J)$ can be written as

$$q_J(\omega, a_J) = |\mathcal{J}| \sum_{j \in J} q_j(\omega, a_j)$$

and

$$q_J(\omega, a_J) p_J(\omega' \mid \omega, a_J) = |\mathcal{J}| \sum_{j \in J} q_j(\omega, a_j) p_j(\omega' \mid \omega, a_j),$$

where $q_j: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{R}^+ \cup \{0\}$ and $p_j: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{P}(\Omega)$.

To a first-order approximation, Assumption 3 builds up the per-period payoff $u_i^\Delta(\omega, J, a_J)$ of player i from the flow payoff $u_{i,j}(\omega, a_j)$ by summing over the players $j \in J$ who have the move. By taking action a_j in state ω , player j “contributes” $|\mathcal{J}|u_{i,j}(\omega, a_j)\Delta$ to the per-period payoff of player i in proportion to the length of a period Δ . This restricts complementarities between players’ actions and other non-separabilities to the higher-order term $\mathcal{O}(\Delta^2)$.⁷ We discuss below our reason for scaling by the number of elements of the partition \mathcal{J} .

While at first glance Assumption 3 may seem to trivialize the strategic interactions between players, it does not. Importantly, Assumption 3 does not restrict how per-period payoffs and state-

⁷More explicitly, we assume that there exists $\bar{c} > 0$ and $\bar{\Delta} > 0$ such that $\|u_i^\Delta(\omega, J, a_J) - |\mathcal{J}| \sum_{j \in J} u_{i,j}(\omega, a_j) \Delta\| \leq \bar{c} \Delta^2$ for all $\Delta < \bar{\Delta}$.

to-state transitions depend on the state. Because players are forward looking, this allows strategic interactions to be channeled through continuation values. As mentioned in Section 1 and further discussed in Section 2.1, examples of dynamic stochastic games that satisfy Assumption 3 range from investment games over the recent continuous-time stochastic games with moves at random times to asynchronously repeated games.

A basic requirement for comparing equilibrium behavior under different protocols of moves is that the game remains well-defined as the protocol of moves changes. Assumption 3 ensures that the given per-period payoffs and state-to-state transitions suffice to fully specify the game for any protocol of moves that Assumption 1 admits. We return to this point in Example 6.

Assumption 3 is rather mild on its own. Indeed, it can be satisfied in many dynamic stochastic games by simply redefining the state space.⁸ However, Assumption 3 has bite in conjunction with Assumption 2. In particular, Assumption 2 precludes the state from instantaneously recording players' actions, thereby limiting attempts to redefine the state space.

In conjunction with Assumption 2, Assumption 3 builds up the components of the transition probability $\Pr^\Delta(\omega'|\omega, J, a_J)$ from the player-specific hazard rate $q_j(\omega, a_j)$ and the transition probability $p_j(\omega'|\omega, a_j)$ by summing over the players $j \in J$ who have the move. Because it imposes a competing hazards model on the transition probability, a change in the state is with high probability due to the action taken by one of the players having the move.

In what follows, we call the above game a *separable dynamic game with noisy transitions* and denote it by $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$. We view the function $u_{i,j}: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{R}$ as a vector $u_{i,j} \in \mathbb{R}^{\sum_{\omega \in \Omega} |A_j(\omega)|}$ and denote $u_i = (u_{i,j})_{j=1}^N \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$ and $u = (u_i)_{i=1}^N \in \mathbb{R}^{N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$. We further denote the collection of hazard rates and transition probabilities $q = (q_j(\omega, a_j))_{\omega \in \Omega, j=1, \dots, N, a_j \in A_j(\omega)}$ and $p = (p_j(\omega'|\omega, a_j))_{\omega \in \Omega, j=1, \dots, N, a_j \in A_j(\omega)}$.

A stationary Markovian strategy for player i is a function $\sigma_i: \Omega \rightarrow \cup_{\omega \in \Omega} \mathcal{P}(A_i(\omega))$ with $\sigma_i(\omega) \in \mathcal{P}(A_i(\omega))$ for all ω , where $\mathcal{P}(A_i(\omega))$ is the set of probability distributions over $A_i(\omega)$. Because \mathcal{J} is a partition of $\{1, 2, \dots, N\}$, Assumption 1 ensures that player i always has the move in conjunction with the same rivals. Hence, while the state of the system comprises both the physical state ω and the protocol state J , it suffices to consider Ω as the domain of σ_i . We use $\sigma_i(a_i \mid \omega)$ to denote the probability that action $a_i \in A_i(\omega)$ is played in state ω . We explore more general strategies in

⁸Consider the canonical dynamic stochastic game in Shapley (1953). If the state space takes the form $\Omega = \prod_{i=1}^N \Omega_i$ and player i independently controls the evolution of $\omega_i \in \Omega_i$, then the game can be considered separable even if the per-period payoff of player i is $v_i(\omega, a)\Delta$. To see this, redefine the state in period t as $\tilde{\omega}^t = (\omega^{t-1}, a^{t-1}, \omega^t)$ and let the flow payoff of player i be $u_{i,i}(\tilde{\omega}^t, a_i^t) = \frac{1}{\delta} v_i(\omega^{t-1}, a^{t-1})$ and $u_{i,j}(\tilde{\omega}^t, a_j^t) = 0$ if $j \neq i$, where δ is the discount factor.

Example 8 and Section 4.2.

From hereon, we denote by Σ_i the set of stationary Markovian strategies for player i and $\Sigma = \prod_{i=1}^N \Sigma_i$ the set of strategy profiles. To account for mixed strategies, we extend the flow payoff $u_{i,j}(\omega, \sigma_j(\omega)) = \sum_{a_j \in A_j(\omega)} u_{i,j}(\omega, a_j) \sigma_j(a_j | \omega)$ and transition probability

$$\Pr^\Delta(\omega' | \omega, J, \sigma_J(\omega)) = \sum_{a_J \in \prod_{j \in J} A_j(\omega)} \left(\Pr^\Delta(\omega' | \omega, J, a_J) \prod_{j \in J} \sigma_j(a_j | \omega) \right).$$

A profile of stationary Markovian strategies $\sigma = (\sigma_i)_{i=1}^N$ is a stationary Markov perfect equilibrium if it is a subgame perfect equilibrium of the game Γ . The set of Markov perfect equilibria of the game Γ is denoted $\text{Equil}(\Gamma)$. This set is nonempty (Shapley 1953).⁹

Our main interest is to compare equilibrium behavior under different protocols of moves. Assumption 1 lets us compare two models $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$ and $\bar{\Gamma} = \langle \Delta, \bar{\mathcal{J}}, \bar{\mathcal{P}}, u, p, q, \rho \rangle$ that differ only in the protocol of moves by ensuring that all players move with the same frequency. The scale factor $|\mathcal{J}|$ in Assumption 3 further ensures that a player's action brings about the same payoffs and chances of changing the state in the two models. To see this, contrast a game with simultaneous moves Γ with a game with alternating moves $\bar{\Gamma}$. In the game with simultaneous moves Γ , player j takes an action a_j every Δ units of time, yielding the payoff $u_{i,j}(\omega, a_j)\Delta$ and the hazard rate $q_j(\omega, a_j)\Delta$ (neglecting the higher-order term $\mathcal{O}(\Delta^2)$). Over a stretch of $N\Delta$ units of time, the action a_j thus yields the payoff $u_{i,j}(\omega, a_j)N\Delta$ and the hazard rate $q_j(\omega, a_j)N\Delta$. In the game with alternating moves $\bar{\Gamma}$, in contrast, player j has the move only once every $N\Delta$ units of time. According to Assumption 3, if player j takes an action a_j , then this yields the payoff $|\bar{\mathcal{J}}|u_{i,j}(\omega, a_j)\Delta = Nu_{i,j}(\omega, a_j)\Delta$ and the hazard rate $|\bar{\mathcal{J}}|q_j(\omega, a_j)\Delta = Nq_j(\omega, a_j)\Delta$. Hence, per-period payoffs and state-to-state transitions in the game with alternating moves $\bar{\Gamma}$ are comparable to those in the game with simultaneous moves Γ .¹⁰

⁹Shapley (1953) establishes existence for dynamic stochastic games with simultaneous moves. To apply his result, we view the game Γ as a dynamic stochastic game with simultaneous moves in which the players that do not have the move have no impact on per-period payoffs and state-to-state transitions.

¹⁰Instead of scaling by the number of elements of the partition \mathcal{J} in Assumption 3, we can assume that interactions occur at time $t = 0, \Delta/|\mathcal{J}|, 2\Delta/|\mathcal{J}|, \dots$. This alternative formulation ensures that a player has the move on average once every Δ units of time. Our results immediately carry over.

2.1 Examples

In the remainder of this section we discuss how prominent examples of dynamic stochastic games from the literature can be cast as special cases of our model. A reader interested in the statement of our protocol-invariance theorem and its proof may skip this section. Note that Examples 1 and 5 are used later in the paper to illustrate some results and extensions.

Example 1 (Entry Games and R&D Races) Consider $N = 2$ firms that may enter a new market. To enter the market, firm i must complete K steps. For example, to build a cement plant and enter the market, a firm needs to find a location, design the plant, obtain environmental permits, negotiate with contractors, etc. Alternatively, consider an R&D race in which a firm gradually discovers an invention and obtains a patent through a series of intermediate steps (Fudenberg, Gilbert, Stiglitz, and Tirole 1983, Grossman and Shapiro 1987, Harris and Vickers 1987).

Let $K \geq 1$ be the number of required steps and $\omega_i \in \Omega_i = \{0, 1, \dots, K\}$ the number of steps that firm i has already completed. The state of the game is $\omega = (\omega_1, \omega_2) \in \Omega_1 \times \Omega_2 = \Omega$. To take the next step, firm i can make an investment, denoted by $a_i = 1$, at cost $c_i > 0$. Action $a_i \in A_i(\omega) = \{0, 1\}$ induces the hazard rate

$$q_i(\omega, a_i) = \begin{cases} \lambda a_i & \text{if } \omega_i \leq K - 1, \\ 0 & \text{if } \omega_i = K. \end{cases}$$

This means that a firm can change its state only if it makes an investment ($a_i = 1$). The transition probability is¹¹

$$p_i(\omega' | \omega, a_i) = \begin{cases} 1 & \text{if } \omega'_i = \omega_i + 1, \omega'_{-i} = \omega_{-i}, \omega_i \leq K - 1, \\ 1 & \text{if } \omega'_i = 0, \omega'_{-i} = \omega_{-i}, \omega_i = K, \\ 0 & \text{otherwise.} \end{cases}$$

Hence, firm i progresses to the next step ($\omega'_i = \omega_i + 1$) unless it has already completed all steps ($\omega_i = K$).

Once firm i has completed all steps it enters the new market (or obtains the patent) and, depending on whether its rival has also completed all steps, obtains the monopoly profit $B_i > 0$ or

¹¹Recall that we normalize $p_J(\omega | \omega, a_J) = 0$. Conditional on a jump occurring, we therefore specify a transition from $\omega_i = K$ to $\omega'_i = 0$ with probability one. This is immaterial, however, because no jump occurs as $q_i(\omega, a_i) = 0$ if $\omega_i = K$.

the duopoly profit $b_i < B_i$. Its flow payoff is

$$u_{i,i}(\omega, a_i) = \begin{cases} B_i - c_i a_i & \text{if } \omega_i = K, \omega_{-i} \leq K - 1, \\ b_i - c_i a_i & \text{if } \omega_i = \omega_{-i} = K, \\ -c_i a_i & \text{otherwise} \end{cases}$$

and $u_{i,j}(\omega, a_j) = 0$ if $j \neq i$. Assumptions 2 and 3 are satisfied.

The literature has recognized the multiplicity of equilibria in standard entry games as a challenge for empirical research (Bresnahan and Reiss 1990, Berry 1992) and has in some cases resorted to asynchronous decisions to obtain uniqueness (Mazzeo 2002, Einav 2010, Cleeren, Verboven, Dekimpe, and Gielens 2010). Lippman and McCardle (1987) study the continuous-time limit of our alternative entry game as $\Delta \rightarrow 0$ and show that the equilibrium is generically unique. As Section 3 shows, generic uniqueness extends to discrete-time games under a variety of protocols of moves, thereby obviating the need for an ad hoc assumption that is difficult to verify in the data.

Example 2 (Industry Dynamics) Ericson and Pakes (1995) develop a discrete-time model of industry dynamics. In their model and the large literature following it (see Doraszelski and Pakes (2007) for a survey), incumbent firms decide on investment and exit and compete in the product market; potential entrants decide on entry. Depending on the application, firm i 's state variable $\omega_i \in \Omega_i$ encodes its current product quality, production capacity, marginal cost, etc. It further encodes whether firm i is currently an incumbent firm that competes in the product market or a potential entrant. The state of the game is $\omega = (\omega_1, \omega_2, \dots, \omega_N) \in \prod_{i=1}^N \Omega_i = \Omega$.

Incumbent firm i earns a profit $\pi_i(\omega)$ from competing in the product market (price or quantity competition, depending on the application) that, following the literature, we treat as a reduced-form input into the model. While $\pi_i(\omega)$ depends on the current state of the game ω , it does not depend on the current investment and exit decisions. The cost of investment $c_i(\omega, a_i)$ as well as any cost or benefit pertaining to exit are simply added to $\pi_i(\omega)$. As a result, per-period payoffs are separable in the sense of Assumption 3.

In many applications of the Ericson and Pakes (1995) model, firm i has exclusive control over the evolution of ω_i through its investment, exit, and entry decisions (e.g., Besanko and Doraszelski 2004, Chen 2009, Doraszelski and Markovich 2007). Because the decisions of firm i affect its own state variable but not its rivals' state variables, the transition probabilities are separable in the sense of Assumption 3. In other applications, there is in addition a common shock such as an increase in the quality of the outside good or an industry-wide depreciation shock (e.g., Berry and Pakes 1993,

Gowrisankaran 1999, Fershtman and Pakes 2000, de Roos 2004, Markovich 2008). Assumption 3 accommodates a common shock because transitions effected by “nature” can be subsumed into those effected by one of the players.¹²

Because investment may or may not result in a favorable outcome, transitions due to investment decisions are noisy as required by Assumption 2. Transitions due to entry and exit decisions present a difficulty because in the Ericson and Pakes (1995) model, an incumbent firm can exit the industry for sure and a potential entrant can enter the industry for sure. Doraszelski and Judd (2012) show how to formulate exit and entry with finite hazard rates either by way of exit and entry intensities or by way of randomly drawn, privately observed scrap values and setup costs (as in Doraszelski and Satterthwaite 2010). Their formulation satisfies Assumption 2.

Example 3 (Continuous-Time Stochastic Games with Moves at Random Times) Arcidiacono, Bayer, Blevins, and Ellickson (2016) develop a continuous-time stochastic game in which a player is given the move at random times. Decisions are asynchronous as the probability that more than one player has the move at a given time is zero. Ambrus and Lu (2015), Ambrus and Ishii (2015), Calcagno, Kamada, Lovo, and Sugaya (2014), and Kamada and Kandori (2017) develop closely related continuous-time stochastic games with moves at random times.

Arcidiacono, Bayer, Blevins, and Ellickson (2016) endow player i with a Poisson process with a constant hazard rate λ . The time between jumps in this process is therefore exponentially distributed. If process i is the first of the N processes to jump, then player i is given the move and chooses an action a_i . The state of the game then changes from ω to ω' with probability $l_i(\omega' \mid \omega, a_i)$, with $l_i(\cdot \mid \omega, a_i) \in \mathbb{P}(\Omega)$.

We can formulate this process in our framework by defining the hazard rate

$$q_i(\omega, a_i) = \lambda(1 - l_i(\omega \mid \omega, a_i))$$

and the transition probability¹³

$$p_i(\omega' \mid \omega, a_i) = \begin{cases} \frac{1}{(1 - l_i(\omega \mid \omega, a_i))} l_i(\omega' \mid \omega, a_i) & \text{if } \omega' \neq \omega, \\ 0 & \text{if } \omega' = \omega. \end{cases}$$

¹²One may alternatively represent nature by an extra player 0.

¹³If $l_i(\omega \mid \omega, a_i) = 1$, then $p_i(\cdot \mid \omega, a_i)$ can be defined arbitrarily.

Finally, the flow payoff of player i is

$$u_{i,j}(\omega, a_j) = \begin{cases} s_i(\omega) + \lambda\pi_i(\omega, a_i) & \text{if } i = j, \\ 0 & \text{if } i \neq j, \end{cases}$$

where $s_i(\omega)$ is a baseline payoff and $\pi_i(\omega, a_i)$ an additional payoff that player i receives if he is given the move. To account for the likelihood that player i is given the move, $\pi_i(\omega, a_i)$ is multiplied by λ in the flow payoff. The flow payoff and transition probability in Arcidiacono, Bayer, Blevins, and Ellickson (2016) clearly conform to Assumptions 2 and 3.

Example 4 (Dynamic Public Contribution Games) Consider N players that contribute towards completing a public project (Marx and Matthews 2000, Compte and Jehiel 2004, Georgiadis 2015). Completing the project requires K steps and $\omega \in \Omega = \{0, 1, \dots, K\}$ indicates the number of steps that have been completed. Player i 's contribution $a_i \in A_i(\omega) \subseteq \mathbb{R}$ induces a hazard rate $q_i(\omega, a_i)$ which is strictly increasing in a_i if $\omega \neq K$, while $q_i(\omega, a_i) = 0$ if $\omega = K$. The transition probability is

$$p_i(\omega' | \omega, a_i) = \begin{cases} 1 & \text{if } \omega' = \omega + 1, \omega \leq K - 1, \\ 1 & \text{if } \omega' = 0, \omega = K, \\ 0 & \text{otherwise.} \end{cases}$$

In this model, a player's action determines the probability that the project advances by one step. The public project is completed once state $\omega = K$ is reached and results in flow payoffs B_i for player i . The cost of contribution is $c_i(\omega, a_i)$ for player i . We therefore specify its flow payoff as

$$u_{i,i}(\omega, a_i) = \begin{cases} B_i - c_i(\omega, a_i) & \text{if } \omega = K, \\ -c_i(\omega, a_i) & \text{otherwise,} \end{cases}$$

and $u_{i,j}(\omega, a_j) = 0$ if $j \neq i$. Assumptions 2 and 3 are satisfied.

Example 5 (Asynchronously Repeated Games) The literature on repeated games has paid particular attention to the timing of decisions. Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) study discrete-time repeated games with asynchronous moves. These games can be re-cast to satisfy our assumptions.

We set $N = 2$, $\Omega = \Omega_1 \times \Omega_2$, $A_i(\omega) = \Omega_i$. The flow payoff of player i takes the form $u_{i,j}(\omega, a_j) = s_i(\omega)$. Depending on the application, $s_i(\omega)$ is the payoff from Bertrand competition (Maskin and

Tirole 1988a), Cournot competition (Maskin and Tirole 1988b), or a coordination game (Lagunoff and Matsui 1997). The hazard rate is

$$q_i(\omega, a_i) = \begin{cases} 0 & \text{if } a_i = \omega_i, \\ \lambda & \text{otherwise} \end{cases}$$

and the transition probability is

$$p_i(\omega' | \omega, a_i) = \begin{cases} 1 & \text{if } a_i = \omega'_i, a_i \neq \omega_i, \\ 0 & \text{otherwise.} \end{cases}$$

We thus obtain a game in which a player's state simply records the last chosen action.¹⁴

A separable dynamic game with noisy transitions built from these primitives under a protocol of alternating moves is similar to the asynchronously repeated games in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) in that changes in the payoff-relevant state do not occur at the same time.¹⁵ An important difference is that in our setting Assumption 2 precludes the state from instantaneously recording players' actions whereas in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) the player who has the move can change the state with probability one. We come back to this difference in Section 4.1, where we extend our protocol-invariance theorem to show that in contrast to Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) the protocol of moves ceases to matter for equilibrium behavior.

3 Protocol-Invariance Theorem

Consider the separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$. We are interested in exploring how the set of Markov perfect equilibria $\text{Equil}(\Gamma)$ of the game Γ changes as we change the protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$.

We endow the set of all flow payoffs $u \in \mathbb{R}^N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$ with the Lebesgue measure and say that a property is *generic* if it does not hold at most on a closed subset of measure zero. In this case we say that the property holds for *almost all* $u \in \mathbb{R}^N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$.

¹⁴This game is a special case of Example 3 obtained by setting $l_i(\omega' | \omega, a_i) = 1$ if and only if $a_i = \omega'_i$.

¹⁵Similar to Maskin and Tirole (1988a, 1988b), we restrict attention to Markov perfect equilibria from the outset. Lagunoff and Matsui (1997) show that there is no loss in doing so in their model, as any outcome that can be sustained by a subgame perfect equilibrium can also be sustained by a Markov perfect equilibrium.

The main result of the paper is a protocol-invariance theorem:

Theorem 1 (Protocol-Invariance Theorem) *Fix p, q , and ρ . For almost all u , all $\langle \mathcal{J}, \mathcal{P} \rangle$ and $\langle \overline{\mathcal{J}}, \overline{\mathcal{P}} \rangle$, and all $\varepsilon > 0$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$ and $\sigma \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$, there exists $\bar{\sigma} \in \text{Equil}(\langle \Delta, \overline{\mathcal{J}}, \overline{\mathcal{P}}, u, p, q, \rho \rangle)$ such that $\|\sigma - \bar{\sigma}\| < \varepsilon$.*

In words, for any Markov perfect equilibrium σ of a game with a protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$, the game with another protocol $\langle \overline{\mathcal{J}}, \overline{\mathcal{P}} \rangle$ has a Markov perfect equilibrium $\bar{\sigma}$ that is arbitrarily close to σ provided that periods are sufficiently short. Theorem 1 thus shows that the set of Markov perfect equilibria of separable dynamic games with noisy transitions is generically almost independent of the protocol of moves.

The proof of Theorem 1 proceeds in two steps. First, we take the equilibrium conditions for a separable dynamic game with noisy transitions as the period length goes to zero and show that the limit conditions are independent of the protocol of moves. Second, we show that the set of equilibria of a separable dynamic game with noisy transitions is, irrespective of the protocol of moves, arbitrarily close to the solutions to the limit conditions provided that periods are sufficiently short.

To begin establishing Theorem 1, consider a Markov perfect equilibrium $\sigma^\Delta = (\sigma_i^\Delta)_{i=1}^N$ of the separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$. Let $V_i^\Delta(\omega, J)$ be the continuation value of player i if players $J \in \mathcal{J}$ have the move and the state is $\omega \in \Omega$. The discrete-time Bellman equation for a period length of Δ is

$$V_i^\Delta(\omega, J) = u_i^\Delta(\omega, J, \sigma_J^\Delta(\omega)) + \exp(-\rho\Delta) \sum_{\omega' \in \Omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J'|J) \Pr^\Delta(\omega'|\omega, J, \sigma_J^\Delta(\omega)),$$

where the player discounts payoffs accruing in the subsequent period by $\exp(-\rho\Delta)$ and $\sigma_J^\Delta(\omega) = (\sigma_j^\Delta(\omega))_{j \in J}$. Under Assumptions 2 and 3 this becomes

$$\begin{aligned} V_i^\Delta(\omega, J) = & |\mathcal{J}| \sum_{j \in J} u_{i,j}(\omega, \sigma_j^\Delta(\omega))\Delta + \exp(-\rho\Delta) \left\{ \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J'|J) \left(1 - |\mathcal{J}| \sum_{j \in J} q_j(\omega, \sigma_j^\Delta(\omega))\Delta \right) \right. \\ & \left. + \sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J'|J) \left(|\mathcal{J}| \sum_{j \in J} \varphi_j(\omega'|\omega, \sigma_j^\Delta(\omega))\Delta \right) \right\} + \mathcal{O}(\Delta^2), \end{aligned} \quad (3.1)$$

where we use the shorthand notation $\varphi_j(\omega' | \omega, a_j) = q_j(\omega, a_j)p_j(\omega'|\omega, a_j)$ and $\varphi_j(\omega' | \omega, \sigma_j(\omega)) = \sum_{a_j \in A_j(\omega)} \varphi_j(\omega' | \omega, a_j)\sigma_j(a_j | \omega)$.

Let $V^\Delta = (V_i^\Delta)_{i=1}^N$ be the profile of value functions corresponding to the Markov perfect equilibrium σ^Δ . Consider a sequence $(\sigma^\Delta, V^\Delta)$ indexed by the period length Δ . Assuming that $(\sigma^\Delta, V^\Delta) \rightarrow (\sigma^0, V^0)$ (where convergence is possibly through a subsequence Δ_n) and taking the limit of equation (3.1) as $\Delta \rightarrow 0$, we deduce that

$$V_i^0(\omega, J) = \sum_{J' \in \mathcal{J}} V_i^0(\omega, J') \Pr(J'|J).$$

Stacking this equation for all $J \in \mathcal{J}$ yields the system of linear equations $\mathcal{P}x = x$, where x is a $|\mathcal{J}|$ -dimensional column vector with entries $V_i^0(\omega, J)$. Assumption 1 implies that $V_i^0(\omega, J) = V_i^0(\omega, J')$ for all $J, J' \in \mathcal{J}$.¹⁶ This means that in the limit the equilibrium continuation value of player i is independent of the identity of the players who have the move and equals $V_i^0(\omega)$: having the move does not imply a higher or lower payoff. From hereon, let $V_i^0: \Omega \rightarrow \mathbb{R}$ be the value function of player i and $V^0 = (V_i^0)_{i=1}^N$ be the profile of value functions in the limit as $\Delta \rightarrow 0$.

Equation (3.1) can equivalently be written as

$$\begin{aligned} & \frac{1}{\Delta} V_i^\Delta(\omega, J) - \frac{\exp(-\rho\Delta)}{\Delta} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J'|J) = \mathcal{O}(\Delta) + |\mathcal{J}| \sum_{j \in J} u_{i,j}(\omega, \sigma_j^\Delta(\omega)) \\ & + \exp(-\rho\Delta) |\mathcal{J}| \sum_{j \in J} \left(\sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J'|J) \varphi_j(\omega' | \omega, \sigma_j^\Delta(\omega)) - \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J'|J) q_j(\omega, \sigma_j^\Delta(\omega)) \right). \end{aligned}$$

Summing this equation for all $J \in \mathcal{J}$ and using the fact that, under Assumption 1, $\sum_{J \in \mathcal{J}} \Pr(J'|J) = 1$ to simplify the left-hand side¹⁷ yields

$$\begin{aligned} & \frac{1 - \exp(-\rho\Delta)}{\Delta} \sum_{J \in \mathcal{J}} V_i^\Delta(\omega, J) = \mathcal{O}(\Delta) + |\mathcal{J}| \sum_{J \in \mathcal{J}} \sum_{j \in J} u_{i,j}(\omega, \sigma_j^\Delta(\omega)) \\ & + \exp(-\rho\Delta) |\mathcal{J}| \sum_{J \in \mathcal{J}} \sum_{j \in J} \left(\sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J'|J) \varphi_j(\omega' | \omega, \sigma_j^\Delta(\omega)) - \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J'|J) q_j(\omega, \sigma_j^\Delta(\omega)) \right). \end{aligned}$$

¹⁶The vector $y = (1, \dots, 1)'$ is always a right eigenvector since \mathcal{P} is stochastic. Since \mathcal{P} is irreducible, the Perron-Frobenius theorem implies that both the left and right eigenvectors associated to the eigenvalue 1 are unique, up to scalar multiplication. It follows that for any solution to the system $\mathcal{P}x = x$, $x_i = x_j$ for all i and j .

¹⁷Recall that if the transition matrix \mathcal{P} is irreducible and its unique stationary distribution is uniform on \mathcal{J} , then \mathcal{P} is doubly stochastic. Hence, the left-hand side becomes $\frac{1}{\Delta} \sum_{J \in \mathcal{J}} V_i^\Delta(\omega, J) - \frac{\exp(-\rho\Delta)}{\Delta} \sum_{J \in \mathcal{J}} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J'|J) = \frac{1 - \exp(-\rho\Delta)}{\Delta} \sum_{J \in \mathcal{J}} V_i^\Delta(\omega, J)$.

Taking the limit as $\Delta \rightarrow 0$ yields the continuous-time Bellman equation

$$\begin{aligned} \rho V_i^0(\omega) &= \sum_{J \in \mathcal{J}} \sum_{j \in J} u_{i,j}(\omega, \sigma_j^0(\omega)) + \sum_{J \in \mathcal{J}} \sum_{j \in J} \left(\sum_{\omega' \neq \omega} V_i^0(\omega') \varphi_j(\omega' | \omega, \sigma_j^0(\omega)) - V_i^0(\omega) q_j(\omega, \sigma_j^0(\omega)) \right) \\ &= \sum_{j=1}^N u_{i,j}(\omega, \sigma_j^0(\omega)) + \sum_{j=1}^N \sum_{\omega' \neq \omega} (V_i^0(\omega') - V_i^0(\omega)) \varphi_j(\omega' | \omega, \sigma_j^0(\omega)), \end{aligned} \quad (3.2)$$

where the last equality uses that, under Assumption 1, there exists a unique $J \in \mathcal{J}$ such that $j \in J$ and the fact that $\sum_{\omega' \neq \omega} p_j(\omega' | \omega, a_i) = 1$. Importantly, condition (3.2) is independent of the protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ used to pass from discrete to continuous time.

The discrete-time optimality condition for a period length of Δ is

$$\begin{aligned} \sigma^\Delta(a_i | \omega) > 0 \Rightarrow \\ a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_i^\Delta(\omega, J, \tilde{a}_i, \sigma_{J \setminus \{i\}}^\Delta(\omega)) + \exp(-\rho\Delta) \sum_{\omega' \in \Omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J' | J) \Pr^\Delta(\omega' | \omega, J, \tilde{a}_i, \sigma_{J \setminus \{i\}}^\Delta(\omega)). \end{aligned}$$

Since $\sigma^\Delta \rightarrow \sigma^0$, $\sigma_i^0(a_i | \omega) > 0$ implies $\sigma_i^\Delta(a_i | \omega) > 0$ if the period length Δ is sufficiently small. Dividing by Δ , rearranging terms, and taking the limit as $\Delta \rightarrow 0$ (as we did in the previous paragraph) thus yields the continuous-time optimality condition

$$\sigma_i^0(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_{i,i}(\omega, \tilde{a}_i) + \sum_{\omega' \neq \omega} (V_i^0(\omega') - V_i^0(\omega)) \varphi_i(\omega' | \omega, \tilde{a}_i). \quad (3.3)$$

Condition (3.3) is again independent of the protocol of moves. It formalizes that player i faces the same tradeoff between current and future payoffs under any protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ and that this tradeoff is not directly affected by his rivals' actions.

The intuition is best seen by contrasting two protocols of moves in a separable dynamic game with noisy transitions. With alternating moves, if player i has the move, then to choose an action a_i he must consider the contribution $u_{i,i}(\omega, a_i)\Delta$ to his per-period payoff that his action yields and the impact his action has on state-to-state transitions through $q_i(\omega, a_i)p_i(\omega' | \omega, a_i)\Delta$ (neglecting the higher-order term $\mathcal{O}(\Delta^2)$). With simultaneous moves, two additional considerations arise. First, player i must consider how his rivals' actions change the contribution to his per-period payoff that his action yields and the impact his action has on state-to-state transitions. However, because Assumption 3 restricts complementarities between players' actions and other non-separabilities in per-period payoffs and state-to-state transitions to the higher-order term $\mathcal{O}(\Delta^2)$, player i can

neglect his rivals' actions if the period length Δ is sufficiently small. Second, player i must consider the possibility that his rivals' actions further change the state of the game. The probability that two or more players cause the state to change is, however, negligible by Assumption 2 if the period length Δ is sufficiently small. Assumption 1 finally ensures that irrespective of the protocol of moves all players move with the same frequency over a sufficiently large number of periods. Thus, in the limit as $\Delta \rightarrow 0$, player i faces the same tradeoff between current and future payoffs.

Conditions (3.2) and (3.3) are the limit as $\Delta \rightarrow 0$ of the equilibrium conditions for the separable dynamic game with noisy transitions Γ . We provide economic interpretations of these conditions in Section 3.1. Here we merely observe that they impose restrictions on the limit strategy and continuation value profiles (σ^0, V^0) . Noting that the limit conditions (3.2) and (3.3) may admit multiple solutions and that V^0 is entirely determined by σ^0 using condition (3.2), we denote the set of strategy profiles $\sigma^0 \in \Sigma$ satisfying condition (3.3) as $\text{Equil}^0(\langle u, p, q, \rho \rangle)$. This set does not depend on the protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ used to pass to the limit.

We summarize the above discussion in a lemma:

Lemma 1 *Consider a sequence (σ^Δ) with $\sigma^\Delta \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$. If $\sigma^\Delta \rightarrow \sigma^0$, then $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$.*

Unfortunately, Lemma 1 does not imply Theorem 1 because conditions (3.2) and (3.3) may admit multiple solutions. In this case, taking the limit of a sequence of equilibria under different protocols of moves may potentially lead to different solutions of the limit conditions. To overcome this difficulty, we would ideally show that all solutions to the limit conditions (3.2) and (3.3) can be approximated by the Markov perfect equilibria of a separable dynamic game with noisy transitions and an arbitrary protocol of moves provided that periods are sufficiently short.

An example makes plain that this cannot always be done. We consider the entry game in Example 1 with $K = 1$ and $(c_i, b_i, B_i) = (c, b, B)$ for all i . We restrict attention to parameters that render firms indifferent between investing and not investing in continuous time. Thus, some solutions to the limit conditions cannot be approximated in discrete time. Formally, we take $\lambda B/\rho = c$ and $b < 0$ so that a duopolist incurs a loss. A firm's only nontrivial decision is whether to invest in state $(0, 0)$. Define the pure strategy profile σ^0 in which both firms invest by $\sigma_i^0(1 | (0, 0)) = 1$. Let us verify that $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$. The limit conditions (3.2) and (3.3) become

$$\rho V_i^0(0, 0) = -c + \lambda \left(\frac{B}{\rho} - V_i^0(0, 0) \right) + \lambda (0 - V_i^0(0, 0))$$

and

$$-c + \lambda\left(\frac{B}{\rho} - V_i^0(0,0)\right) \geq 0, \quad (3.4)$$

where we use the fact that $V_i^0(\omega) = B/\rho$ when $\omega_i \neq \omega_{-i} = 0$ and $V_i^0(\omega) = 0$ when $\omega_i \neq \omega_{-i} = 1$. We deduce that $V_i^0(0,0) = \frac{1}{\rho+2\lambda}(-c + \lambda\frac{B}{\rho}) = 0$ and σ^0 is indeed a solution to the limit conditions.

However, this solution cannot be approximated by a discrete-time game with simultaneous moves. In this game, a firm has an incentive to invest if and only if

$$-c + e^{-\rho\Delta}\Delta\lambda^2\frac{b\Delta}{1-e^{-\rho\Delta}} + e^{-\rho\Delta}\lambda(1-\lambda\Delta)\frac{B\Delta}{1-e^{-\rho\Delta}} \geq 0. \quad (3.5)$$

Because $\lambda B/\rho = c$ and $b < 0$, this condition holds in the limit as $\Delta \rightarrow 0$ but not for $\Delta > 0$.¹⁸

This example illustrates the problem we have to solve in establishing Theorem 1. While the expected net present value of the stream of payoffs in a discrete-time game is virtually independent of the protocol of moves and arbitrarily close to its continuous-time counterpart, equilibrium behavior is governed by differences in continuation values. The mere fact that continuation values converge does not ensure that the sign of these differences coincide in discrete and continuous time. In the example, condition (3.4) shows that in the limit as $\Delta \rightarrow 0$ the payoff from investing is greater than or equal to the payoff from not investing. Yet, in the discrete-time game the fact that condition (3.5) does not hold for $\Delta > 0$ shows that the payoff from investing is less than the payoff from not investing.

We proceed as follows. To rule out the above example, we first restrict attention to solutions $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ that are *regular*. Regularity ensures that the limit conditions (3.2) and (3.3) form a well-posed system of equations with locally unique solutions.¹⁹ Regularity is easy to define, but rather tedious to check in practice. The formal definition of regularity is in Appendix A.1 and allows for both pure and mixed strategy profiles.²⁰ We note that a pure strategy profile is regular if it is strict, i.e., if the maximization problem in condition (3.3) admits a unique solution. Our key insight is that for a regular solution σ^0 of the limit conditions (3.2) and (3.3), the differences in continuation values that govern equilibrium behavior have the same sign as their discrete-time counterparts under any protocol of moves. As a result, σ^0 can be approximated by a Markov

¹⁸This solution also cannot be approximated by a discrete-time game with alternating moves. In Appendix B.3, we complement this example by showing that the limit solution can be approximated by a slightly modified discrete-time game with alternating moves but not with simultaneous moves.

¹⁹At a regular solution, a player puts strictly positive weight on all actions that maximize the expected net present value of his payoffs.

²⁰The limit conditions (3.2) and (3.3) may not admit a solution in pure strategies. An example is available upon request.

perfect equilibrium of a separable dynamic game with noisy transitions and an arbitrary protocol of moves. We then use differential topology tools to establish that for almost all flow payoffs u , the restriction to regular solutions is without loss of generality. This yields the following lemma:

Lemma 2 *Fix p , q , and ρ . For almost all u , all $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$, all $\langle \mathcal{J}, \mathcal{P} \rangle$, and all $\varepsilon > 0$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, there exists $\sigma \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$ such that $\|\sigma - \sigma^0\| < \varepsilon$.*

Lemmas 1 and 2 finally combine to yield Theorem 1.

The proof of Lemma 2 draws on and expands ideas in Harsanyi (1973a, 1973b). Existing genericity results for dynamic stochastic games (Haller and Lagunoff 2000, Doraszelski and Escobar 2010) do not apply to our setting because the limit conditions (3.2) and (3.3) are in continuous time and separable dynamic games with noisy transitions restrict per-period payoffs and state-to-state transitions and are therefore a subset of measure zero of the dynamic stochastic games considered in the literature.²¹

3.1 Interpretations of Limit Conditions

We offer two economic interpretations of the limit conditions (3.2) and (3.3). First, they can be interpreted as the equilibrium conditions for a continuous-time stochastic game along the lines of Doraszelski and Judd (2012). In this game, moves are simultaneous and the state follows a continuous-time Markov process that players can influence through their actions. Properly defining mixed strategies in continuous time is, however, subtle because it requires working with a continuum of independent and identically distributed random variables that satisfy a law of large numbers. As in Bolton and Harris (1999), we can use time to “purify” these strategies and avoid the continuum of independent and identically distributed random variables. Beyond this observation, we follow the literature and alert the reader that a rigorous foundation for mixed strategies in continuous time is an open problem (Bolton and Harris 1999, Faingold and Sannikov 2011).

Second, the limit conditions (3.2) and (3.3) can be interpreted as the equilibrium conditions for a dynamic stochastic game with random moves. The following construction, known as uniformization (Serfoso 1979), is adapted from single-agent decision problems. Fix $B > N \max_{j=1, \dots, N, \omega \in \Omega, a_j \in A_j(\omega)} q_j(\omega, a_j)$.

²¹Our proof also shows that $\text{Equil}^0(\langle u, p, q, \rho \rangle)$ consists of a finite number of isolated solutions. This generalizes results on the generic finiteness of the set of Markov perfect equilibria in Haller and Lagunoff (2000) and Doraszelski and Escobar (2010) to continuous-time stochastic games.

We use the parameter B to ensure that our construction yields a well-defined transition probability and discount factor in discrete time. Define the per-period payoff $\tilde{u}_{i,j}(\omega, a_j) = \frac{N}{\rho+B}u_{i,j}(\omega, a_j)$, the discount factor $\beta = \frac{B}{\rho+B} < 1$, and the transition probability

$$\tilde{\varphi}_j(\omega' | \omega, a_j) = \begin{cases} \frac{N}{B}\varphi_j(\omega' | \omega, a_j) & \text{if } \omega' \neq \omega, \\ 1 - \frac{N}{B}q_j(\omega, a_j) & \text{if } \omega' = \omega. \end{cases}$$

Note that $\tilde{\varphi}_j(\cdot | \omega, a_j) \in \mathbf{P}(\Omega)$ because we normalize using the parameter B . Now formulate a dynamic stochastic game with random moves in which in any period one player $j \in \{1, \dots, N\}$ is randomly and uniformly selected to make a decision $a_j \in A_j(\omega)$. Each player strives to maximize the expected net present value of his stream of payoffs and discounts future payoffs using the discount factor β . Denote by $\text{Equil}^R(\tilde{u}, \tilde{\varphi}, \beta)$ the set of Markov perfect equilibria of this dynamic stochastic game.

The following proposition shows that the Markov perfect equilibria of the dynamic stochastic game with random moves constructed above are the solutions of the limit conditions (3.2) and (3.3):

Proposition 1 $\text{Equil}^R(\langle \tilde{u}, \tilde{\varphi}, \beta \rangle) = \text{Equil}^0(\langle u, p, q, \rho \rangle)$.

The proof of Proposition 1 is simple and illustrative. The equilibrium conditions for $\sigma \in \text{Equil}^R(\langle \tilde{u}, \tilde{\varphi}, \beta \rangle)$ are

$$V_i(\omega) = \sum_{j=1}^N \frac{1}{N} \left(\tilde{u}_{i,j}(\omega, \sigma_j(\omega)) + \beta \sum_{\omega' \in \Omega} V_i(\omega') \tilde{\varphi}_j(\omega' | \omega, \sigma_j(\omega)) \right) \quad (3.6)$$

and

$$\sigma_i(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} \tilde{u}_{i,i}(\omega, \tilde{a}_i) + \beta \sum_{\omega' \in \Omega} V_i(\omega') \tilde{\varphi}_i(\omega' | \omega, \tilde{a}_i). \quad (3.7)$$

These conditions can be equivalently written as

$$V_i(\omega) = \sum_{j=1}^N \frac{1}{N} \left\{ \frac{N}{\rho+B}u_{i,j}(\omega, \sigma_j(\omega)) + \frac{B}{\rho+B} \left(\sum_{\omega' \neq \omega} V_i(\omega') \frac{N}{B}\varphi_j(\omega' | \omega, \sigma_j(\omega)) + (1 - \frac{N}{B}q_j(\omega, \sigma_j(\omega)))V_i(\omega) \right) \right\}$$

and

$$\sigma_i(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} \frac{N}{\rho + B} u_{i,i}(\omega, \tilde{a}_i) + \frac{B}{\rho + B} \left(\sum_{\omega' \neq \omega} V_i(\omega') \frac{N}{B} \varphi_i(\omega' | \omega, \tilde{a}_i) + \left(1 - \frac{N}{B} q_i(\omega, \tilde{a}_i)\right) V_i(\omega) \right)$$

Rearranging terms, the limit conditions (3.2) and (3.3) are therefore identical to the equilibrium conditions (3.6) and (3.7) for the dynamic stochastic game with random moves constructed above.

While dynamic stochastic games with random moves are sparsely used, several important papers study repeated games with alternating moves. For example, Maskin and Tirole (1988a) explore a repeated Bertrand game with alternating moves and show how Edgeworth cycles can arise. Maskin and Tirole (1988b) show that an analog to limit pricing can arise in a model of repeated Cournot competition with alternating moves and large fixed costs. Lagunoff and Matsui (1997) show how players can coordinate on the efficient outcome in a repeated coordination game with alternating moves. These results are driven by the fact that a player remains committed to his previously chosen action over a stretch of time.²² The dynamic stochastic game with random moves constructed above shares this feature. Similarly rich dynamic phenomena thus appear in the continuous-time stochastic game that we obtain as we pass to the limit and, by Theorem 1, in separable dynamic games with noisy transitions and arbitrary protocols of moves provided that periods are sufficiently short.

3.2 Discussion of Assumptions

To illustrate the tightness of our assumptions, we provide a series of examples showing that protocol invariance may fail if any one of them is relaxed.

Example 6 (Separability) *The literature provides a number of examples in which complementarities between players' actions and other non-separabilities in per-period payoffs preclude protocol invariance. Our example with non-separable per-period payoffs is inspired by Lagunoff and Matsui (1997) and Wen (2002). In Appendix B.4, we present a closely related example with non-separable state-to-state transitions.*

Consider a coordination game with the following payoff matrix:

²²As Lagunoff and Matsui (1997) point out, what matters is that moves are asynchronous “rather than the specific structure of asynchronous choice” (p. 1473).

	L	R
T	2, 2	0, 0
B	0, 0	1, 1

Denote by $b(a) = (b_1(a), b_2(a))$ the payoff profile given the action profile $a = (a_1, a_2) \in \{T, B\} \times \{L, R\}$.

We construct a dynamic stochastic game with a trivial state space $|\Omega| = 1$ (that we omit along with specifying the transition probability) and contrast the set of Markov perfect equilibria under simultaneous and alternating moves. In the game with simultaneous moves, the per-period payoff of player i is $u_i^\Delta(\omega, \{1, 2\}, a) = b_i(a)\Delta$. Irrespective of the period length Δ , there are two Markov perfect equilibria, namely $\sigma_1(T) = \sigma_2(L) = 1$ and $\tilde{\sigma}_1(B) = \tilde{\sigma}_2(R) = 1$.

As per-period payoffs are not separable, we require auxiliary assumptions to construct the game with alternating moves. We assume a repeated extensive-form game along the lines of Rubinstein and Wolinsky (1995). In violation of Assumption 3 the per-period payoff of player i is $u_i^\Delta(\omega, \{1\}, a_1) = 0$ and $u_i^\Delta(\omega, \{2\}, a) = b_i(a)\Delta$, meaning that payoffs “materialize” after player 2 moves.²³ Since player 1’s action a_1 is payoff relevant for player 2, a Markovian strategy for player 2 includes a_1 as a state variable. Irrespective of the period length Δ , player 2 has an incentive to coordinate and match his action to that of player 1. Thus, the unique Markov perfect equilibrium is $\sigma_1(T) = 1$ and $\sigma_2(T | L) = \sigma_2(B | R) = 1$.

Example 6 can also be used to show how protocol-invariance may fail when the assumption of noisy transitions is relaxed. We discuss this point in more generality in Section 4.1, where we explore a model in which the current state is a record of past actions and fully determines per-period payoffs.

Example 7 (Noisy Transitions) Consider the entry game in Example 1 with $K = 1$. Two firms may enter a new market in a single step. We further set $b_i < 0$ for all i , so that a duopolist incurs a loss.

We change Example 1 by assuming that the transition probability takes the form

$$\Pr^\Delta(\omega' | \omega, J, a_J) = \begin{cases} 1 & \text{if } \omega'_J = a_J, \omega = (0, 0), \\ 1 & \text{if } \omega' = \omega, \omega_i = 1 \text{ for some } i, \\ 0 & \text{otherwise.} \end{cases}$$

²³We may alternatively formulate the game with alternating moves following Lagunoff and Matsui (1997) so that payoffs “materialize” after any player moves. This does not change the message of the example.

Hence, irrespective of the protocol of moves, if firm i takes action $a_i = 1$, then its state changes for sure from $\omega_i = 0$ to $\omega'_i = 1$. It is clear that this transition probability does not satisfy Assumption 2. In the game with alternating moves, if Δ is sufficiently small, then the unique Markov perfect equilibrium outcome is that the firm that moves first enters whereas its rival never enters. In contrast, in the game with simultaneous moves, there exists a Markov perfect equilibrium in which both firms enter with positive probability.

Comparing Example 7 to Example 1 makes clear that the assumption of noisy transitions rules out uncoordinated entry. It is easy to see that in Example 1 with $K = 1$ and $b_i < 0$ for all i , only one firm enters in the unique Markov perfect equilibrium regardless of the protocol of moves provided moves are frequent.

Finally, we show that protocol invariance does not extend beyond Markov perfect equilibria to more general equilibrium concepts. For a separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$, we say that a strategy σ_i^T for player i has *finite memory* $T \geq 0$ if $\sigma_i^T(h) = \sigma_i^T(\tilde{h})$ for any two histories h and \tilde{h} (perhaps of different length) that coincide in the current state and the outcomes of the previous T rounds of interactions between players. If $T = 0$, then we recover the definition of a stationary Markovian strategy in Section 2.

Example 8 (Markov Perfect Equilibrium) Consider a partnership game and construct a separable dynamic game with noisy transitions and a trivial state space (that we again omit). There are $N = 2$ players. The set of actions of player i is $A_i = \{0, 1\}$, his flow payoff is

$$u_{i,j}(a_j) = \begin{cases} -a_j & \text{if } j = i, \\ 2a_j & \text{if } j \neq i, \end{cases}$$

and the discount rate is ρ . Intuitively, by choosing action $a_i = 1$, player i incurs a cost of 1 but confers a positive externality of 2 on player j . Irrespective of the protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ and the period length Δ , the unique Markov perfect equilibrium of this game is $\sigma_1(0) = \sigma_2(0) = 1$ and has players repeating $(0, 0)$.

We show that this is not the case for strict subgame perfect equilibria in finite memory strategies. In the game with simultaneous moves, consider a finite memory strategy σ_i^T with $T \geq 1$ for player i such that player i chooses $a_i^t = 1$ in period t if $t = 0$ or if the players have chosen the same action over the last $\min\{T, t\}$ rounds: $a_1^{\tilde{t}} = a_2^{\tilde{t}}$ for all $\tilde{t} \in \{t - 1, \dots, t - \min\{t, T\}\}$. The strategy profile $\sigma^T = (\sigma_1^T, \sigma_2^T)$ is a strict subgame perfect equilibrium if $\frac{1 - e^{-\rho\Delta T}}{1 - e^{-\rho\Delta}} > e^{\rho\Delta}$. This condition holds if

$T \geq 2$ and the period length Δ is sufficiently small. Hence, there exists a strict subgame perfect equilibrium in finite memory strategies in which players repeatedly play $(1, 1)$.

Turning to the game with alternating moves, consider a finite memory strategy σ_i^T with $T \geq 1$ for player i . We argue that for any strategy profile σ^T to be a strict subgame perfect equilibrium it must be a Markov perfect equilibrium. Hence, the unique strict subgame perfect equilibrium in finite memory strategies is the Markov perfect equilibrium in which players repeat $(0, 0)$.

To complete the argument, suppose player i moves in round t . Because player $-i$ moves after player i and conditions his decision on the previous T periods of interactions, the continuation value of player i depends on a_i^t and the previous $T - 1$ periods of interactions. The current payoff of player i moreover depends only on a_i^t . Since σ^T is a strict subgame perfect equilibrium, the maximization problem of player i admits a unique solution which depends, at most, on the previous $T - 1$ rounds of interactions. This means that σ_i^T actually conditions on the previous $T - 1$ periods of interactions. Continuing iteratively, we deduce that the strategy profile σ^T cannot condition on any previous interactions.

4 Applications and Extensions

We apply and extend our main result in three ways. We first provide an extension to arbitrarily large hazard rates. Then we offer a rationale for focusing on Markov perfect equilibria and discuss computing these equilibria.

4.1 Protocol Invariance with Arbitrarily Large Hazard Rates

We modify the model in Section 2 by assuming that hazard rates are of form $\tilde{q}_i(\omega, a_i) = \lambda q_i(\omega, a_i)$, where $\lambda \geq 1$ is a parameter. As $\lambda \rightarrow \infty$, hazard rates become arbitrarily large. We explore how $\text{Equil}^0(\langle u, p, \tilde{q}, \rho \rangle)$ changes in response.

Consider $\sigma^{0,\lambda} \in \text{Equil}^0(\langle u, p, \tilde{q}, \rho \rangle)$ and assume that $\sigma^{0,\lambda} \rightarrow \sigma^{0,\infty}$ as $\lambda \rightarrow \infty$. Taking the limit of condition (3.2) yields the Bellman equation

$$\rho V_i^{0,\lambda}(\omega) = \sum_{j=1}^N u_{i,j}(\omega, \sigma_j^{0,\lambda}(\omega)) + \sum_{j=1}^N \sum_{\omega' \neq \omega} \left(V_i^{0,\lambda}(\omega') - V_i^{0,\lambda}(\omega) \right) \lambda \varphi_j(\omega' | \omega, \sigma_j^{0,\lambda}(\omega)),$$

where $\varphi_j(\omega' | \omega, a_j) = q_j(\omega, a_j) p_j(\omega' | \omega, a_j)$. Fix a state $\omega^0 \in \Omega$ and define the function $h_i^\lambda(\omega) =$

$\lambda(V_i^{0,\lambda}(\omega) - V_i^{0,\lambda}(\omega_0))$. Assuming that $h_i^\lambda(\omega) \rightarrow h_i(\omega)$ for some $h_i: \Omega \rightarrow \mathbb{R}$, the Bellman equation becomes

$$v_i = \sum_{j=1}^N u_{i,j}(\omega, \sigma_j^{0,\infty}(\omega)) + \sum_{j=1}^N \sum_{\omega' \neq \omega} (h_i(\omega') - h_i(\omega)) \varphi_j(\omega' | \omega, \sigma_j^{0,\infty}(\omega)), \quad (4.1)$$

where $v_i \in \mathbb{R}$ does not depend on ω .²⁴ Analogously, taking the limit of condition (3.3) as $\lambda \rightarrow \infty$ yields the optimality condition

$$\sigma_i^{0,\infty}(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_{i,i}(\omega, \tilde{a}_i) + \sum_{\omega' \neq \omega} (h_i(\omega') - h_i(\omega)) \varphi_i(\omega' | \omega, \tilde{a}_i). \quad (4.2)$$

Conditions (4.1) and (4.2) extend conditions (3.2) and (3.3) to arbitrarily large hazard rates and characterize the solutions to the limit conditions.

The following proposition summarizes the discussion:

Proposition 2 (Protocol Invariance with Arbitrarily Large Hazard Rates) *Consider a sequence $(\sigma^{0,\lambda})$ with $\sigma^{0,\lambda} \in \text{Equil}^0(\langle u, p, \tilde{q}, \rho \rangle)$. Assume $\sigma^{0,\lambda} \rightarrow \sigma^{0,\infty}$ as $\lambda \rightarrow \infty$. Fix $\omega^0 \in \Omega$ and assume that, for all i , $\lambda(V_i^{0,\lambda}(\omega) - V_i^{0,\lambda}(\omega_0))$ is uniformly bounded.²⁵ Then, for all i , there exists $h_i: \Omega \rightarrow \mathbb{R}$ and $v_i \in \mathbb{R}$, such that $\sigma^{0,\infty}$ satisfies conditions (4.1) and (4.2).*

Proposition 2 extends protocol invariance to the limiting case of deterministic transitions and provides a novel dynamic programming characterization of separable dynamic games with noisy transitions as moves become arbitrarily frequent and hazard rates arbitrarily large.²⁶

Given our previous effort in Example 7 to show that protocol invariance fails if transitions are deterministic, Proposition 2 may seem puzzling. Example 7 and Proposition 2 can be reconciled by noting that there is a discontinuity in the set of Markov perfect equilibria as moves become arbitrarily frequent and hazard rates arbitrarily large.

²⁴To see this, note that for all $\epsilon > 0$, there exists $\bar{\lambda}$ such that for all $\lambda > \bar{\lambda}$, $|V_i^\lambda(\omega) - V_i^\lambda(\omega^0) - h(\omega)/\lambda| < \epsilon/\lambda$. As a result, $\lim_{\lambda \rightarrow \infty} V_i^\lambda(\omega) = \lim_{\lambda \rightarrow \infty} V_i^\lambda(\omega^0)$.

²⁵Note that this is a restriction on endogenous variables. To provide a sufficient condition, fix a strategy profile σ and define $\tau = \inf\{t \mid \omega^t = \omega_0\}$. If $\mathbb{E}_\sigma[\tau \mid \omega^{t=0} = \omega]$ is finite for all $\omega \in \Omega$, all $\sigma \in \Sigma$, and all λ sufficiently large, then $\lambda(V_i^{0,\lambda}(\omega) - V_i^{0,\lambda}(\omega_0))$ is uniformly bounded. This type of condition appears in dynamic programming problems without discounting (Arapostathis, Borkar, Fernández-Gaucherand, Ghosh, and Marcus 1993).

²⁶This result contributes to the well-known literature drawing often subtle connections between discrete- and continuous-time stochastic games with infinite hazard rates (Fudenberg and Tirole 1985, Simon and Stinchcombe 1989). Fudenberg and Tirole (1985), in particular, show that passing to the limit is non-trivial in games with infinite hazard rates even if strategies are restricted to be Markovian. Proposition 2, in contrast, provides a quite tractable model of a dynamic game with arbitrarily frequent moves and arbitrarily large hazard rates.

Many examples from the extant literature where equilibrium behavior hinges on the protocol of moves even if periods are short (as in Maskin and Tirole 1988a, 1988b and Lagunoff and Matsui 1997) can be seen as a manifestation of this discontinuity. In the rest of this section, we expose this discontinuity and illustrate Proposition 2 using a dynamic coordination game similar to the one studied by Lagunoff and Matsui (1997).

We consider a two-player game as the one in Example 5. The set of states is $\Omega = \Omega_1 \times \Omega_2$, with $|\Omega_i| \geq 2$ for all i . The set of actions of player i is $A_i(\omega) = \Omega_i$. We now assume that the flow payoffs of the players coincide:

$$s_i(\omega, a_j) = \frac{1}{2}\pi(\omega).$$

Denote $\omega^* = \arg \max_{\omega \in \Omega} \pi(\omega)$ (we assume the argmax is unique). Use the function $\pi: \Omega \rightarrow \mathbb{R}$ to define the payoff matrix of a normal-form game in which players simultaneously choose actions $a_i \in \Omega_i$. We assume that this normal-form game has a Nash equilibrium ω^{NE} such that $\omega_i^{NE} \neq \omega_i^*$ for all i . Because ω^* is the unique maximizer of the function π , ω^* is also a Nash equilibrium. An example of such game is the payoff matrix in Example 6 with $\Omega_1 = \{T, B\}$, $\Omega_2 = \{L, R\}$ and $\pi(T, L) = 2$, $\pi(B, R) = 1$, and $\pi(\omega) = 0$ otherwise. In this example, $\omega^* = (T, L)$ and $\omega^{NE} = (B, R)$.²⁷

To illustrate Proposition 2 and protocol invariance with arbitrarily large hazard rates, define $\sigma_i^*(\omega) = \omega_i^*$ for all i and all ω and observe that

$$\lim_{\lambda \rightarrow \infty} \text{Equil}^0(\langle u, p, q, \rho \rangle) = \lim_{\lambda \rightarrow \infty} \lim_{\Delta \rightarrow 0} \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle) = \{\sigma^*\}.$$

This means that with arbitrarily frequent moves and arbitrarily large hazard rates, players coordinate on the efficient state ω^* .²⁸

To expose the discontinuity in the set of Markov perfect equilibria, consider a protocol of simultaneous moves $\langle \mathcal{J}^{sim}, \mathcal{P}^{sim} \rangle$. Imposing $\lambda\Delta = 1$, players determine the state with probability

²⁷The formulation of the game here differs from Example 6. In Example 6, the game is not separable because players' actions are complements in per-period payoffs. In contrast, here we rule out complementarities in per-period payoffs by assuming that a player's action determines the distribution over states but not period payoffs.

²⁸The logic of this result follows from Lagunoff and Matsui (1997). If player i 's state is $\omega_i = \omega_i^*$, then player $-i$ has an incentive to choose $a_{-i} = \omega_{-i}^*$ to obtain $\pi(\omega^*)$. In a given state ω , player i thus knows that if his state changes to ω_i^* , then his rival will switch to action $a_{-i} = \omega_{-i}^*$ relatively soon as long as λ is sufficiently large. The unique limit solution is therefore σ^* .

one when they move. It is relatively simple to show that

$$\{\sigma^*, \sigma^{NE}\} \subseteq \text{Equil}(\langle \Delta, \mathcal{J}^{sim}, \mathcal{P}^{sim}, u, p, q, \rho \rangle),$$

where $\sigma_i^{NE}(\omega) = \omega_i^{NE}$ for all i and all ω . Intuitively, if $\lambda\Delta = 1$, then players can coordinate on one of the two Nash equilibria ω^{NE} and ω^* . We conclude that

$$\lim_{\lambda \rightarrow \infty} \lim_{\Delta \rightarrow 0} \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle) \subsetneq \liminf_{\lambda\Delta=1, \Delta \rightarrow 0} \text{Equil}(\langle \Delta, \mathcal{J}^{sim}, \mathcal{P}^{sim}, u, p, q, \rho \rangle)$$

and thus there is a discontinuity in the joint limit as moves become arbitrarily frequent and hazard rates arbitrarily large.

Finally, there is a sense in which we can replicate the results in Lagunoff and Matsui (1997) that the protocol matters when moves are frequent. We show that the limit of the set of Markov perfect equilibria as $\Delta \rightarrow 0$ keeping $\lambda\Delta$ constant depends on the protocol of moves. Consider a protocol of alternating moves $\langle \mathcal{J}^{alt}, \mathcal{P}^{alt} \rangle$. Imposing $\lambda\Delta = 1/2$, a player determines his state with probability one when he moves.²⁹ From Theorem 1 in Lagunoff and Matsui (1997), if the period length Δ is sufficiently small, then³⁰

$$\text{Equil}(\langle \Delta, \mathcal{J}^{alt}, \mathcal{P}^{alt}, u, p, q, \rho \rangle) = \{\sigma^*\}.$$

This implies that

$$\lim_{\lambda\Delta=1/2, \Delta \rightarrow 0} \text{Equil}(\langle \Delta, \mathcal{J}^{alt}, \mathcal{P}^{alt}, u, p, q, \rho \rangle) \subsetneq \liminf_{\lambda\Delta=1, \Delta \rightarrow 0} \text{Equil}(\langle \Delta, \mathcal{J}^{sim}, \mathcal{P}^{sim}, u, p, q, \rho \rangle).$$

In this sense, we can replicate the results in Lagunoff and Matsui (1997) by taking the joint limit $\lambda\Delta \rightarrow 1$ and $\Delta \rightarrow 0$, but this is just one of many ways of taking the joint limit in our setting. Proposition 2, in contrast, shows that protocol invariance arises if we first take moves to be arbitrarily frequent and then take hazard rates to be arbitrarily large.

²⁹Recall from Assumption 3 that the hazard rates are scaled by $|\mathcal{J}|$. This is the reason we take $\lambda\Delta = 1$ with simultaneous moves and $\lambda\Delta = 1/2$ with alternating moves.

³⁰The restriction to a discount factor close to one in Lagunoff and Matsui (1997) translates into a period length Δ close to zero in our setting.

4.2 Justification of Markov Perfect Equilibria

We apply our main result to justify Markov perfect equilibria in a class of dynamic stochastic games, thereby complementing several arguments given for a variety of dynamic models in favor of Markov perfect equilibria (Maskin and Tirole 2001, Bhaskar and Vega-Redondo 2002, Sannikov and Skrzypacz 2007, Faingold and Sannikov 2011, Bhaskar, Mailath, and Morris 2013, Bohren 2014). Provided that periods are sufficiently short and a robustness requirement is satisfied, we show that the set of Markov perfect equilibrium payoffs in separable dynamic games with noisy transitions and simultaneous moves almost coincides with the set of payoffs that can be attained under more general equilibrium concepts.

We focus on strict subgame perfect equilibria in finite memory strategies. By definition, a strict equilibrium involves only pure strategies. Strictness is a natural robustness requirement. In repeated public monitoring games only strict subgame perfect equilibria in finite memory strategies are robust to private monitoring (Mailath and Morris 2002, Mailath and Samuelson 2006, Bhaskar, Mailath, and Morris 2013). Equilibria that fail to be strict are also fragile to perturbations of payoffs and information (Harsanyi 1973a, Harsanyi 1973b, Doraszelski and Escobar 2010).

As we change the protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ of a separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$, the sets of histories change and are therefore difficult to compare. To circumvent this difficulty, we explore how the set of payoff profiles $\text{Payoffs}^F(\Gamma) \subseteq \mathbb{R}^N$ associated with strict subgame perfect equilibria in finite memory strategies changes as we change the protocol of moves. We also define the set of payoff profiles $\text{Payoffs}^M(\Gamma) \subseteq \mathbb{R}^N$ corresponding to the set of Markov perfect equilibria $\text{Equil}(\Gamma)$.

Let $\Gamma^{sim} = \langle \Delta, \mathcal{J}^{sim}, \mathcal{P}^{sim}, u, p, q, \rho \rangle$ denote a separable dynamic game with noisy transitions under a protocol of simultaneous moves $\langle \mathcal{J}^{sim}, \mathcal{P}^{sim} \rangle$, with $\mathcal{J}^{sim} = \{\{1, \dots, N\}\}$. We say that the payoff profile $v \in \text{Payoffs}^F(\Gamma^{sim})$ is *approachable* if for all $\epsilon > 0$ there exists some protocol of asynchronous moves $\langle \mathcal{J}^{asy}, \mathcal{P}^{asy} \rangle$, with $\mathcal{J}^{asy} = \{\{1\}, \{2\}, \dots, \{N\}\}$, and a payoff profile $w \in \text{Payoffs}^F(\langle \Delta, \mathcal{J}^{asy}, \mathcal{P}^{asy}, u, p, q, \rho \rangle)$ such that $\|v - w\| < \epsilon$. In words, focusing on strict subgame perfect equilibria in finite memory strategies, an equilibrium payoff profile of the game with simultaneous moves is approachable if there exists a nearby equilibrium payoff profile of the game for some asynchronous protocol of moves.

The following proposition shows that an approachable equilibrium payoff profile of the game with simultaneous moves almost coincides with a payoff profile corresponding to a Markov perfect equilibrium provided that periods are sufficiently short:

Proposition 3 *Fix p , q , and ρ . For almost all u , and all $\epsilon > 0$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, if $v \in \text{Payoffs}^F(\Gamma^{sim})$ is approachable, then there exists $w \in \text{Payoffs}^M(\Gamma^{sim})$ such that $\|v - w\| < \epsilon$.*

Proposition 3 implies that there is no loss in restricting attention to Markov perfect equilibria in separable dynamic games with noisy transitions and simultaneous moves and thus provides a rationale for doing so.

To prove Proposition 3, we build on related results for dynamic stochastic games with asynchronous moves by Bhaskar and Vega-Redondo (2002) and Bhaskar, Mailath, and Morris (2009) and combine them with our Theorem 1. The proof of Proposition 3 draws on the insight from Example 8 that although some payoff profiles can be attained with strict subgame perfect equilibria in finite memory strategies when moves are simultaneous, these payoff profiles cannot be attained when moves are alternating.

Approachability is conceptually similar to purifiability in Bhaskar, Mailath, and Morris (2013) in that both are robustness requirements: approachability says that equilibrium payoffs should survive changes in the protocol of moves, whereas purifiability says that equilibrium strategies should survive the introduction of private information. We show that only Markov perfect equilibria are approachable in our separable dynamic games with noisy transitions and simultaneous moves, whereas Bhaskar, Mailath, and Morris (2013) show that only Markov perfect equilibria are purifiable in dynamic stochastic games with asynchronous moves.

Proposition 3 also limits possible extensions of Theorem 1. By showing that an equilibrium payoff that is robust to alternative specifications of the protocol of moves must be a Markov perfect equilibrium payoff, Proposition 3 implies that the assumption of Markov perfection is not only sufficient (as shown in Theorem 1) but also necessary for protocol invariance.³¹

4.3 Computation of Markov Perfect Equilibria

Dynamic stochastic games are often not very tractable analytically and thus call for the use of numerical methods. Our main result has a number of implications for computing Markov perfect

³¹Note, however, that Proposition 3 applies only when strategies have finite memory. In the tightly specified model in Example 8, under arbitrary protocols of moves there exists a subgame perfect equilibrium with unbounded memory in which $a_i^t = 1$ for all i and after all on-path histories. Without restrictions on strategies, the properties of the set of equilibrium payoffs as $\Delta \rightarrow 0$ are generally not well understood in the literature. The existing results consider either the limit $\Delta \rightarrow 0$ with simultaneous moves (Peski and Wiseman 2015) or the limit $\rho \rightarrow 0$ (Dutta 1995, Yoon 2001, Hörner, Sugaya, Takahashi, and Vieille 2011).

equilibria.

First, Doraszelski and Judd (2018) show that the computational burden can vary by orders of magnitude with the protocol of moves. For the class of separable dynamic games with noisy transitions, Theorem 1 justifies imposing the protocol of moves that is most convenient from a computational perspective.

Second, Doraszelski and Judd (2012) contrast the burden of computing Markov perfect equilibria in discrete- and continuous-time stochastic games with simultaneous moves. They argue that, under widely used laws of motion for the evolution of the state, computing the expectation over successor states ω' in a continuous-time stochastic game does not suffer from the curse of dimensionality that plagues the discrete-time stochastic game, and that this can reduce the computational burden by orders of magnitude. While passing from discrete to continuous time is computationally advantageous, a natural question is if this changes the nature of the strategic interactions among players. Theorem 1 answers this question for the class of separable dynamic games with noisy transitions. Moreover, the techniques we develop allow us to more broadly establish a tight link between discrete- and continuous-time stochastic games even in the absence of Assumption 3.

Consider a dynamic stochastic game with noisy transitions and simultaneous moves. The per-period payoff is $u_i^\Delta(\omega, \{1, \dots, N\}, a) = u_i(\omega, a)\Delta + \mathcal{O}(\Delta^2)$. The probability that the state changes is $q(\omega, a)\Delta$ in proportion to the length of a period Δ and, conditional on the state changing, the probability that it changes from ω to ω' is $p(\omega'|\omega, a)$. Hence, while Assumptions 1 and 2 are satisfied, Assumption 3 is not. Overloading notation, let $\text{Equil}(\langle\Delta, u, p, q, \rho\rangle)$ be the set of Markov perfect equilibria of this game and $\text{Equil}^0(\langle u, p, q, \rho\rangle)$ the set of solutions to the analog of the limit conditions (3.2) and (3.3).³²

Proposition 4 *Fix $p, q,$ and ρ . For almost all u , $\lim_{\Delta \rightarrow 0} \text{Equil}(\langle\Delta, u, p, q, \rho\rangle) = \text{Equil}^0(\langle u, p, q, \rho\rangle)$.*

In words, provided that periods are sufficiently short the Markov perfect equilibria of the discrete-time stochastic game with simultaneous moves almost coincide with those of the continuous-time stochastic game, although the latter are much easier to compute than the former. We note that Proposition 4 does not carry over from simultaneous to alternating moves. We also note that

³²Given a sequence (A_ν) indexed by $\nu \in \mathbb{N}$, with $A_\nu \subseteq \mathbb{R}^n$, we define

$$\liminf_{\nu \rightarrow \infty} A_\nu = \{x \in \mathbb{R}^n \mid \limsup_{\nu \rightarrow \infty} d(x, A_\nu) = 0\} \quad \text{and} \quad \limsup_{\nu \rightarrow \infty} A_\nu = \{x \in \mathbb{R}^n \mid \liminf_{\nu \rightarrow \infty} d(x, A_\nu) = 0\},$$

where $d(x, A) = \inf \{\|y - x\| \mid y \in A\}$. If both limits coincide, we denote their common value by $\lim_{\nu \rightarrow \infty} A_\nu$.

with a continuum of actions, a version of Proposition 4 (and of Theorem 1) can be obtained by considering approximate equilibria in pure strategies as in Fudenberg and Levine (1986).

5 Conclusions

The timing of decisions is an essential ingredient into modeling many strategic situations. Yet, determining the protocol of moves that is most realistic and appropriate for the application at hand can be challenging. While the literature abounds with examples where the protocol of moves matters crucially for equilibrium behavior, our paper is a first attempt to show that the implications and predictions of a fairly general and widely used class of dynamic models are independent of the timing of decisions and thus more robust for the purposes of applied work.

We introduce separable dynamic games with noisy transitions and establish that they are protocol invariant provided that periods are sufficiently short and moves are therefore sufficiently frequent. Separable dynamic games with noisy transitions are well-suited for situations in which a player primarily influences his rivals' payoffs by taking action to change the state and there is some residual uncertainty if the taken action brings about such a change. A particular highlight of this class of dynamic stochastic games is that per-period payoffs and state-to-state transitions can depend arbitrarily on the state. We show that investment games, R&D races, models of industry dynamics, dynamic public contribution games, the recent continuous-time stochastic games with moves at random times, asynchronously repeated games, and many other models from the extant literature can be cast as special cases of separable dynamic games with noisy transitions.

In addition to alleviating the burden of determining the most realistic and appropriate protocol of moves, our main result and its extensions have a number of implications for applied work. They provide a new justification for focusing on Markov perfect equilibria in dynamic stochastic games and facilitate computing these equilibria. They further clarify a driving force behind some of the well-known and important examples in the literature where equilibrium behavior hinges on the protocol of moves. Many of these models, including entry games and the asynchronously repeated Bertrand, Cournot, and coordination games in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997), are separable but transitions are assumed to be deterministic. If hazard rates are instead finite, the protocol of moves ceases to matter.

Our paper opens up several questions. First, extending our protocol-invariance theorem to games with a continuum of states that follow a jointly-controlled diffusion process would be useful,

e.g., in finance applications. Yet, a continuum of states brings new challenges to ensure equilibrium existence and, outside special cases such as Abreu and Gul (2000) and Caruana and Einav (2008), to establish the local uniqueness and continuity of solutions to the limit conditions (3.2) and (3.3) using differential topology tools (as in Lemma 2). Second, with Theorem 1 we provide sufficient conditions for protocol invariance. While our examples show that protocol invariance may fail if any of our assumptions is relaxed, a characterization of the class of dynamic stochastic games that are robust to the timing of decisions would be informative. Deriving necessary conditions for protocol invariance requires formulating a model that encompasses a wider class of games under various protocols of moves. We leave these questions for future research.

A Appendix: Proofs

Appendix A.1 provides the proof of Theorem 1 and Appendix A.2 provides the proofs for Section 4.

A.1 Proof of Theorem 1

A.1.1 Notation and Preliminary Definitions

Enumerate the state space as $\Omega = \{\omega^1, \dots, \omega^{|\Omega|}\}$ and the set of actions for player i as $A_i(\omega) = \{a_i^1, \dots, a_i^{|A_i(\omega)|}\}$. Given a strategy profile $\sigma = (\sigma_i)_{i=1}^N \in \Sigma$, define the matrix $\mathcal{P}_\sigma \in \mathbb{R}^{|\Omega| \times \sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)|}$ as

$$\begin{pmatrix} \sigma_1(a_1^1 | \omega^1) & \dots & \sigma_1(a_1^{|A_1(\omega)|} | \omega^1) & \sigma_2(a_2^1 | \omega^1) & \dots & \sigma_N(a_N^{|A_N(\omega)|} | \omega^1) & 0 & \dots & 0 & \dots \\ 0 & \dots & \dots & \dots & \dots & \dots & \sigma_1(a_1^1 | \omega^2) & \dots & \sigma_N(a_N^{|A_N(\omega)|} | \omega^2) & \dots \\ 0 & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots & \dots \end{pmatrix}.$$

Define the matrix $Q \in \mathbb{R}^{\sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)| \times |\Omega|}$ as

$$\begin{pmatrix} q_1(a_1, \omega^1) & 0 & \dots & \\ \vdots & & & \\ q_1(a_1^{|A_1(\omega)|}, \omega^1) & 0 & \dots & \\ \vdots & & & \\ q_N(a_N^{|A_N(\omega)|}, \omega^1) & 0 & & \\ 0 & q_1(a_1, \omega^2) & 0 & \dots \\ \vdots & & \vdots & \end{pmatrix}$$

and the matrix $P \in \mathbb{R}^{\sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)| \times |\Omega|}$ as

$$P_{(i, a_i, \omega), \omega'} = \begin{cases} \varphi_i(\omega' | a_i, \omega) & \text{if } \omega' \neq \omega, \\ 0 & \text{if } \omega' = \omega. \end{cases}$$

Given a player $i \in \{1, \dots, N\}$, limit condition (3.2) can be written as

$$\left(\rho \mathbb{1} + \mathcal{P}_\sigma(Q - P) \right) V_i^0 = \mathcal{P}_\sigma u_i,$$

where $\mathbb{1}$ is the identity matrix, $V_i^0 \in \mathbb{R}^{|\Omega|}$, and $u_i \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$. The matrix $\rho \mathbb{1} + \mathcal{P}_\sigma(Q - P)$ is

strictly dominant diagonal and therefore invertible.³³ We emphasize the dependence of the unique solution to limit condition (3.2) by writing $V_i^0(\cdot) = V_i^0(\cdot, \sigma)$. This solution is

$$V_i^0(\cdot, \sigma) = \left(\rho \mathbb{1} + \mathcal{P}_\sigma(Q - P) \right)^{-1} \mathcal{P}_\sigma u_i.$$

Given $i \in \{1, \dots, N\}$ and $u_i \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$, consider the vector

$$\begin{aligned} u_i + (P - Q)V_i^0(\cdot, \sigma) &= u_i + (P - Q) \left(\rho \mathbb{1} + \mathcal{P}_\sigma(Q - P) \right)^{-1} \mathcal{P}_\sigma u_i \\ &= u_i + (P - Q) \frac{1}{\rho} \left(\mathbb{1} - \frac{1}{\rho} \mathcal{P}_\sigma(Q - P) + \frac{1}{\rho^2} (\mathcal{P}_\sigma(Q - P))^2 \dots \right) \mathcal{P}_\sigma u_i \\ &= \left(\mathbb{1} - \frac{1}{\rho} (Q - P) \mathcal{P}_\sigma + \frac{1}{\rho^2} ((Q - P) \mathcal{P}_\sigma)^2 - \frac{1}{\rho^3} ((Q - P) \mathcal{P}_\sigma)^3 + \dots \right) u_i \\ &= \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} u_i, \end{aligned}$$

where the inversion is justified by strict diagonal dominance. The map

$$u_i \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|} \mapsto \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} u_i \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$$

is invertible.

The above results have been presented for a given strategy profile $\sigma \in \Sigma$. Following Appendix A.1 in Doraszelski and Escobar (2010), we construct an open set $\Sigma^\varepsilon \subset \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$ that strictly contains Σ such that all the preceding operations are valid for any $\sigma \in \Sigma^\varepsilon$.

A.1.2 Regularity

We begin by providing a formal definition of regularity and establishing the key technical point that for almost all flow payoffs $u \in \mathbb{R}^{\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$, the restriction to regular solutions is without loss of generality.

Given $i \in \{1, \dots, N\}$, $\omega \in \Omega$, $a_i \in A_i(\omega)$, and $\sigma \in \Sigma^\varepsilon$, define the function

$$\mathcal{U}_i(\omega, a_i, \sigma) = u_{i,i}(\omega, a_i) + \sum_{\omega' \neq \omega} (V_i^0(\omega', \sigma) - V_i^0(\omega, \sigma)) \varphi_i(\omega' \mid \omega, a_i).$$

In light of limit condition (3.3), we interpret $\mathcal{U}_i(a_i, \omega, \sigma)$ as the objective function that player

³³A strictly dominant diagonal matrix X is a square matrix with entries X_{ij} such that $|X_{ii}| > \sum_{j \neq i} |X_{ij}|$ for all i .

$i \in \{1, \dots, N\}$ maximizes over $a_i \in A_i(\omega)$ given state $\omega \in \Omega$ and continuation play $\sigma \in \Sigma^\varepsilon$.

Consider $\bar{\sigma} \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$. Choose a_i^ω such that $\bar{\sigma}_i(a_i^\omega | \omega) > 0$ for all $i \in \{1, \dots, N\}$ and all $\omega \in \Omega$. Given $a_i \neq a_i^\omega$ and $\sigma \in \Sigma^\varepsilon$, define the continuously differentiable function $f: \Sigma^\varepsilon \times \mathbb{R}^N \sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)| \rightarrow \mathbb{R} \sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)|$ such that its (i, a_i, ω) component is

$$f_{i, a_i, \omega}(\sigma) = \sigma_i(a_i | \omega) \left(\mathcal{U}_i(a_i, \omega, \sigma) - \mathcal{U}_i(a_i^\omega, \omega, \sigma) \right)$$

while

$$f_{i, a_i^\omega, \omega}(\sigma) = \sum_{a_i \in A_i(\omega)} \sigma_i(a_i | \omega) - 1.$$

By definition, $f(\bar{\sigma}) = 0$. In what follows, we sometimes emphasize the dependence of f on u by writing $f(\sigma, u)$.

Definition 1 $\bar{\sigma} \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ is regular if the Jacobian of f with respect to σ , $\frac{\partial f}{\partial \sigma}(\bar{\sigma})$, has full rank $\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$.

We present two preliminary lemmas. We say that a strategy profile $\sigma \in \Sigma^\varepsilon$ is *completely mixed* if $\sigma_i(a_i | \omega) > 0$ for all $i \in \{1, \dots, N\}$, $\omega \in \Omega$, and all $a_i \in A_i(\omega)$.

Lemma 3 If $\sigma \in \Sigma^\varepsilon$ is completely mixed, then the Jacobian of f with respect to (σ, u) , $\frac{\partial f}{\partial(\sigma, u)}(\sigma, u)$, has full rank $\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$.

Proof. Define the matrix $M(\sigma, i) \in \mathbb{R}^{\sum_{\omega \in \Omega} (|A_i(\omega)| - 1) \times \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$ such that, for all $a_i \neq \bar{a}_i^\omega$, its (i, a_i, ω) row equals 0 in all components save for the (i, a_i, ω) column, where we write $\sigma_i(a_i | \omega)$, and for the $(i, \bar{a}_i^\omega, \omega)$ column, where we write $-\sigma_i(a_i | \omega)$. The function f can be expressed as

$$f_i(\sigma, u) = \begin{cases} \sum_{a_i \in A_i(\omega)} \sigma_i(a_i | \omega^1) - 1, \\ \vdots \\ \sum_{a_i \in A_i(\omega)} \sigma_i(a_i | \omega^{|\Omega|}) - 1, \\ M(\sigma, i) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} u_i. \end{cases}$$

Up to permutation (which are irrelevant to determine the rank of the Jacobian), we can write

$$\frac{\partial f(\sigma, u)}{\partial(\sigma, u)} = \left(\begin{array}{c|c|c|c|c|c|c|c} \sigma_1 & \sigma_2 & \dots & \sigma_N & u_1 & u_2 & \dots & u_N \\ \hline X_1 & 0 & 0 & 0 & 0 & 0 & \dots & 0 \\ 0 & X_2 & 0 & 0 & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & X_N & 0 & 0 & \dots & 0 \\ \hline & & & & Z_1 & 0 & \dots & 0 \\ & & & & 0 & Z_2 & \dots & 0 \\ Y_1 & Y_2 & \dots & Y_N & \vdots & \vdots & \ddots & \vdots \\ & & & & 0 & 0 & \dots & Z_N \end{array} \right),$$

where X_i equals

$$\left(\begin{array}{c|c|c|c} \sigma_i(\cdot | \omega^1) & \sigma_i(\cdot | \omega^2) & \dots & \sigma_i(\cdot | \omega^{|\Omega|}) \\ \hline 1 & \dots & 1 & 0 & \dots & 0 & \dots & 0 & \dots & 0 \\ 0 & \dots & 0 & 1 & \dots & 1 & \dots & 0 & \dots & 0 \\ \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 & \dots & 1 & \dots & 1 \end{array} \right)$$

and has rank $|\Omega|$, while $Z_i = M(\sigma, i) \left(\mathbb{1} + \frac{1}{\rho}(P-Q)\mathcal{P}_\sigma \right)^{-1}$. Since $M(\sigma, i)$ has full rank $\sum_{\omega \in \Omega} (|A_i(\omega)| - 1)$ and $(\mathbb{1} + \frac{1}{\rho}(P-Q)\mathcal{P}_\sigma)^{-1}$ has full rank $\sum_{\omega \in \Omega} \sum_{j=1}^N |A_j(\omega)|$, Z_i has rank $\sum_{\omega \in \Omega} (|A_i(\omega)| - 1)$. We deduce that the Jacobian has full rank $\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$ and the lemma follows. ■

Given $\sigma \in \Sigma$, $i \in \{1, \dots, N\}$, and $\omega \in \Omega$, define the *best reply* as

$$B_i(\sigma, \omega) = \arg \max_{a_i \in A_i(\omega)} \mathcal{U}_i(\omega, a_i, \sigma)$$

and the *carrier* as

$$C_i(\sigma, \omega) = \left\{ a_i \in A_i(\omega) \mid \sigma_i(a_i | \omega) > 0 \right\}.$$

Using this notation, $\sigma \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ if and only if $C_i(\sigma, \omega) \subseteq B_i(\sigma, \omega)$ for all $i \in \{1, \dots, N\}$.

We say that $\sigma \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ is *quasi-strict* if $C_i(\sigma, \omega) = B_i(\sigma, \omega)$ for all $i \in \{1, \dots, N\}$.

Lemma 4 *For almost all $u \in \mathbb{R}^N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$, any $\sigma \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ is quasi-strict.*

Proof. Given $i \in \{1, \dots, N\}$, consider correspondences $B_i^*: \Omega \rightarrow \cup_{\omega \in \Omega} A_i(\omega)$ and $C_i^*: \Omega \rightarrow \cup_{\omega \in \Omega} A_i(\omega)$, with $C_i^*(\omega) \subseteq B_i^*(\omega) \subseteq A_i(\omega)$ for all $\omega \in \Omega$. Define $G(B^*, C^*)$ as the set of all u having some $\sigma \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ with best replies $B^* = (B_i^*)_{i=1}^N$ and carriers $C^* = (C_i^*)_{i=1}^N$. Formally,

$$G(B^*, C^*) = \left\{ u \mid \exists \sigma \in \text{Equil}^0(\langle u, p, q, \rho \rangle) \text{ with } B_i(\sigma, \cdot) = B_i^* \text{ and } C_i(\sigma, \cdot) = C_i^* \text{ for all } i = 1, \dots, N \right\}.$$

Consider first $\bar{\sigma} \in \text{Equil}^0(\langle \bar{u}, p, q, \rho \rangle)$ such that $B_i(\bar{\sigma}, \omega) = B_i^*(\omega)$ for all $\omega \in \Omega$. Fix a_i^ω such that $\bar{\sigma}_i(a_i^\omega \mid \omega) > 0$ and note that the indifference condition $\mathcal{U}_i(a_i, \omega, \bar{\sigma}) - \mathcal{U}_i(a_i^\omega, \omega, \bar{\sigma}) = 0$ holds for all $i \in \{1, \dots, N\}$ and all $a_i \in B_i^*(\omega)$. For all $\omega \in \Omega$ and all $i \in \{1, \dots, N\}$, define the matrix $P_i(\sigma) \in \mathbb{R}^{\sum_{\omega \in \Omega} (|B_i^*(\omega)| - 1) \times \sum_{\omega \in \Omega} |A_i(\omega)|}$, such that for all $a_i \in B_i^*(\omega)$, its (ω, a_i) row equals 0 save for the (ω, a_i) component, where it equals 1, and the (ω, a_i^ω) component, where it equals -1. We can therefore stack all the indifference conditions by writing

$$M(\sigma, u) = \begin{pmatrix} P_1(\sigma) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} u_1 \\ \vdots \\ P_N(\sigma) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} u_N \end{pmatrix}$$

and note that $M(\bar{\sigma}, \bar{u}) = 0$. The Jacobian $\frac{\partial M}{\partial u}(\sigma, u)$ can be computed as

$$\frac{\partial M}{\partial u}(\sigma, u) = \begin{pmatrix} P_1(\sigma) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} & 0 & \dots & 0 \\ 0 & P_2(\sigma) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & \dots & P_N(\sigma) \left(\mathbb{1} + \frac{1}{\rho} (P - Q) \mathcal{P}_\sigma \right)^{-1} \end{pmatrix}.$$

Since $P_i(\sigma)$ has full rank $\sum_{\omega \in \Omega} (|B_i^*(\omega)| - 1)$, the Jacobian $\frac{\partial M}{\partial u}(\sigma, u)$ has rank $\sum_{i=1}^N \sum_{\omega \in \Omega} (|B_i^*(\omega)| - 1)$. In particular, since $M(\bar{\sigma}, \bar{u}) = 0$, we can construct open sets $\mathcal{N}, \mathcal{N}_1 \subseteq \mathbb{R}^{\sum_{i=1}^N \sum_{\omega \in \Omega} (N|A_i(\omega)| - |B_i^*(\omega)| + 1)}$, $\mathcal{N}_2 \subseteq \mathbb{R}^{\sum_{i=1}^N \sum_{\omega \in \Omega} (|B_i^*(\omega)| - 1)}$, with $\bar{\sigma} \in \mathcal{N}$ and $\bar{u} \in \mathcal{N}_1 \times \mathcal{N}_2$, and a continuously differentiable function Φ such that for all $(\sigma, u_1) \in \mathcal{N} \times \mathcal{N}_1$ there exists a unique $u_2 = \Phi(\sigma, u_1) \in \mathcal{N}_2$ which is a solution to $M(\sigma, (u_1, u_2)) = 0$. Without loss, all these open sets are balls with rational centers and radii and we emphasize their dependence on $(\bar{\sigma}, \bar{u})$ by writing $\mathcal{N}_1^{\bar{\sigma}, \bar{u}}, \mathcal{N}_2^{\bar{\sigma}, \bar{u}}$, and $\mathcal{N}^{\bar{\sigma}, \bar{u}}$.

Now take C^* such that for some $i \in \{1, \dots, N\}$ and some $\omega \in \Omega$, $C_i^*(\omega) \subsetneq B_i^*(\omega)$. Consider the

set

$$R^{\bar{\sigma}, \bar{u}}(B^*, C^*) = \left\{ u \in \mathcal{N}_1^{\bar{\sigma}, \bar{u}} \times \mathcal{N}_2^{\bar{\sigma}, \bar{u}} \mid \text{there exists } (\sigma, u^1) \in (\mathcal{N}^{\bar{\sigma}} \cap A(C^*)) \times \mathcal{N}_1^{\bar{\sigma}, \bar{u}} \text{ such that } u_2 = \Phi(\sigma, u_1) \right\} \\ \subseteq \mathbb{R}^{N \sum_{\omega \in \Omega} \sum_{j=1}^N |A_j(\omega)|},$$

where $A(C^*) = \{\sigma \in \Sigma \mid C_i(\cdot, \omega) = C_i^*(\omega) \text{ for all } i = 1, \dots, N\}$. Note that the dimension of $(\mathcal{N}^{\bar{\sigma}} \cap A(C^*)) \times \mathcal{N}_1^{\bar{\sigma}, \bar{u}}$ equals $N \sum_{\omega \in \Omega} \sum_{j=1}^N |A_j(\omega)| - \sum_{i=1}^N \sum_{\omega \in \Omega} |B_i^*(\omega)| + \sum_{i=1}^N \sum_{\omega \in \Omega} |C_i^*(\omega)| < N \sum_{\omega \in \Omega} \sum_{j=1}^N |A_j(\omega)|$. Therefore, $M^{\bar{\sigma}, \bar{u}}(B^*, C^*)$ has measure zero. Since we are choosing the neighborhoods from a countable set, it follows that $G(B^*, C^*) \subseteq \cup_{n \in \mathbb{N}} Q^n$, where $Q^n = R^{\bar{\sigma}_n, \bar{u}_n}(B^*, C^*)$, has measure zero as well. ■

The following is the main result of this subsection.

Proposition 5 *For almost all $u \in \mathbb{R}^{N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$, all $\bar{\sigma} \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ are regular.*

Proof. From Lemma 4, we can rule out games $u \in \mathbb{R}^{N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$ having non quasi-strict solutions and focus on games having only quasi-strict solutions. Since there is a finite number of correspondences $B_i^*: \Omega \rightarrow A_i$, it is enough to prove that the set of games having a non-regular equilibrium σ with

$$B_i(\sigma, \omega) = C_i(\sigma, \omega) = B_i^*(\omega)$$

for all $i \in \{1, \dots, N\}$ and all $\omega \in \Omega$ has measure zero. Considering the submatrix $\bar{J}(\sigma, u)$ obtained from $\frac{\partial f}{\partial \sigma}(\sigma, u)$ by crossing out all rows and columns corresponding to components (a_i, ω) with $a_i \notin B_i^*(\omega)$, it follows that $\bar{J}(\sigma)$ has full rank if and only if so does $\frac{\partial f}{\partial \sigma}(\sigma, u)$. Noting that $\bar{J}(\sigma, u)$ is the Jacobian of a completely mixed solution, without loss of generality we can therefore assume that $B^*(\omega)$ does not depend on ω and restrict attention to completely mixed solutions. Using Lemma 3 and the transversality theorem, we deduce that for almost all games, all completely mixed equilibria are regular. ■

A.1.3 Establishing Lemma 2

Fix a game $u \in \mathbb{R}^{N \sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)|}$ and a regular solution $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$. Let $\langle \mathcal{J}, \mathcal{P} \rangle$ be a protocol of moves and $\Delta > 0$ the period length. We establish that the regular solution σ^0 can be approximated by a Markov perfect equilibrium of a separable dynamic game with noisy transitions and an arbitrary protocol of moves if the period length Δ is sufficiently small. To do so, we apply a version of the implicit function theorem to the limit conditions.

Proof of Lemma 2. In the separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$, write the continuation value of player $i \in \{1, \dots, N\}$ if players $J \in \mathcal{J}$ have the move and the state is $\omega \in \Omega$ as $V_i^\Delta(\omega, J)$. Note that the value function $V_i^\Delta: \Omega \times \mathcal{J} \rightarrow \mathbb{R}$ is uniquely determined by the strategy profile $\sigma \in \Sigma$. We therefore write $V_i^\Delta(\cdot, \cdot) = V_i^\Delta(\cdot, \cdot, \sigma)$. Note that $V_i^\Delta(\cdot, \cdot, \sigma)$ is a continuous function of (σ, Δ) and its differential with respect to σ at $\Delta = 0$ exists. In particular, for all $J \in \mathcal{J}$ and all $\sigma \in \Sigma$, $V_i^\Delta(\cdot, J, \sigma) \rightarrow V_i^0(\cdot, \sigma)$ as $\Delta \rightarrow 0$.

A strategy profile σ^Δ is a Markov perfect equilibrium of the separable dynamic game with noisy transitions $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$ if for all $i = 1, \dots, N$, $\omega \in \Omega$, and all $a_i \in A_i(\omega)$

$$\sigma_i^\Delta(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} \mathcal{U}_i^\Delta(\omega, \tilde{a}_i, \sigma^\Delta)$$

with

$$\mathcal{U}_i^\Delta(\omega, a_i, \sigma^\Delta) = u_{i,i}(\omega, a_i) + \exp(-\rho\Delta) \sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} \left(V_i^\Delta(\omega', J', \sigma^\Delta) - V_i^\Delta(\omega, J', \sigma^\Delta) \right) \varphi_i(\omega' | \omega, a_i) \Pr(J' | J) + \mathcal{O}(\Delta)$$

and $J \in \mathcal{J}$ is such that $i \in J$.

Consider the profile $(a_i^\omega)_{i=1, \dots, N, \omega \in \Omega}$ that is used in the construction of the function f in Appendix A.1.2 for which σ^0 is regular. Abusing notation, construct the function $f: [0, 1] \times \Sigma^\varepsilon \rightarrow \mathbb{R}^{\sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)|}$ such that for all $a_i \neq a_i^\omega$

$$f_{i, a_i, \omega}(\Delta, \sigma) = \sigma_i(a_i, \omega) \left(\mathcal{U}_i^\Delta(\omega, a_i, \sigma) - \mathcal{U}_i^\Delta(\omega, a_i^\omega, \sigma) \right)$$

while

$$f_{i, a_i^\omega, \omega}(\Delta, \sigma) = \sum_{a_i \in A_i(\omega)} \sigma_i(a_i | \omega) - 1.$$

Observe that $f(\Delta, \sigma)$ is a continuous function, with a well-defined differential with respect to σ , $D_\sigma f(0, \sigma)$, at $(0, \sigma)$. Moreover, $f(0, \sigma^0) = 0$ and $D_\sigma f(0, \sigma^0)$ has full rank $\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$. A version of the implicit function theorem (see Lemma 5 below) implies that for all $r > 0$ there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, there exists $\sigma^\Delta \in \Sigma^\varepsilon$ with $\|\sigma^0 - \sigma^\Delta\| < r$ such that $f(\Delta, \sigma^\Delta) = 0$. Moreover, we can take $\bar{\Delta}$ and r small enough so that (i) $\sigma_i^\Delta(a_i, \omega) > 0$ whenever $\sigma_i^0(a_i, \omega) > 0$, and (ii) $\mathcal{U}_i^\Delta(\omega, a_i, \sigma^\Delta) < \mathcal{U}_i^\Delta(\omega, a_i^\omega, \sigma^\Delta)$ whenever $\mathcal{U}_i(\omega, a_i, \sigma^0) < \mathcal{U}_i(\omega, a_i^\omega, \sigma^0)$.

To prove that σ^Δ is a Markov perfect equilibrium of the separable dynamic game with noisy transitions $\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$, consider first $a_i \in A_i(\omega)$ and $\omega \in \Omega$ such that $\sigma^0(a_i | \omega) = 0$. Since σ^0 is regular, it is also quasi-strict and therefore $\mathcal{U}_i(\omega, a_i, \sigma^0) < \mathcal{U}_i(\omega, a_i^\omega, \sigma^0)$. From (ii),

$\mathcal{U}_i^\Delta(\omega, a_i, \sigma^\Delta) < \mathcal{U}_i^\Delta(\omega, a_i^\omega, \sigma^\Delta)$. Since $f(\Delta, \sigma^\Delta) = 0$, it follows that $\sigma^\Delta(a_i | \omega) = 0$. Next consider $a_i \in A_i(\omega)$ and $\omega \in \Omega$ such that $\sigma^0(a_i | \omega) > 0$. We can use (i) to deduce that $\sigma^\Delta(a_i | \omega) > 0$ and, since $f(\Delta, \sigma^\Delta) = 0$, $\mathcal{U}_i^\Delta(\omega, a_i, \sigma^\Delta) = \mathcal{U}_i^\Delta(\omega, a_i^\omega, \sigma^\Delta)$. These observations prove that $\sigma_i^\Delta(a_i | \omega) \geq 0$ for all $a_i \in A_i(\omega)$, all $i \in \{1, \dots, N\}$, and all $\omega \in \Omega$. Moreover, whenever $\sigma_i^\Delta(a_i, \omega) > 0$, a_i solves $\max_{\tilde{a}_i \in A_i(\omega)} \mathcal{U}_i^\Delta(\omega, \tilde{a}_i, \sigma^\Delta)$. Since $f(\Delta, \sigma^\Delta) = 0$, it follows that $\sum_{a_i \in A_i(\omega)} \sigma_i^\Delta(a_i | \omega) = 1$ for all $i \in \{1, \dots, N\}$ and all $\omega \in \Omega$. Therefore, $\sigma^\Delta \in \Sigma$ is a Markov perfect equilibrium of the separable dynamic game with noisy transitions $\langle \Delta, \mathcal{J}, \mathcal{P}, u, , q, \rho \rangle$.

Since for almost all $u \in \mathbb{R}^N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$, all $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ are regular, Lemma 2 follows. ■

It remains to prove the implicit function theorem we used above. The textbook presentation of the implicit function theorem (Section M.E in Mas-Colell, Whinston, and Green 1995) applies to continuously differentiable functions defined on open sets. In our setup, the set of parameters $\Delta \in [0, 1]$ is closed and, moreover, we are interested in the boundary case $\Delta = 0$. The following result is a modification of Theorem A in Halkin (1974).

Lemma 5 (Implicit Function Theorem) *Assume $f: [0, 1] \times \Sigma^\varepsilon \rightarrow \mathbb{R}^{\sum_{\omega \in \Omega} \sum_{i=1}^N |A_i(\omega)|}$ is a continuous function such that its differential with respect to $\sigma \in \Sigma^\varepsilon$ at $\Delta = 0$, $D_\sigma f(0, \sigma)$, exists. Let $\sigma^0 \in \Sigma$ be such that $f(0, \sigma^0) = 0$ and $D_\sigma f(0, \sigma^0)$ has full rank $\sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|$. Then, for all $r > 0$ there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, there exists σ^Δ such that $\|\sigma^0 - \sigma^\Delta\| < r$ and $f(\Delta, \sigma^\Delta) = 0$.*

Proof. Consider the function $\phi(\sigma, \Delta) = \sigma - [D_\sigma f(0, \sigma^0)]^{-1} f(\Delta, \sigma)$ and note that the problem of finding σ^Δ such that $f(\Delta, \sigma^\Delta) = 0$ reduces to the problem of finding a fixed point of $\phi(\cdot, \Delta)$. Note that $D_\sigma \phi(0, \sigma^0) = 0$ and therefore we can assume, without loss, that $r > 0$ is small enough so that for all $\|\sigma - \sigma^0\| < r$, $\sigma \in \Sigma^\varepsilon$ and

$$\frac{\|\phi(\sigma, 0) - \phi(\sigma^0, 0)\|}{\|\sigma - \sigma^0\|} < \frac{1}{2}.$$

Since $\phi(\sigma^0, 0) = \sigma^0$, we can therefore deduce that for all $\|\sigma - \sigma^0\| \leq r$, $\|\phi(\sigma, 0) - \sigma^0\| \leq r/2$.

Define now $m(\Delta) = \max_{\{\sigma \mid \|\sigma - \sigma^0\| \leq r\}} \|\phi(\sigma, \Delta) - \phi(\sigma, 0)\|$. Berge's maximum theorem (Theorem 17.31 in Aliprantis and Border 2006) implies that m is continuous in $\Delta \in [0, 1]$. Since $m(0) = 0$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, $m(\Delta) < r/2$. We thus deduce that for all σ such that

$$\|\sigma - \sigma^0\| \leq r \text{ and } \Delta < \bar{\Delta}$$

$$\begin{aligned} \|\phi(\sigma, \Delta) - \sigma^0\| &\leq \|\phi(\sigma, \Delta) - \phi(\sigma, 0)\| + \|\phi(\sigma, 0) - \sigma^0\| \\ &\leq m(\Delta) + \|\phi(\sigma, \Delta) - \sigma^0\| \\ &\leq r. \end{aligned}$$

It follows that for all $\Delta < \bar{\Delta}$, the continuous function $\phi(\cdot, \Delta)$ maps the convex and compact set $\{\sigma \mid \|\sigma - \sigma^0\| \leq r\}$ into itself. For any such $\Delta < \bar{\Delta}$, Brouwer's fixed point theorem (Theorem M.I.1 in Mas-Colell, Whinston, and Green 1995) implies the existence of σ^Δ within distance r of σ^0 such that $\phi(\sigma^\Delta, \Delta) = \sigma^\Delta$. ■

A.1.4 Proof of Theorem 1

Proof of Theorem 1. From Lemma 2, take one of the generic flow payoffs $u \in \mathbb{R}^N \sum_{\omega \in \Omega} \sum_{j=1}^N |A_j(\omega)|$ and any two protocols of moves as in the statement of Theorem 1. From Lemma 1, there exists $\tilde{\Delta} > 0$ such that for all $\Delta < \tilde{\Delta}$ and all $\sigma^\Delta \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$, there exists $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ such that $\|\sigma^\Delta - \sigma^0\| < \varepsilon/2$. From Lemma 2 we can find $\hat{\Delta} > 0$ such that for all $\Delta < \hat{\Delta}$, there exists $\hat{\sigma}^\Delta \in \text{Equil}(\langle \Delta, \bar{\mathcal{J}}, \bar{\mathcal{P}}, u, p, q, \rho \rangle)$ such that $\|\hat{\sigma}^\Delta - \sigma^0\| < \varepsilon/2$. Taking $\bar{\Delta} = \min\{\tilde{\Delta}, \hat{\Delta}\} > 0$, Theorem 1 follows from the triangle inequality. ■

A.2 Proofs for Section 4

Proof of Proposition 3. Take any u as in Theorem 1. Define $\text{Payoffs}^0(\langle u, p, q, \rho \rangle) \subseteq \mathbb{R}^N$ to be the set of payoff profiles associated with solutions in pure strategy profiles to the limit conditions (3.2) and (3.3):

$$\begin{aligned} \text{Payoffs}^0(\langle u, p, q, \rho \rangle) = \\ \left\{ V^0(\omega^{t=0}) \in \mathbb{R}^N \mid (V^0, \sigma^0) \text{ solves (3.2) and (3.3) for some pure strategy profile } \sigma^0 \right\}, \end{aligned}$$

where $\omega^{t=0} \in \Omega$ is the initial state of the game. For $\varepsilon > 0$ take $\bar{\Delta}$ such that for all $\Delta < \bar{\Delta}$, and all protocols $\langle \mathcal{J}, \mathcal{P} \rangle$, the Hausdorff distance between $\text{Payoffs}^0(\langle u, p, q, \rho \rangle)$ and $\text{Payoffs}^F(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$ is less than $\varepsilon/3$.

Take $v \in \text{Payoffs}^F(\Gamma^{sim})$. Because v is approachable, for all $n \geq 1$, there exists an asynchronous protocol $\langle \mathcal{J}^n, \mathcal{P}^n \rangle$ and $w^n \in \text{Payoffs}^F(\langle \Delta, \mathcal{J}^n, \mathcal{P}^n, u, p, q, \rho \rangle)$ such that $\|v - w^n\| < 1/n$. Restrict

the sequence such that $1/n < \epsilon/3$. From Bhaskar, Mailath, and Morris (2013), we can actually take $w^n \in \text{Payoffs}^M(\langle \Delta, \mathcal{J}^n, \mathcal{P}^n, u, p, q, \rho \rangle)$. By construction, for any such w^n we can find $\tilde{w}^n \in \text{Payoffs}^0(\langle u, p, q, \rho \rangle)$ such that $\|w^n - \tilde{w}^n\| < \epsilon/3$. Since $\text{Payoffs}^0(\langle u, p, q, \rho \rangle)$ has a finite number of elements, we can assume that $\tilde{w}^n = \tilde{w}$ does not depend on n (perhaps, by taking a subsequence). Now, take $w \in \text{Payoffs}^M(\langle \Delta, \mathcal{J}^{sim}, \mathcal{P}^{sim}, u, p, q, \rho \rangle)$ such that $\|w - \tilde{w}\| < \epsilon/3$. It follows that

$$\|v - w\| \leq \|v - w^n\| + \|w^n - \tilde{w}\| + \|\tilde{w} - w\| < \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3},$$

which proves the result. ■

Proof Sketch of Proposition 4. The proof follows from the analysis in Appendix A.1 and arguments in Doraszelski and Escobar (2010). Details are available upon request. To provide a sketch, consider the analog to the limit conditions (3.2) and (3.3) that arise without Assumption 3:

$$\rho V_i(\omega) = u_i(\omega, \sigma(\omega)) + \sum_{\omega' \neq \omega} (V_i(\omega') - V_i(\omega)) \varphi(\omega' \mid \omega, \sigma(\omega))$$

and

$$\sigma_i(a_i \mid \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_i(\omega, \tilde{a}_i, \sigma_{-i}(\omega)) + \sum_{\omega' \neq \omega} (V_i(\omega') - V_i(\omega)) \varphi(\omega' \mid \omega, \tilde{a}_i, \sigma_{-i}(\omega)).$$

From these limit conditions, we can construct a function f (as we did in Appendix A.1) such that all solutions are zeros of f and, moreover, for almost all flow payoffs u , all solutions are regular. We then apply the implicit function theorem to deduce the result. ■

B Appendix: Examples and Extensions

Appendices B.1 and B.2 generalize our notion of a protocol of moves and provide extensions of Theorem 1, Appendix B.3 presents a slightly modified version of the entry game in Section 3, and Appendix B.4 provides a counterexample complementing Example 6.

B.1 Generalized Protocol of Moves

We relax Assumption 1 by generalizing our notion of a protocol of moves. We allow the evolution of the protocol state J to depend on the physical state ω . We maintain that \mathcal{J} is a partition of

the set of players, but allow for a non-uniform stationary distribution. We show that Theorem 1 remains valid.

Assumption 4 (Generalized Protocol of Moves) *Let \mathcal{J} be a partition of $\{1, 2, \dots, N\}$ and $\mathcal{P} = (\Pr(J'|J, \omega))_{J, J' \in \mathcal{J}}$ a $|\mathcal{J}| \times |\mathcal{J}|$ transition matrix for all $\omega \in \Omega$. Assume that $\mathcal{P}(\cdot | \cdot, \omega)$ is irreducible for all $\omega \in \Omega$ and its unique stationary distribution $\pi = (\pi(J))_{J \in \mathcal{J}} \in \mathcal{P}(\mathcal{J})$ is independent of ω .*

We call $\langle \mathcal{J}, \mathcal{P} \rangle$ a *generalized protocol of moves*. Under a generalized protocol of moves, the current physical state may make a transition from one protocol state to another more likely, but on average all protocol states are visited with frequencies that are independent of physical states.³⁴ The remaining aspects of the model are unchanged.

Denote the above game by $\Gamma = \langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle$ and consider a Markov perfect equilibrium $\sigma^\Delta = (\sigma_i^\Delta)_{i=1}^N$. Under Assumptions 2, 3, and 4, the discrete-time Bellman equation for a period length of Δ is

$$\begin{aligned} V_i^\Delta(\omega, J) &= |\mathcal{J}| \sum_{j \in J} u_{i,j}(\omega, \sigma_j^\Delta(\omega)) \Delta \\ &+ \exp(-\rho \Delta) \left\{ \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \sum_{a_J \in A_J(\omega)} \sigma_J^\Delta(a_J | \omega) \Pr(J'|J, \omega) \left(1 - |\mathcal{J}| \sum_{j \in J} q_j(\omega, a_j) \Delta \right) \right. \\ &\left. + \sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \sum_{a_J \in A_J(\omega)} \sigma_J^\Delta(a_J | \omega) \Pr(J'|J, \omega) \left(|\mathcal{J}| \sum_{j \in J} \varphi_j(\omega' | \omega, a_j) \Delta \right) \right\} + \mathcal{O}(\Delta^2). \quad (\text{B.1}) \end{aligned}$$

Taking the limit as $\Delta \rightarrow 0$, we deduce that $V_i^0(\omega, J) = \sum_{J' \in \mathcal{J}} V_i^0(\omega, J') \Pr(J'|J, \omega)$. Assumption 4 implies that the transition matrix $\Pr(J' | J, \omega)$ has a unique (and uniform) right eigenvector so that $V_i^0(\omega, J) = V_i^0(\omega, J')$ for all $J, J' \in \mathcal{J}$. Let $V_i^0: \Omega \rightarrow \mathbb{R}$ be the value function of player i and $V^0 = (V_i^0)_{i=1}^N$ be the profile of value functions in the limit as $\Delta \rightarrow 0$.

³⁴An example of a generalized protocol of moves is the following. For all $\omega \in \Omega$, consider a transition matrix P_ω over \mathcal{J} with a uniform invariant distribution. Taking $\Pr(J' | J, \omega) = P_\omega(J' | J)$ defines a generalized protocol of moves.

Rewriting the Bellman equation, multiplying by $\pi(J)$, and summing over $J \in \mathcal{J}$ yields

$$\begin{aligned} & \frac{1}{\Delta} \sum_{J \in \mathcal{J}} \pi(J) V_i^\Delta(\omega, J) - \frac{\exp(-\rho\Delta)}{\Delta} \sum_{J \in \mathcal{J}} \pi(J) \sum_{J' \in \mathcal{J}} \sum_{a_J \in A_J(\omega)} \sigma_J^\Delta(a_J | \omega) V_i^\Delta(\omega, J') \Pr(J' | J, \omega) \\ &= |\mathcal{J}| \sum_{J \in \mathcal{J}} \pi(J) \sum_{j \in J} u_{i,j}(\omega, \sigma_j^\Delta(\omega)) \\ &+ \exp(-\rho\Delta) |\mathcal{J}| \sum_{J \in \mathcal{J}} \pi(J) \sum_{j \in J} \left\{ \sum_{\omega' \neq \omega} \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega', J') \Pr(J' | J, \omega) \varphi_j(\omega' | \omega, \sigma_j^\Delta(\omega)) \right. \\ &\quad \left. - \sum_{J' \in \mathcal{J}} V_i^\Delta(\omega, J') \Pr(J' | J, \omega) q_j(\omega, \sigma_j^\Delta(\omega)) \right\} + \mathcal{O}(\Delta^2). \end{aligned}$$

Using the facts that $\sum_{J \in \mathcal{J}} \Pr(J' | J, \omega) \pi(J) = \pi(J')$ and $\sum_{J' \in \mathcal{J}} \Pr(J' | J, \omega) = 1$, and taking the limit as $\Delta \rightarrow 0$, we obtain the continuous-time Bellman equation

$$\rho V_i^0(\omega) = |\mathcal{J}| \sum_{J \in \mathcal{J}} \pi(J) \sum_{j \in J} u_{i,j}(\omega, \sigma_j^0(\omega)) + |\mathcal{J}| \sum_{J \in \mathcal{J}} \pi(J) \sum_{j \in J} \left(\sum_{\omega' \neq \omega} V_i^0(\omega') \varphi_j(\omega' | \omega, \sigma_j^0(\omega)) - V_i^0(\omega) q_j(\omega, \sigma_j^0(\omega)) \right). \quad (\text{B.2})$$

The discrete-time optimality condition for a period length Δ is

$$\begin{aligned} & \sigma^\Delta(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_i^\Delta(\omega, J, \tilde{a}_i, \sigma_{J \setminus \{i\}}^\Delta(\omega)) \\ &+ \exp(-\rho\Delta) \sum_{\omega' \in \Omega} \sum_{J' \in \mathcal{J}} \sum_{a_{J \setminus \{i\}}} \sigma_{J \setminus \{i\}}^\Delta(a_{J \setminus \{i\}} | \omega) V_i^\Delta(\omega', J') \Pr(J' | J, \omega) \Pr^\Delta(\omega' | \omega, J, \tilde{a}_i, a_{J \setminus \{i\}}). \end{aligned}$$

Dividing by Δ , rearranging terms, and taking the limit as $\Delta \rightarrow 0$, we deduce the continuous-time optimality condition

$$\sigma_i^0(a_i | \omega) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_{i,i}(\omega, \tilde{a}_i) + \sum_{\omega' \neq \omega} (V_i^0(\omega') - V_i^0(\omega)) \varphi_i(\omega' | \omega, \tilde{a}_i). \quad (\text{B.3})$$

Conditions (B.2) and (B.3) are the analogs of conditions (3.2) and (3.3) for a generalized protocol of moves.

Consider the generalized protocols of moves $\langle \mathcal{J}_1, \mathcal{P}_1 \rangle$ and $\langle \mathcal{J}_2, \mathcal{P}_2 \rangle$ with stationary distributions π_1 and π_2 . For all $j = 1, \dots, N$, define $J_1(j)$ to be the unique element in \mathcal{J}_1 such that $j \in J_1(j)$. Define $J_2(j)$ analogously. We say that the generalized protocols of moves $\langle \mathcal{J}_1, \mathcal{P}_1 \rangle$ and $\langle \mathcal{J}_2, \mathcal{P}_2 \rangle$ are *comparable* if $|\mathcal{J}_1| \pi_1(J_1(j)) = |\mathcal{J}_2| \pi_2(J_2(j))$ for all $j = 1, \dots, N$. Given a protocol $\langle \mathcal{J}, \mathcal{P} \rangle$, a player j moves a fraction $\pi(J(j))$ of the time and has an impact on payoffs and transitions which

is scaled by $|\mathcal{J}|$. Thus, comparability means that the total impact of a player's strategy on payoffs and transitions does not depend on the particular protocol that we use in the model. All protocols that satisfy Assumption 1 are comparable.

Theorem 1 remains valid for generalized protocols of moves that are comparable.³⁵

Theorem 2 (Generalized Protocol-Invariance Theorem) *Fix p , q , and ρ . For almost all u , all generalized protocols of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ and $\langle \overline{\mathcal{J}}, \overline{\mathcal{P}} \rangle$ that are comparable, and all $\varepsilon > 0$, there exists $\overline{\Delta} > 0$ such that for all $\Delta < \overline{\Delta}$ and $\sigma \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$, there exists $\overline{\sigma} \in \text{Equil}(\langle \Delta, \overline{\mathcal{J}}, \overline{\mathcal{P}}, u, p, q, \rho \rangle)$ such that $\|\sigma - \overline{\sigma}\| < \varepsilon$.*

B.2 Non-Partition Protocol of Moves

We now provide a partial extension to Theorem 1 that applies when the protocol of moves is not a partition. We thus relax Assumption 1 and assume that \mathcal{J} is not a partition of the set of players but contains subsets $J \subseteq \{1, \dots, N\}$ such that for all $i = 1, \dots, N$, there exists $J \in \mathcal{J}$ such that $i \in J$. This allows player i to have the move in conjunction with different sets of rivals. To simplify the exposition, we assume that $|\{J \in \mathcal{J} \mid i \in J\}| = \kappa$ for all $i = 1, \dots, N$. As before, there is an irreducible Markov chain \mathcal{P} defined on \mathcal{J} that has a unique stationary distribution that is uniform on \mathcal{J} . We call $\langle \mathcal{J}, \mathcal{P} \rangle$ a *non-partition protocol of moves*.

With a non-partition protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$, the per-period payoff $u_i^\Delta(\omega, J, a_J)$ is written as

$$u_i^\Delta(\omega, J, a_J) = \frac{|\mathcal{J}|}{\kappa} \sum_{j \in J} u_{i,j}(\omega, a_j) \Delta + \mathcal{O}(\Delta^2),$$

and the hazard rate $q_J(\omega, a_J)$ and transition probability $p_J(\omega' \mid \omega, a_J)$ are written as

$$q_J(\omega, a_J) = \frac{|\mathcal{J}|}{\kappa} \sum_{j \in J} q_j(\omega, a_j)$$

and

$$q_J(\omega, a_J) p_J(\omega' \mid \omega, a_J) = \frac{|\mathcal{J}|}{\kappa} \sum_{j \in J} q_j(\omega, a_j) p_j(\omega' \mid \omega, a_j),$$

where $q_j: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{R}^+ \cup \{0\}$ and $p_j: \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \rightarrow \mathbb{P}(\Omega)$. The remaining aspects of the model are unchanged.

³⁵Strictly speaking, here we only show that Lemma 1 remains valid. The proof that Lemma 2 remains valid is available upon request.

With a non-partition protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$, the identity of the players that have the move in conjunction with player i is a state variable. Thus, a Markovian strategy for player i is a function $\sigma_i: \Omega \times \{J \in \mathcal{J} \mid i \in J\} \rightarrow \cup_{\omega \in \Omega} \mathbf{P}(A_i(\omega))$. Overloading notation, we use $\text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$ to denote the set of Markov perfect equilibria. We say that a Markov perfect equilibrium σ is *simple* if $\sigma_i(a_i \mid \omega, J) = \sigma_i(a_i \mid \omega, \tilde{J})$ for all $i = 1, \dots, N$, $\omega \in \Omega$, $a_i \in A_i$, and all $J, \tilde{J} \in \mathcal{J}$. In this case, we write $\sigma_i(a_i \mid \omega)$.

The following proposition partially extends Theorem 1 to a non-partition protocol of moves:³⁶

Proposition 6 *Assume $\text{Equil}^0(\langle u, p, q, \rho \rangle)$ only contains strict solutions. Then there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, $\sigma \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$ is simple and*

$$\text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle) = \text{Equil}^0(\langle u, p, q, \rho \rangle).$$

In contrast to Theorem 1, Proposition 6 restricts attention to strict and thus pure solutions. When mixed solutions are considered, the limit conditions may have a continuum of solutions if players use the payoff-irrelevant realization of J to randomize over actions and our differential topology tools therefore cannot be applied.

Proof. Consider a sequence (σ^Δ) with $\sigma^\Delta \in \text{Equil}(\langle \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho \rangle)$ and $\sigma^\Delta \rightarrow \sigma^0$ (possibly through a subsequence) as $\Delta \rightarrow 0$. Let V^Δ be the profile of value functions corresponding to σ^Δ and assume it converges to V^0 . Similar to Section 3, we can deduce that $V^0(\omega, J)$ does not depend on $J \in \mathcal{J}$ and simply write $V^0(\omega)$. We can also follow Section 3 to deduce that

$$\rho V_i^0(\omega) = \frac{1}{\kappa} \sum_{J \in \mathcal{J}} \sum_{j \in J} \left(u_{i,j}(\omega, \sigma_j(\omega, J)) + \sum_{\omega' \neq \omega} (V_i^0(\omega') - V_i^0(\omega)) \varphi_j(\omega' \mid \omega, \sigma_j(\omega, J)) \right)$$

and

$$\sigma_i^0(a_i \mid \omega, J) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_{i,i}(\omega, \tilde{a}_i) + \sum_{\omega' \neq \omega} (V_i^0(\omega') - V_i^0(\omega)) \varphi_i(\omega' \mid \omega, \tilde{a}_i). \quad (\text{B.4})$$

Define $\tilde{\sigma}_i(\cdot \mid \omega) = \frac{1}{\kappa} \sum_{J \in \mathcal{J}} \sigma_i^0(\cdot \mid \omega, J)$ for all $i = 1, \dots, N$ and all $\omega \in \Omega$, and note that $(\tilde{\sigma}, V^0)$ is a solution to the limit conditions (3.2) and (3.3). Since $\text{Equil}^0(\langle u, p, q, \rho \rangle)$ only contains strict solutions, the profile $\tilde{\sigma} = (\tilde{\sigma}_i)_{i=1}^N$ must be a strict solution and thus the maximization problem in

³⁶We have been unable to find an example showing Theorem 1 cannot be fully extended to a non-partition protocol of moves.

equation (B.4) has a unique solution. Therefore $\sigma_i^0(a_i | \omega, J)$ does not depend on J and σ^0 is simple. In particular, $\sigma^0 \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$ and therefore there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, $\sigma^\Delta \in \text{Equil}^0(\langle u, p, q, \rho \rangle)$. To see the converse, note that $\text{Equil}^0(\langle u, p, q, \rho \rangle)$ has a finite number of pure solutions that are all strict. For any solution (σ^0, V^0) to the limit conditions (3.2) and (3.3), σ^0 satisfies the conditions for a (simple) Markov perfect equilibrium of a separable dynamic game with noisy transitions and the non-partition protocol of moves $\langle \mathcal{J}, \mathcal{P} \rangle$ since the continuation values in such a game converge to V^0 and, as a result, the incentive constraints are satisfied if $\Delta > 0$ is sufficiently small. ■

B.3 Modified Entry Game

We slightly modify the entry game following Lemma 1 in Section 3 to show that the solution to the limit conditions (3.2) and (3.3) can be approximated by a discrete-time game with alternating moves but not with simultaneous moves. In particular, we assume that the expected net present value of the stream of payoffs to a monopolist is $\frac{B(\Delta + \Delta^2)}{1 - e^{-\rho\Delta}}$ instead of $\frac{B\Delta}{1 - e^{-\rho\Delta}}$. Clearly, this does not change the limit conditions (3.2) and (3.3) and the fact that σ^0 is a solution to these conditions provided that $\lambda B/\rho = c$ and $b < 0$.

We first show that this solution can be approximated by the discrete-time game with alternating moves provided that $\rho < 2$. A firm has an incentive to invest if and only if

$$-c + e^{-\rho\Delta} \lambda B \frac{\Delta + \Delta^2}{1 - e^{-\rho\Delta}} \geq 0.$$

This condition holds for Δ sufficiently small. To see this, note that the inequality holds with equality at $\Delta = 0$ while the derivative of the left-hand side at $\Delta = 0$ equals $-\lambda B + \lambda B(1/2 + 1/\rho) > 0$.

Next we show that this solution cannot be approximated by the discrete-time game with simultaneous moves. A firm has an incentive to invest if and only if

$$-c + e^{-\rho\Delta} \Delta(\lambda)^2 b \frac{\Delta}{1 - e^{-\rho\Delta}} + e^{-\rho\Delta} \lambda(1 - \lambda\Delta) B \frac{\Delta + \Delta^2}{1 - e^{-\rho\Delta}} \geq 0.$$

This condition holds in the limit as $\Delta \rightarrow 0$ but not for $\Delta > 0$. To see this, note that the inequality holds with equality at $\Delta = 0$ while the derivative of the left-hand side at $\Delta = 0$ equals $\lambda B(1/\rho - 1/2 - \lambda) + \lambda^2 b/\rho < 0$.

B.4 Non-Separable State-to-State Transitions

Consider a dynamic stochastic game with $N = 2$ players, $\Omega = \{0, 1\}$, and $A_i = \{0, 1\}$. The hazard rate in state $\omega = 0$ is $q(0, a_1, a_2) = 1$ if and only if $a_1 = a_2 = 1$, and $q(0, a_1, a_2) = 0$ otherwise, whereas in state $\omega = 1$, $q(1, a) = 0$. The transition probability satisfies $p(1 | 0, (a_1, a_2)) = 1$. State $\omega = 1$ is thus absorbing. Flow payoffs do not depend on actions and take the form $u_{i,i}(\omega) = \omega$. In the game with simultaneous moves, it is simple to see that there exist two Markov perfect equilibria in pure strategies. In one of them, the state is stuck in $\omega = 0$. In contrast, in the game with alternating moves and transitions “materializing” only once both players have made a decision (in a violation of Assumption 3 similar to Example 6), the unique Markov perfect equilibrium is $\sigma_1^* = 1$, $\sigma_2^*(a_1) = a_1$, and the state eventually jumps to $\omega' = 1$.

References

- ABREU, D., AND F. GUL (2000): “Bargaining and Reputation,” *Econometrica*, 68(1), 85–117.
- ALIPRANTIS, C., AND K. BORDER (2006): *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer Science & Business Media.
- AMBRUS, A., AND Y. ISHII (2015): “On Asynchronicity of Moves and Coordination,” Working paper, Duke University.
- AMBRUS, A., AND S. LU (2015): “A Continuous-Time Model of Multilateral Bargaining,” *American Economic Journal: Microeconomics*, 7, 208–249.
- ARAPOSTATHIS, A., V. S. BORKAR, E. FERNÁNDEZ-GAUCHERAND, M. K. GHOSH, AND S. I. MARCUS (1993): “Discrete Time Controlled Markov Processes with Average Cost Criterion: A Survey,” *SIAM Journal on Control and Optimization*, 31(2), 282–344.
- ARCIDIACONO, P., P. BAYER, J. BLEVINS, AND P. ELLICKSON (2016): “Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition,” *Review of Economic Studies*, 83(3), 889–931.
- BASU, K., AND J. WEIBULL (1991): “Strategy Subsets Closed under Rational Behavior,” *Economics Letters*, 36(2), 141–146.
- BERRY, S. (1992): “Estimation of a Model of Entry in the Airline Industry,” *Econometrica*, 60(4), 889–917.

- BERRY, S., AND A. PAKES (1993): “Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis,” *American Economic Review*, 83(2), 247–252.
- BESANKO, D., AND U. DORASZELSKI (2004): “Capacity Dynamics and Endogenous Asymmetries in Firm Size,” *Rand Journal of Economics*, 35(1), 23–49.
- BHASKAR, V., G. MAILATH, AND S. MORRIS (2009): “A Foundation for Markov Equilibria in Infinite Horizon Perfect Information Games,” Working paper, Department of Economics, University of Pennsylvania.
- (2013): “A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory,” *Review of Economic Studies*, 80(3), 925–948.
- BHASKAR, V., AND F. VEGA-REDONDO (2002): “Asynchronous Choice and Markov Equilibria,” *Journal of Economic Theory*, 103(2), 334–350.
- BOHREN, A. (2014): “Stochastic Games in Continuous Time: Persistent Actions in Long-Run Relationships,” Working paper, Department of Economics, University of Pennsylvania.
- BOLTON, P., AND C. HARRIS (1999): “Strategic Experimentation,” *Econometrica*, 67(2), 349–374.
- BRESNAHAN, T., AND P. REISS (1990): “Entry into Monopoly Markets,” *Review of Economic Studies*, 57, 531–553.
- CALCAGNO, R., Y. KAMADA, S. LOVO, AND T. SUGAYA (2014): “Asynchronicity and Coordination in Common and Opposing Interest Games,” *Theoretical Economics*, 9(2), 409–434.
- CARUANA, G., AND L. EINAV (2008): “A Theory of Endogenous Commitment,” *Review of Economic Studies*, 75(1), 99–116.
- CHEN, J. (2009): “The Effects of Mergers with Dynamic Capacity Accumulation,” *International Journal of Industrial Organization*, 27, 92–109.
- CLEEREN, K., F. VERBOVEN, M. DEKIMPE, AND K. GIELENS (2010): “Intra- and Interformat Competition Among Discounters and Supermarkets,” *Marketing Science*, 29(3), 456–473.
- COMPTE, O., AND P. JEHIEL (2004): “Gradualism in Bargaining and Contribution Games,” *Review of Economic Studies*, 71(4), 975–1000.
- COURNOT, A. (1838): *Recherches sur les Principes Mathématiques de la Théorie des Richesses*. Hachette, Paris.

- CYERT, R., AND M. DEGROOT (1970): “Multi-period Decision Models with Alternating Choice as a Solution to the Duopoly Problem,” *Quarterly Journal of Economics*, 84(3), 410–299.
- DE ROOS, N. (2004): “A Model of Collusion Timing,” *International Journal of Industrial Organization*, 22, 351–387.
- DOCKNER, E. (1992): “A Dynamic Theory of Conjectural Variations,” *The Journal of Industrial Economics*, pp. 377–395.
- DORASZELSKI, U., AND J. ESCOBAR (2010): “A Theory of Regular Markov Perfect Equilibria in Dynamic Stochastic Games: Genericity, Stability, and Purification,” *Theoretical Economics*, 5, 369–402.
- DORASZELSKI, U., AND K. JUDD (2012): “Avoiding the Curse of Dimensionality in Dynamic Stochastic Games,” *Quantitative Economics*, 3(1), 53–93.
- (2018): “Dynamic Stochastic Games with Random Moves,” *Quantitative Economics and Marketing*, forthcoming.
- DORASZELSKI, U., AND S. MARKOVICH (2007): “Advertising Dynamics and Competitive Advantage,” *Rand Journal of Economics*, 38(3), 557–592.
- DORASZELSKI, U., AND A. PAKES (2007): “A Framework for Applied Dynamic Analysis in IO,” in *Handbook of Industrial Organization*, vol. 3, pp. 1887–1966. Elsevier.
- DORASZELSKI, U., AND M. SATTERTHWAITTE (2010): “Computable Markov Perfect Industry Dynamics,” *Rand Journal of Economics*, 41(2), 215–243.
- DUTTA, P. (1995): “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, 66(1), 1–32.
- EINAV, L. (2010): “Not All Rivals Look Alike: Estimating an Equilibrium Model of the Release Date Timing Game,” *Economic Inquiry*, 48(2), 369–390.
- ELFVING, G. (1937): “Zur Theorie der Markoffschen Ketten,” *Acta Societatis Scientiarum Fennicae Nova Series A*, 2(8), 1–17.
- ERICSON, R., AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62(1), 53–82.
- FAINGOLD, E., AND Y. SANNIKOV (2011): “Reputation in Continuous-Time Games,” *Econometrica*, 79(3), 773–876.

- FERSHTMAN, C., AND A. PAKES (2000): “A Dynamic Oligopoly with Collusion and Price Wars,” *Rand Journal of Economics*, 31, 294–326.
- FUDENBERG, D., R. GILBERT, J. STIGLITZ, AND J. TIROLE (1983): “Preemption, leapfrogging and competition in patent races,” *European Economic Review*, 22, 3–31.
- FUDENBERG, D., AND D. LEVINE (1986): “Limit Games and Limit Equilibria,” *Journal of Economic Theory*, 38(2), 261–279.
- FUDENBERG, D., AND J. TIROLE (1983): “Capital as Commitment: Strategic Investment to Deter Mobility,” *Journal of Economic Theory*, 31(2), 227–250.
- (1985): “Preemption and Rent Equalization in the Adoption of New Technology,” *Review of Economic Studies*, 52(3), 383–401.
- GEORGIADIS, G. (2015): “Projects and Team Dynamics,” *Review of Economic Studies*, 82(1), 187–218.
- GIOVANNETTI, E. (2001): “Perpetual Leapfrogging in Bertrand Duopoly,” *International Economic Review*, 42(3), 671–969.
- GOWRISANKARAN, G. (1999): “A Dynamic Model of Endogenous Horizontal Mergers,” *Rand Journal of Economics*, 30(1), 56–83.
- GROSSMAN, G., AND C. SHAPIRO (1987): “Dynamic R&D competition,” *Economic Journal*, 97(386), 372–387.
- HALKIN, H. (1974): “Implicit Functions and Optimization Problems without Continuous Differentiability of the Data,” *SIAM Journal on Control*, 12(2), 229–236.
- HALLER, H., AND R. LAGUNOFF (2000): “Genericity and Markovian Behavior in Stochastic Games,” *Econometrica*, 68(5), 1231–1248.
- HANIG, M. (1986): “Differential Gaming Models of Oligopoly,” Ph.D. thesis, Massachusetts Institute of Technology.
- HARRIS, C., AND J. VICKERS (1987): “Racing with uncertainty,” *Review of Economic Studies*, 54(1), 1–21.
- HARSANYI, J. (1973a): “Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *International Journal of Game Theory*, 2(1), 1–23.
- (1973b): “Oddness of the Number of Equilibrium Points: A New Proof,” *International Journal of Game Theory*, 2(4), 235–250.

- HÖRNER, J., T. SUGAYA, S. TAKAHASHI, AND N. VIEILLE (2011): “Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem,” *Econometrica*, 79(4), 1277–1318.
- IFRACH, B., AND G. WEINTRAUB (2017): “A framework for dynamic oligopoly in concentrated industries,” *Review of Economic Studies*, 84(3), 1106–1150.
- ISKHAKOV, F., J. RUST, AND B. SCHJERNING (2017): “The Dynamics of Bertrand Price Competition with Cost-Reducing Investments,” *International Economic Review*, forthcoming.
- KALAI, E. (2004): “Large Robust Games,” *Econometrica*, 72(6), 1631–1665.
- KALAI, E., AND D. SAMET (1984): “Persistent Equilibria in Strategic Games,” *International Journal of Game Theory*, 13(3), 129–144.
- KAMADA, Y., AND M. KANDORI (2017): “Revision Games Part I: Theory,” Discussion paper, Harvard University.
- LAGUNOFF, R., AND A. MATSUI (1997): “Asynchronous Choice in Repeated Coordination Games,” *Econometrica*, pp. 1467–1477.
- (2001): “Are “Anti Folk” Theorems in Repeated Games Nongeneric,” *Review of Economic Dynamics*, 6, 397–412.
- LIPPMAN, S., AND K. MCCARDLE (1987): “Dropout Behavior in R&D Races with Learning,” *Rand Journal of Economics*, 18(2), 287–295.
- MAILATH, G., AND S. MORRIS (2002): “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory*, 102(1), 189–228.
- MAILATH, G., AND L. SAMUELSON (2006): *Repeated Games and Reputations*. Oxford University Press, Oxford.
- MARKOVICH, S. (2008): “Snowball: A Dynamic Oligopoly Model with Indirect Network Effects,” *Journal of Economic Dynamics and Control*, 32, 909–938.
- MARX, L., AND S. MATTHEWS (2000): “Dynamic Voluntary Contribution to a Public Project,” *Review of Economic Studies*, 67(2), 327–358.
- MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic Theory*. Oxford University Press, Oxford.
- MASKIN, E., AND J. TIROLE (1987): “A Theory of Dynamic Oligopoly, III: Cournot Competition,” *European Economic Review*, 31(4), 947–968.

- (1988a): “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles,” *Econometrica*, 56(3), 571–599.
- (1988b): “A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs,” *Econometrica*, pp. 549–569.
- (2001): “Markov Perfect Equilibrium, I: Observable Actions,” *Journal of Economic Theory*, 100(2), 191–219.
- MAZZEO, M. (2002): “Product Choice and Oligopoly Market Structure,” *Rand Journal of Economics*, 33(2), 1–22.
- PAKES, A., AND P. MCGUIRE (1994): “Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model,” *Rand Journal of Economics*, 25(4), 555–589.
- (2001): “Stochastic algorithms, symmetric Markov perfect equilibrium, and the “curse” of dimensionality,” *Econometrica*, 69(5), 1261–1281.
- PESKI, M., AND T. WISEMAN (2015): “A Folk Theorem for Stochastic Games with Infrequent State Changes,” *Theoretical Economics*, 10(1), 131–173.
- REINGANUM, J. (1982): “A Dynamic Game of R and D: Patent Protection and Competitive Behavior,” *Econometrica*, 50, 671–688.
- (1989): “The Timing of Innovation: Research, Development, and Diffusion,” in *Handbook of industrial organization*, ed. by R. Schmalensee, and R. Willig. North-Holland, Amsterdam.
- REYNOLDS, S. (1987): “Capacity Investment, Preemption and Commitment in an Infinite Horizon Model,” *International Economic Review*, pp. 69–88.
- REYNOLDS, S. (1991): “Dynamic Oligopoly with Capacity Adjustment Costs,” *Journal of Economic Dynamics and Control*, 15(3), 491–514.
- RIORDAN, M., AND D. SALANT (1994): “Preemptive Adoptions of an Emerging Technology,” *Journal of Industrial Economics*, 42(3), 247–261.
- ROSENTHAL, R. (1991): “A Note on Robustness of Equilibria with Commitment Opportunities,” *Games and Economic Behavior*, 3, 237–243.
- RUBINSTEIN, A., AND A. WOLINSKY (1995): “Remarks on Infinitely Repeated Extensive-Form Games,” *Games and Economic Behavior*, 9(1), 110–115.

- SANNIKOV, Y., AND A. SKRZYPACZ (2007): “Impossibility of Collusion Under Imperfect Monitoring with Flexible Production,” *American Economic Review*, 97(5), 1794–1823.
- SERFOZO, R. (1979): “An Equivalence Between Continuous and Discrete Time Markov Decision Processes,” *Operations Research*, 27(3), 616–620.
- SHAPLEY, L. (1953): “Stochastic Games,” *Proceedings of the National Academy of Sciences*, 39(10), 1095–1100.
- SIMON, L., AND M. STINCHCOMBE (1989): “Extensive form Games in Continuous Time: Pure Strategies,” *Econometrica*, pp. 1171–1214.
- SPENCE, A. (1979): “Investment Strategy and Growth in a New Market,” *Bell Journal of Economics*, pp. 1–19.
- VAN DAMME, E., AND S. HURKENS (1996): “Commitment Robust Equilibria and Endogenous Timing,” *Games and Economic Behavior*, 15, 290–311.
- VON STACKELBERG, H. (1934): *Marktform und Gleichgewicht*. Springer, Wien.
- WEINTRAUB, G., L. BENKARD, AND B. VAN ROY (2008): “Markov perfect industry dynamics with many firms,” *Econometrica*, 76(6), 1375–1411.
- WEN, Q. (2002): “A Folk Theorem for Repeated Sequential Games,” *Review of Economic Studies*, 69(2), 493–512.
- YOON, K. (2001): “A Folk Theorem for Asynchronously Repeated Games,” *Econometrica*, 69(1), 191–200.