DISCUSSION PAPER SERIES

DP11447

PROTOCOL INVARIANCE AND THE TIMING OF DECISIONS IN DYNAMIC GAMES

Ulrich Doraszelski and Juan Escobar

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

PROTOCOL INVARIANCE AND THE TIMING OF DECISIONS IN DYNAMIC GAMES

Ulrich Doraszelski and Juan Escobar

Discussion Paper DP11447
Published 20 August 2016
Submitted 20 August 2016

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Ulrich Doraszelski and Juan Escobar

PROTOCOL INVARIANCE AND THE TIMING OF DECISIONS IN DYNAMIC GAMES

Abstract

The timing of decisions is an essential ingredient in modelling any strategic situation. Yet, determining the most realistic and appropriate protocol of moves can be challenging. We introduce a class of dynamic stochastic games that we call separable dynamic games with noisy transitions and establish that they 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. We also show that the equilibria have a remarkably simple structure.

JEL Classification: N/A

Keywords: Protocol of moves, Dynamic games, Markov perfect equilibrium

Ulrich Doraszelski - ulrichd@wharton.upenn.edu
University of Pennsylvania and CEPR

Juan Escobar - jescobar@dii.uchile.cl
University of Chile

Protocol Invariance and the Timing of Decisions in Dynamic Games*

Ulrich Doraszelski† Juan F. Escobar‡

July 29, 2016

Abstract

The timing of decisions is an essential ingredient in modelling any strategic situation. Yet, determining the most realistic and appropriate protocol of moves can be challenging. We introduce a class of dynamic stochastic games that we call separable dynamic games with noisy transitions and establish that they 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. We also show that the equilibria have a remarkably simple structure.

*We thank Yuichiro Kamada, George Mailath, and Juuso Toikka 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, Ministerio de Economía.

†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 in modelling any strategic situation. 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. Yet, 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.

The timing of decisions is not only difficult to determine in practice but also often consequential. From the basically static models in Cournot (1838) and von Stackelberg (1934) to the genuinely dynamic models in Cyert and DeGroot (1970), Maskin and Tirole (1987, 1988a, 1988b), Cahuc and Kempf (1997), Noel (2008), and Iskhakov, Rust, and Schjerning (2016) 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. In fact, it arguably matters more than the perturbations of payoffs that are traditionally the focus of robustness analysis (Harsanyi 1973a, Doraszelski and Escobar 2010). 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.

Against this backdrop, we present a fairly general class of dynamic models that are 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 ensures that the implications and predictions of the model are independent of the timing of decisions and thus more robust for the purposes of applied work.

We focus on infinite-horizon dynamic stochastic games and their stationary Markov perfect equilibria (henceforth Markov perfect equilibria for short). Dating back to Shapley (1953), these models have a long tradition in economics and are central to the analysis of strategic interactions among forward-looking players in dynamic environments. A dynamic stochastic game is a dynamic system that can be in different states at different times. The evolution of the state from one period to the next is governed by a discrete-time Markov process that players can influence through their actions. Each player strives to maximize the expected net present value of his stream of payoffs. Applications of dynamic stochastic games abound and range from public finance (Bernheim and Ray 1989) and political economy (Acemoglu and Robinson 2001) to industrial organization (Ericson and Pakes 1995, Bergemann and Välimäki 1996).

The main contribution of this paper is to introduce a new class of dynamic stochastic games that we call *separable dynamic games with noisy transitions* and to establish that these models are protocol invariant

provided that periods are sufficiently short. In addition, we show that the Markov perfect equilibria of separable dynamic stochastic games with noisy transitions have a remarkably simple structure.

While per-period payoffs and state-to-state transitions in a general dynamic stochastic game depend arbitrarily on the state and on players' actions, in a separable dynamic game per-period payoffs and state-to-state transitions depend on players' actions in an additive manner: to a first-order approximation, they are build from parts that depend on the actions taken by subsets of players. As a consequence, the strategic situation that the N players face in a given state of the dynamic system, holding fixed the value of continued play, is akin to N independent optimization problems. While at first glance this may seem to trivialize the game, the separability assumption does not restrict how per-period payoffs and state-to-state transitions depend on the state. Because players are forward looking, this allows strategic interactions to be channeled through continuation values and ensures that the separability assumption is not overly onerous. Indeed, our framework is quite rich.

The assumption of noisy transitions precludes that there is an action that a player can take to guarantee a change in the state. To impose noisy transitions 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. This captures the idea that in many real-world settings players often take a certain amount of time to finalize their actions (Iijima and Kasahara 2015).

Separablility and noisy transitions doubtlessly rule out some interesting applications. Many dynamic stochastic games in the literature are nevertheless 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), and models of industry dynamics (Ericson and Pakes 1995). The recent continuous-time stochastic games with moves at random times (Arcidiacono, Bayer, Blevins, and Ellickson 2015, Ambrus and Lu 2015, Calcagno, Kamada, Lovo, and Sugaya 2014) also fit into our framework.

Our first main result, Theorem 1, 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 alternatingly, 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. It follows that the protocol of moves is almost immaterial to the

firm's decision.

The proof of Theorem 1 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. Hence, when the limit conditions admit a unique solution, our protocol-invariance theorem is immediate. However, our framework is quite rich and the limit conditions may admit multiple solutions. To handle the difficulties arising from multiplicity, we introduce differential topology tools to study the limit conditions.[1] Drawing on ideas in Harsanyi (1973a, 1973b) for normal-form games and in Doraszelski and Escobar (2010) for dynamic stochastic games, we prove that generically the limit conditions have a finite number of solutions and that all solutions 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.

Our assumptions of separability and noisy transitions are tight in the sense that counterexamples to protocol invariance can be constructed if any one of them is relaxed. Moreover, we show that protocol invariance does not extend beyond Markov perfect equilibria to other equilibrium concepts.

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 the richness of our framework and clarifies the strategic implications of assuming separability and noisy transitions.

Our second main result, Theorem 2, shows that the Markov perfect equilibria of a separable dynamic game with noisy transitions have a remarkably simple structure provided that periods are sufficiently short. In particular, we show that the number of actions that a player uses with positive probability in a given state cannot exceed the number of players in the game. The proof of Theorem 2 capitalizes on the analytic tractability of the limit conditions and on Theorem 1 to extend the analysis to separable dynamic games with noisy transitions.

Taken together, our main results facilitate and inform applied work in a number of ways. First and most important, they alleviate the burden of determining the protocol of moves that is most realistic and appropriate for the application at hand. Protocol invariance ensures that equilibrium behavior in a separable dynamic game with noisy transitions is independent of the timing of decisions provided that periods are

[1] While these tools have been used in economic theory to prove generic robustness to changes in information and specification of payoffs (Harsanyi 1973a, Harsanyi 1973b), we seem to be the first ones using them to explore the robustness of dynamic game predictions to modifications in the timing and the protocols of moves.

sufficiently short. Second, they caution against the presumption that imposing asynchronous instead of synchronous decisions on a dynamic stochastic game reduces the number of equilibria. Third, 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. Fourth, dynamic stochastic games are often not very tractable analytically and thus call for the use of numerical methods. Doraszelski and Judd (2007) show that the computational burden can vary by orders of magnitude with the protocol of moves. Protocol invariance justifies imposing the protocol of moves that is most convenient from a computational perspective.

We apply and extend our main results in three ways. First, we provide a new justification for focusing on Markov perfect equilibria in dynamic stochastic games. Provided periods are sufficiently short, we show that if a strict finite-memory equilibrium payoff profile in a separable dynamic game with noisy transitions is protocol invariant, then it must be attained by a Markov perfect equilibrium. In this sense, Markov perfect equilibria are the only equilibria that are robust to changes in the protocol of moves. This result adds to the literature providing foundations for Markov perfect equilibria (Maskin and Tirole 2001, Bhaskar and Vega-Redondo 2002, Bhaskar, Mailath, and Morris 2013, Bohren 2014). Second, as Doraszelski and Judd (2012) argue, the limit conditions that arise as we pass from discrete to continuous time are particularly easy to solve numerically. We provide a justification for doing so by showing that the solutions to the limit conditions 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. We also show that restricting attention to games with simultaneous moves, we can dispense with the separability assumption. Third, we shed light on the numerous examples in the literature in which the protocol of moves matters crucially for equilibrium behavior. We show that there is a discontinuity in the set of Markov perfect equilibria as hazard rates become large and moves become frequent. In particular, the celebrated results in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) are knife edge cases that hinge on assuming infinite hazard rates. We view the fact that the order of limits has a significant effect on equilibrium behavior as a word of caution to researchers working on games with infinite hazard rates and arbitrarily frequent 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. 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. Our work is also related to Kalai (2004), who shows that the Nash equilibria of large anonymous Bayesian games are approximately robust to variations in the extensive-form version of the game. 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 vanishes as periods become short. From a more technical perspective, Kalai (2004) allows for ϵ-equilibria, while we impose exact equilibrium and establish our results for generic payoffs.

Second, previous attempts to exposit dynamic games where the protocol of moves does not matter for equilibrium behavior are few and far between and confined to very specific models. 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. As our framework 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. Sections 3 and 4 develop our main results. Section 5 discusses a number of applications and extensions of our main results and Section 6 concludes. The proofs of our main results are in the Appendix. An Online Appendix provides further examples and proofs.

2 Separable Dynamic Games with Noisy Transitions

We focus on dynamic stochastic games with finite sets of players, states, and actions. Let $\{1, 2, \ldots, N\}$ denote the set of players, Ω the set of states, and $A_i(\omega)$ the set of actions of player i in state ω. Time is discrete and the horizon is infinite. Monitoring is perfect.

We extend the standard notion of a dynamic stochastic game in two ways. First, we allow for a variable period length $\Delta > 0$. Time is thus $t = 0, \Delta, 2\Delta, \ldots$. Each player strives to maximize the expected net present value of his stream of payoffs and discounts future payoffs using the discount factor $\exp(-\rho\Delta)$ corresponding to the discount rate $\rho > 0$. Second, 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 $J^t \subseteq \{1, 2, \ldots, N\}$ who have the move at time t thus becomes part of the state of the system, and we refer to it as the "protocol" state to distinguish it from the familiar "physical" state $\omega^t \in \Omega$. In contrast to the physical state, 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 $a^{t=0}_{J^{t=0}} = \left(a^{t=0}_j\right)_{j \in J^{t=0}}$ simultaneously and independently from each other. Now two things happen, depending on the state $(\omega^{t=0}, J^{t=0})$ and the actions $a^{t=0}_{J^{t=0}}$. First, player i receives a payoff $u^{\Delta}_i(\omega^{t=0}, J^{t=0}, a^{t=0}_{J^{t=0}})$. 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 $\Pr^{\Delta}\left(\omega^{t=\Delta}|\omega^{t=0}, J^{t=0}, a^{t=0}_{J^{t=0}}\right)$ and that from $J^{t=0}$ to $J^{t=\Delta}$ with probability $\Pr\left(J^{t=\Delta}|J^{t=0}\right)$. 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 $a^{t=0}_{J^{t=0}}$ 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 $a^{t=\Delta}_{J^{t=\Delta}}$. Then player i receives a payoff $u^{\Delta}_i(\omega^{t=\Delta}, J^{t=\Delta}, a^{t=\Delta}_{J^{t=\Delta}})$ 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 $\{1, 2, \ldots, N\}$ 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:

Assumption 1 (Protocol of Moves) *Let $\mathcal{J}$ be a partition of $\{1, 2, \ldots, N\}$ and $\mathcal{P} = \left(\Pr\left(J'|J\right)\right)_{J,J' \in \mathcal{J}}$ 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 $< \mathcal{J}, \mathcal{P} >$ in what follows. Because $\mathcal{J}$ is a partition of $\{1, 2, \ldots, N\}$, Assumption 1 ensures that player i always has the move in conjunction with the same rivals.[2] 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.[3]

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 Tirole 1988b, Maskin and Tirole 1988a, Lagunoff and Matsui 1997), and games with random moves (Doraszelski and Judd 2007). In games with simultaneous moves, the partition is $\mathcal{J} = \{\{1, \ldots, N\}\}$ with the trivial 1×1 transition matrix $\mathcal{P}$; in games with alternating moves the partition is $\mathcal{J} = \{\{1\}, \ldots, \{N\}\}$ with the $N \times N$ transition matrix $\mathcal{P}$ with entries $\Pr(\{\mod_N(i+1)\}|\{i\}) = 1$.[4] In games with sequential moves, 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 sequential 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.

[2] In the Online Appendix, we provide an alternative formulation in which the protocol of moves is not modeled with a partition.

[3] Abreu and Gul (2000) and Caruana and Einav (2008) do not require that players move with the same frequency. In their models, hazard rates are infinite so that a player who moves more often than others does not gain much.

[4] The notation $\mod_N(x)$ refers to the modulo N congruence.

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:

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

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

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

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 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 $u_{i,j} \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathbb{R}$. *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 $q_j \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathbb{R}^+ \cup \{0\}$ *and* $p_j \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathtt{P}(\Omega)$.

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 $\mathcal{O}(\Delta^2)$.[5] We discuss below our reason for scaling by the number of elements of the partition $\mathcal{J}$.

Assumption 3, in conjunction with Assumption 2, also 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 denote the above game by $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$. We view the function $u_{i,j} \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathbb{R}$ 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 $q = (q_j(\omega, a_j))_{\omega\in\Omega, j=1,\ldots,N, a_j\in A_j(\omega)}$ and $p = (p_j(\omega'|\omega, a_j))_{\omega\in\Omega, j=1,\ldots,N, a_j\in A_j(\omega)}$.

A stationary Markovian strategy for player i is a function $\sigma_i : \Omega \to \cup_{\omega\in\Omega}\mathtt{P}(A_i(\omega))$ with $\sigma_i(\omega) \in \mathtt{P}(A_i(\omega))$ for all ω, where $\mathtt{P}(A_i(\omega))$ is the set of probability distributions over $A_i(\omega)$. Because $\mathcal{J}$ is a partition of $\{1, 2, \ldots, N\}$, 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.

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 $u_{i,j}(\omega, \sigma_j(\omega)) = \sum_{a_j\in A_j(\omega)} u_{i,j}(\omega, a_j)\sigma_j(a_j \mid \omega)$ and transition probability

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

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 Equil(Γ). This set is nonempty (Shapley 1953).

Our main interest is to compare equilibrium behavior under different protocols of moves. Assumption 1 lets us compare two models $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$ and $\overline{\Gamma} =< \Delta, \overline{\mathcal{J}}, \overline{\mathcal{P}}, u, p, q, \rho >$ that differ only in the protocol of moves by ensuring that all players move with the same frequency. Assumption 3 further ensures that players' ranks between current and future payoffs are comparable in the two models. To see this, contrast a game with simultaneous moves Γ with a game with alternating moves $\overline{\Gamma}$. 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

[5] More 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 $\Delta < \bar{\Delta}$.

alternating moves $\overline{\Gamma}$, 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 in the game with alternating moves, then this yields the payoff $|\overline{\mathcal{J}}|u_{i,j}(\omega, a_j)\Delta = Nu_{i,j}(\omega, a_j)\Delta$ and the hazard rate $|\overline{\mathcal{J}}|q_j(\omega, a_j)\Delta = Nq_j(\omega, a_j)\Delta$. The scale factor $|\mathcal{J}|$ in Assumption 3 therefore ensures that players have comparable payoffs and chances of changing the state in the game with alternating moves and the game with simultaneous moves.[6]

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.

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.*

Let $K \geq 1$ be the number of required steps and $\omega_i \in \Omega_i = \{0, 1, \ldots, K\}$ 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

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

The transition probability is[7]

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

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

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

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

[7] Note that conditional on a jump occurring we 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 $\omega_i = K$.

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

Lippman and McCardle (1987) study the continuous-time limit of this game as $\Delta \to 0$. Under suitable parameter restrictions, the equilibrium is generically unique and equilibrium behavior exhibits a pattern of increasing dominance. In standard entry games, in contrast, multiple equilibria often arise and potentially complicate empirical research (Bresnahan and Reiss 1990, Berry 1992, Quint and Einav 2005).

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 $\omega = (\omega_1, \omega_2, \ldots, \omega_N) \in \prod_{i=1}^{N} \Omega_i = \Omega$.*

Incumbent firm i earns a profit $\pi_i(\omega)$ from competing in the product market. This profit stems from the Nash equilibrium of a static product market game (price or quantity competition, depending on the application). 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.[8]

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 (2015) develop a continuous-time stochastic game in which a player is given the move*

[8] One may alternatively represent nature by an extra player 0.

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), and Calcagno, Kamada, Lovo, and Sugaya (2014) develop closely related continuous-time stochastic games with moves at random times.

Arcidiacono, Bayer, Blevins, and Ellickson (2015) 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 $l_i(\cdot \mid \omega, a_i) \in \mathtt{P}(\Omega)$.

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 probability

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

Finally, the flow payoff of player i is

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

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 (2015) clearly conform to Assumptions 2 and 3.

3 Protocol-Invariance Theorem

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

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* $u \in \mathbb{R}^{N \sum_{j=1}^{N} \sum_{\omega \in \Omega} |A_j(\omega)|}$.

The first 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 $< \mathcal{J}, \mathcal{P} >$ and $< \overline{\mathcal{J}}, \overline{\mathcal{P}} >$, and all $\varepsilon > 0$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$ and $\sigma \in$ Equil$(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$, there exists $\overline{\sigma} \in$ Equil$(< \Delta, \overline{\mathcal{J}}, \overline{\mathcal{P}}, u, p, q, \rho >)$ such that $\|\sigma - \overline{\sigma}\| < \varepsilon$.*

In words, for any Markov perfect equilibrium σ of a game with a protocol of moves $< \mathcal{J}, \mathcal{P} >$, the game with another protocol $< \overline{\mathcal{J}}, \overline{\mathcal{P}} >$ has a Markov perfect equilibrium $\overline{\sigma}$ 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 intuition for Theorem 1 is best seen by contrasting two protocols of moves. In a game 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 $\mathcal{O}(\Delta^2)$). In the game 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 complementarities between players' actions and other non-separabilities in per-period payoffs and state-to-state transitions are restricted 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 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, player i faces the same tradeoff between current and future payoffs. As a result, provided that periods are sufficiently short, the protocol of moves ceases to matter for equilibrium behavior.

To establish Theorem 1, consider a Markov perfect equilibrium $\sigma^\Delta = (\sigma_i^\Delta)_{i=1}^N$ of the separable dynamic game with noisy transitions $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$. 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 is

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

where $\sigma_J^\Delta(\omega) = \left(\sigma_j^\Delta(\omega)\right)_{j \in J}$. Under Assumptions 2 and 3 this becomes

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

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

Let $V^\Delta = \left(V_i^\Delta\right)_{i=1}^N$ be the profile of value functions corresponding to the Markov perfect equilibrium σ^Δ. Consider a family $(\sigma^\Delta, V^\Delta)$ indexed by the period length Δ. Assuming that $(\sigma^\Delta, V^\Delta) \to (\sigma^0, V^0)$ (where

convergence is possibly through a subsequence Δ_n) and taking the limit of equation (3.1) as $\Delta \to 0$, 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 $J, J' \in \mathcal{J}$.[9] This means that in equilibrium the continuation value of player i is almost 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 $V_i^0 \colon \Omega \to \mathbb{R}$ 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 $\Delta \to 0$.

Equation (3.1) can equivalently be written as

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

Summing this equation for all $J \in \mathcal{J}$ yields

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

where we use the fact that, under Assumption 1, $\sum_{J \in \mathcal{J}} \Pr(J'|J) = 1$.[10] Taking the limit as $\Delta \to 0$ yields the continuous-time Bellman equation

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

where the last inequality 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

[9] The vector $y = (1, \ldots, 1)'$ is always a right eigenvector since $\mathcal{P}$ is stochastic. Since $\mathcal{P}$ is irreducible, the Perron-Frobenious 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.

[10] Recall 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.

$< \mathcal{J}, \mathcal{P} >$ used to pass from discrete to continuous time.

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

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

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

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

Condition (3.3) is again independent of the protocol of moves. It formalizes the intuition that player i faces the same tradeoff between current and future payoffs under any protocol of moves $< \mathcal{J}, \mathcal{P} >$ and that this tradeoff is not directly affected by his rivals' actions. Holding fixed the value of continued play, the strategic situation that the N players face in a given state is thus akin to N independent optimization problems.

Conditions (3.2) and (3.3) are the limit as $\Delta \to 0$ 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 $\text{Equil}^0(< u, p, q, \rho >)$. This set does not depend on the protocol of moves $< \mathcal{J}, \mathcal{P} >$ used to pass to the limit.

We summarize the above discussion in a lemma:

Lemma 1 *Consider a sequence* $\left(\sigma^{\Delta}\right)$ *with* $\sigma^{\Delta} \in \text{Equil}(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$. *If* $\sigma^{\Delta} \to \sigma^0$, *then* $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$.

Unfortunately, Theorem 1 cannot be established by simply taking the limit of the equilibrium conditions as $\Delta \to 0$ because conditions (3.2) and (3.3) may admit multiple solutions. In other words, Theorem 1 is not implied by Lemma 1 when the limit system has several solutions. To overcome this difficulty, we use tools from differential topology to analyze the limit conditions. We first restrict attention to solutions $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$ that are *regular*. The formal definition of regularity is in the Appendix; here we just note that σ^0 is regular if it is strict, i.e., if the maximization problem in condition (3.3) admits a unique solution. Intuitively, a 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. The key technical point is that for almost all flow payoffs u, the restriction to regular solutions is without loss of generality.

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

The proof of Lemma 2 draws on ideas in Harsanyi (1973a, 1973b) and Doraszelski and Escobar (2010). We note that the results in Doraszelski and Escobar (2010) do not directly apply because 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 Doraszelski and Escobar (2010).

Lemmas 1 and 2 combine to yield Theorem 1. Our proofs also show that $\text{Equil}^0(< u, p, q, \rho >)$ 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.

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, 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 (Serfozo 1979), is adapted from single-agent decision problems. Fix any $B > N \max_{j=1,\ldots,N,\omega\in\Omega,a_j\in A_j(\omega)} q_j(\omega, a_j)$. Define the per-period payoff $\widetilde{u}_{i,j}(\omega, a_j) = \frac{N}{\rho+B} u_{i,j}(\omega, a_j)$, the discount factor $\beta = \frac{B}{\rho+B} < 1$, and the transition probability

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

Note that $\widetilde{\varphi}_j(\cdot \mid \omega, a_j) \in \mathsf{P}(\Omega)$ by construction of B. Now formulate a dynamic stochastic game with random moves in which in any period one player $j \in \{1, \ldots, N\}$ 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 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 $\text{Equil}^R\left(<\tilde{u},\widetilde{\varphi},\beta>\right)=\text{Equil}^0\left(<u,p,q,\rho>\right)$.

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

$$V_i(\omega)=\sum_{j=1}^{N}\frac{1}{N}\left(\widetilde{u}_{i,j}(\omega,\sigma_j(\omega))+\beta\sum_{\omega'\in\Omega}V_i(\omega)\widetilde{\varphi}_j(\omega'\mid\omega,\sigma_j(\omega))\right) \tag{3.4}$$

and

$$\sigma_i(a_i\mid\omega)>0\Rightarrow a_i\in\arg\max_{\tilde{a}_i\in A_i(\omega)}\widetilde{u}_{i,i}(\omega,\tilde{a}_i)+\beta\sum_{\omega'\in\Omega}V_i(\omega')\widetilde{\varphi}_i(\omega'\mid\omega,\tilde{a}_i). \tag{3.5}$$

These conditions can be equivalently written as

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

and

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

Rearranging terms, the limit conditions (3.2) and (3.3) are therefore identical to the equilibrium conditions (3.4) and (3.5) for the dynamic stochastic game with random moves constructed above.

While dynamic stochastic games with random moves are sparsely used, several important papers study dynamic stochastic games with alternating moves. For example, Maskin and Tirole (1988a) explore a dynamic Bertrand game with alternating moves and show how Edgeworth cycles can arise. Lagunoff and Matsui (1997) show how players can coordinate on the efficient outcome in a dynamic 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.[11] 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 4 (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*

[11] As Lagunoff and Matsui (1997) point out, what matters is that moves are asynchronous "rather than the specific structure of asynchronous choice" (p. 1473).

example with non-separable per-period payoffs is inspired by Lagunoff and Matsui (1997) and Wen (2002). In the Online Appendix we present a closely related example with non-separable state-to-state transitions.

Consider a coordination game with the following payoff matrix:

	E	O
E	$2,2$	$0,0$
O	$0,0$	$1,1$

Denote by $b(a) = (b_1(a), b_2(a))$ *the payoff profile given the action profile* $a = (a_1, a_2) \in \{E, O\}^2$.

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* $\sigma_1(E) = \sigma_2(E) = 1$ *and* $\tilde{\sigma}_1(O) = \sigma_2(O) = 1$.

In the game with alternating moves, in violation of Assumption 3 the per-period payoff of player i *is* $u_i^\Delta(\omega, \{1\}, a_1) = 0$ *and* $u_i^\Delta(\omega, \{2\}, a) = 2b_i(a)\Delta$, *meaning that payoffs "materialize" after player 2 moves. 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, *the unique Markov perfect equilibrium is* $\sigma_1(E) = 1$ *and* $\sigma_2(a_1 \mid a_1) = 1$ *for* $a_1 \in \{E, O\}$.

Example 5 (Noisy Transitions) *Consider the entry game in Example 1 with* $K = 1$. *We further restrict* $b_i < 0$ *so that a duopolist incurs a loss.*

We change Example 1 by assuming that if firm i *takes action* $a_i = 1$, *then its state changes for sure from* $\omega_i = 0$ *to* $\omega_i' = 1$. *Irrespective of the protocol of moves, given a set of players* $J \subseteq \{1,2\}$ *who have the move, the transition probability takes the form*

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

and does not satisfy Assumption 2. In the game with alternating moves, if the period length 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.*[12]

[12] Iskhakov, Rust, and Schjerning (2016) study an investment game that satisfies Assumption 3. Contrary to Assumption 2, however, by investing a firm can guarantee increasing its capital stock. Iskhakov, Rust, and Schjerning (2016) show that their game with alternating moves has a unique Markov perfect equilibrium whereas they find a tremendous number of equilibria in the game with simultaneous moves. There is also a well-known literature drawing 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.

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 $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$, 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 periods of interactions between players. If $T = 0$, then we recover the definition of a stationary Markovian strategy in Section 2.

Example 6 (Markov Perfect Equilibrium) *Consider a public contribution 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*

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

and the discount rate is ρ. Irrespective of the protocol of moves $< \mathcal{J}, \mathcal{P} >$ 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 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 $\min\{T, t/\Delta\}$ periods: $a_1^{\tilde{t}} = a_2^{\tilde{t}}$ for all $\tilde{t} \in \{t - \Delta, \ldots, t - \min\{t, \Delta T\}\}$. 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 period 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 $T - 1$ periods 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 Complexity Theorem

The limit conditions (3.2) and (3.3) are analytically more tractable than the equilibrium conditions for the separable dynamic game with noisy transitions from which they are derived. This allows us to characterize the solutions to the limit conditions in more detail and to use Theorem 1 to extend this characterization to separable dynamic games with noisy transitions and frequent moves.

The following lemma shows that the solutions to the limit conditions (3.2) and (3.3) have a remarkably simple structure.

Lemma 3 *Fix p, q, and ρ. For almost all u, all $\sigma \in \text{Equil}^0(< u, p, q, \rho >)$, all i, and all ω, $\sigma_i(\cdot \mid \omega)$ puts positive probability on at most N actions.*

The proof of Lemma 3 adapts and extends some results for so-called additive-reward, additive-transition dynamic stochastic games (Raghavan, Tijs, and Vrieze 1985). While Lemma 3 bounds the complexity of mixed strategies, in the Online Appendix we provide an example showing that the limit conditions (3.2) and (3.3) in general may not admit a solution in pure strategies.

The proof of Lemma 3 proceeds by contradiction. Note that a player's strategy determines his rivals' expected continuation values in a given state as a convex combination of the finite number of the continuation values associated with the actions that the player chooses with positive probability. Because of Assumption 3, some actions can be dispensed with if the player randomizes over more than N actions without affecting his rivals' expected continuation values. We can therefore construct a new solution to the limit conditions (3.2) and (3.3) which puts positive probability on a smaller number of actions. But this new solution cannot be regular because at least one player uses some but not all of his best replies. Lemma 3 follows by noting that generically all solutions are regular, as established in Section 3.

Lemma 3 and Theorem 1 combine to yield the second main result of the paper:

Theorem 2 (Complexity Theorem) *Fix p, q, and ρ. For almost all u, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, for any $\sigma^\Delta \in \text{Equil}(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$, all i, and all ω, $\sigma_i^\Delta(\cdot \mid \omega)$ puts positive probability on at most N actions.*

We note that Theorem 2 applies although the separable dynamic game with noisy transitions $< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$ may admit complementarities between players' actions and other non-separabilities in the higher-order term $\mathcal{O}(\Delta^2)$ that make a direct proof difficult. This again shows the usefulness of the limit conditions (3.2) and (3.3).[13]

[13] Haller and Lagunoff (2010) show that in two-player repeated games with alternating moves, generically no Markov perfect equilibrium is perfectly mixed. Theorem 2 shows that irrespective of the protocol of moves no Markov perfect equilibrium of separable dynamic games with noisy transitions is perfectly mixed only when $|A_i(\omega)| \geq N + 1$ for some i and some ω.

5 Applications and Extensions

5.1 Justification of Markov Perfect Equilibria

We apply and extend our main results to provide a new justification for focusing on Markov perfect equilibria in dynamic stochastic games. We argue that the set of Markov perfect equilibrium payoffs coincides with the set of payoffs that can be attained with more general equilibrium concepts under arbitrary protocols of moves.

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 $< \mathcal{J}, \mathcal{P} >$ of a separable dynamic game with noisy transitions $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$, 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 say that a payoff profile $v \in \mathbb{R}^N$ is *limit protocol-invariant* if for all $\varepsilon > 0$ and for all $< \mathcal{J}, \mathcal{P} >$, there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$ there exists $w \in \text{Payoffs}^F(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$ such that $\|v - w\| < \varepsilon$. In words, the payoff profile v is limit protocol-invariant if it can be almost attained by a strict subgame perfect equilibrium in finite memory strategies irrespective of the protocol of moves provided that periods are sufficiently short.

Define $\text{Payoffs}^0(< u, p, q, \rho >) \subseteq \mathbb{R}^N$ to be the set of payoff profiles associated with solutions in pure strategy profiles to the limit conditions (3.2) and (3.3). Formally

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

where $\omega^{t=0} \in \Omega$ is the initial state of the game. These are all payoff profiles that can be attained in a separable dynamic game with noisy transitions as moves become arbitrarily frequent. From Theorem 1, for almost all u, a payoff profile $v \in \text{Payoffs}^0(< u, p, q, \rho >)$ is limit protocol-invariant. The following proposition shows that this characterization is actually tight:

Proposition 2 (Tight Bound) *Fix p, q, and ρ. For almost all u, $v \in \mathbb{R}^N$ is limit protocol-invariant if and only if $v \in \text{Payoffs}^0(< u, p, q, \rho >)$.*

Proposition 2 provides a rationale for focusing on Markov perfect equilibria in separable dynamic games with noisy transitions and frequent moves: the set of equilibrium payoffs that are robust to the specification

of the protocol of moves that the modeler imposes coincides with the set of Markov perfect equilibrium payoffs. Proposition 2 complements several arguments in favor of Markov perfect equilibria (Maskin and Tirole 2001, Bhaskar and Vega-Redondo 2002, Bhaskar, Mailath, and Morris 2013, Bohren 2014).

To derive a tight bound, Proposition 2 builds on related results for games with asynchronous moves by Bhaskar and Vega-Redondo (2002) and Bhaskar, Mailath, and Morris (2009) and combines them with our Theorem 1. The proof of Proposition 2 draws on the insight from Example 6 that while 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.

It is an open question whether the set of payoff profiles that can be attained under arbitrary protocols of moves exceeds the set $\text{Payoffs}^0(< u, p, q, \rho >)$ if the restriction to strict subgame perfect equilibria in finite memory strategies is relaxed.[14] Without restrictions on strategies, the properties of the set of equilibrium payoffs as $\Delta \to 0$ are generally not well understood in the literature. The existing results consider either the limit $\Delta \to 0$ with simultaneous moves (Peski and Wiseman 2015) or the limit $\rho \to 0$ (Dutta 1995, Yoon 2001, Hörner, Sugaya, Takahashi, and Vieille 2011).

5.2 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 results have a number of implications for computing Markov perfect equilibria.

First, Doraszelski and Judd (2007) show that the computational burden can vary by orders of magnitude with the protocol of moves. For separable dynamic games with noisy transitions and frequent moves, protocol invariance justifies imposing the protocol of moves that is most convenient from a computational perspective.

Second, Assumption 3 facilitates numerically solving the limit conditions (3.2) and (3.3) that arise as we pass from discrete to continuous time. Recall that holding fixed the value of continued play, the strategic situation that the N players face in a given state is akin to N independent optimization problems. Without Assumption 3, the strategic situation is, in contrast, akin to a normal-form game with N players that is harder to solve than N independent optimization problems (see McKelvey and McLennan (1996) and the references therein). Our main results justify solving the limit conditions (3.2) and (3.3) by showing that their 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.

Third, 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. Even without Assump-

[14] In the tightly specified model in Example 6, 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.

tion 3, the techniques we develop allow us to clarify the relationship between discrete- and continuous-time stochastic games.

Consider a dynamic stochastic game with simultaneous moves and noisy transitions. The per-period payoff is $u_i^\Delta(\omega, \{1, \ldots, N\}, a) = u_i(\omega, a)\Delta + \mathcal{O}(\Delta^2)$. 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. Analogously, let $\text{Equil}^0(\langle u, p, q, \rho \rangle)$ be the set of solutions to the limit conditions (3.2) and (3.3).[15]

Proposition 3 *Fix p, q, and ρ. For almost all u, $\lim_{\Delta \to 0} \text{Equil}(\langle \Delta, u, p, q, \rho \rangle) = \text{Equil}^0(\langle u, p, q, \rho \rangle)$.*

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 3 does not carry over from simultaneous to alternating moves. We also note that with a continuum of actions, a version of Proposition 3 (and of Theorem 1) can be obtained by considering approximate equilibria as in Fudenberg and Levine (1986).

5.3 Timing and Commitment with Infinite Hazard Rates and Frequent Moves

We apply and extend our main results to shed light on the numerous examples in the literature in which the protocol of moves matters crucially for equilibrium behavior. We use the canonical model of Lagunoff and Matsui (1997) to expose a discontinuity in the set of Markov perfect equilibria as hazard rates become large and moves become frequent.

We construct a separable dynamic game with noisy transitions as follows. There are $N = 2$ players and 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 $A_i(\omega) = \Omega_i$. Consider a function $\pi \colon \Omega \to \mathbb{R}$ and define the flow payoff as

$$u_{i,j}(\omega, a_j) = \frac{1}{2}\pi(\omega).$$

This means that the flow payoffs of the players coincide. While actions do not determine payoffs, they determine transitions from one state to another. As in Example 3, each player controls a Poisson process with hazard rate $\lambda > 0$ and the distribution over the successor state is $l_i(\omega' \mid \omega, a_i) = 1$ if and only if $\omega_i' = a_i$. As in Example 3, we construct $q_i(\omega, a_i)$ and $p_i(\omega' | \omega, a_i)$ from λ and $l_i(\omega' | \omega, a_i)$.

Denote $\omega^* = \arg\max_{\omega \in \Omega} \pi(\omega)$ and assume it is unique. Use the function $\pi \colon \Omega \to \mathbb{R}$ 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

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

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

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

maximizer of the function π, ω^* is also a Nash equilibrium. An example of such game is the payoff matrix in Example 4 with $\Omega_i = \{E, O\}$ and $\pi(E, E) = 2$, $\pi(O, O) = 1$, and $\pi(\omega) = 0$ otherwise. In this example, $\omega^* = (E, E)$ and $\omega^{NE} = (O, O)$.

Our Theorem 1 shows that irrespective of the protocol of moves $< \mathcal{J}, \mathcal{P} >$, the set of Markov perfect equilibria $\text{Equil}(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$ converges to $\text{Equil}^0(< u, p, q, \rho >)$ as $\Delta \to 0$. We can further characterize the limit as $\lambda \to \infty$:

Proposition 4 *There exists a strategy profile σ^* such that*

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

Moreover, $\sigma_i^*(\omega) = \omega_i^*$ for all i and all ω.

Up to the fact that the limit as $\Delta \to 0$ is a continuous-time stochastic game, the logic of Proposition 4 follows from Lagunoff and Matsui (1997). Note that 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 solution to the joint limit as $\lambda \to \infty$ and $\Delta \to 0$ is therefore ω^*.

Proposition 4 shows that when the hazard rate is infinite, the set of Markov perfect equilibria does not depend on the protocol of moves once moves are frequent. This seemingly contradicts the results in Maskin and Tirole (1988a) and Lagunoff and Matsui (1997) that the set of Markov perfect equilibria depends critically on the protocol of moves when hazard rate are infinite.

The explanation is that there is a discontinuity in the joint limit as $\lambda \to \infty$ and $\Delta \to 0$. Consider a protocol $< \mathcal{J}^{sim}, \mathcal{P}^{sim} >$ of simultaneous moves. We impose that $\lambda\Delta = 1$ so that players determine the state with probability one when they move. It is relatively simple to see that when $\lambda\Delta = 1$

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

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

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

and the order of limits matters for equilibrium behavior.

We view the fact that the order of limits matters as a word of caution to researchers working on games with infinite hazard rates and arbitrarily frequent moves. From this perspective, the celebrated results in Maskin and Tirole (1988a, 1988b) and Lagunoff and Matsui (1997) can be seen as knife edge cases that hinge on specific ways of taking the limits. Indeed, consider a protocol $< \mathcal{J}^{alt}, \mathcal{P}^{alt} >$ of alternating moves.

Imposing $\lambda\Delta = 1/2$, a player determines his state with probability one when he moves.[16] From Theorem 1 in Lagunoff and Matsui (1997), if the period length Δ is sufficiently small, then[17]

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

This implies that

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

Hence, the limit of the set of Markov perfect equilibria as $\Delta \to 0$ keeping $\lambda\Delta$ constant depends on the protocol of moves.

6 Conclusions

The timing of decisions is an essential ingredient in modelling any strategic situation. 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 in which the protocol of moves matters crucially for equilibrium behavior, our paper is the first to present a fairly general class of dynamic models whose implications and predictions are independent of the timing of decisions and thus more robust for the purposes of applied work.

In particular, 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. Protocol invariance means that the set of Markov perfect equilibria of these games 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. We also show that the Markov perfect equilibria of separable dynamic stochastic games with noisy transitions have a remarkably simple structure.

In addition to alleviating the burden of determining the most realistic and appropriate protocol of moves, our main results 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 point to a discontinuity in the set of Markov perfect equilibria as hazard rates become large and moves become frequent, thereby shedding light on the examples in the literature in which the protocol of moves matters crucially for equilibrium behavior.

[16] Recall 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.

[17] The 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.

Appendix

This Appendix consists of three parts. Appendix A.1 provides the proof of Theorem 1, Appendix A.2 provides the proof of Theorem 2, and Appendix A.3 provides the proofs for Section 5.

A.1 Proof of Theorem 1

A.1.1 Notation and Preliminary Definitions

Enumerate the state space as $\Omega = \{\omega^1, \ldots, \omega^{|\Omega|}\}$ and the set of actions for player i as $A_i(\omega) = \left\{a_i^1, \ldots, a_i^{|A_i(\omega)|}\right\}$. 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

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

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

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

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

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

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

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

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 $r\mathbb{1} + \mathcal{P}_\sigma(Q-P)$ is strictly dominant diagonal and therefore invertible.[18] 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

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

[18] A 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.

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

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

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 $u \in \mathbb{R}^{N \sum_{j=1}^{N} \sum_{\omega \in \Omega} |A_j(\omega)|}$, the restriction to regular solutions is without loss of generality.

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

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

In light of limit condition (3.3), we interpret $\mathcal{U}_i(a_i, \omega, \sigma)$ as the objective function that player $i \in \{1, \ldots, N\}$ maximizes over $a_i \in A_i(\omega)$ given state $\omega \in \Omega$ and continuation play $\sigma \in \Sigma^\varepsilon$.

Consider $\bar{\sigma} \in \text{Equil}^0(< u, p, q, \rho >)$. Choose a_i^ω such that $\bar{\sigma}_i(a_i^\omega \mid \omega) > 0$ for all $i \in \{1, \ldots, N\}$ and all $\omega \in \Omega$. Given $a_i \neq a_i^\omega$ and $\sigma \in \Sigma^\varepsilon$, define

$$f_{i,a_i,\omega}(\sigma) = \sigma_i(a_i \mid \omega)\Big(\mathcal{U}_i(a_i, \omega, \sigma) - \mathcal{U}_i(a_i^\omega, \omega, \sigma)\Big)$$

while

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

By definition, $f(\bar{\sigma}) = 0$. In this subsection, we sometimes emphasize the dependence of f on u by writing $f(\sigma, u)$. Note that $f \colon \Sigma^\varepsilon \times \mathbb{R}^{N \sum_{\omega \in \Omega} \sum_{i=1}^{N} |A_i(\omega)|} \to \mathbb{R}^{\sum_{\omega \in \Omega} \sum_{i=1}^{N} |A_i(\omega)|}$ is continuously differentiable.

Definition 1 $\bar{\sigma} \in \text{Equil}^0(< u, p, q, \rho >)$ *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 $\sigma_i(a_i \mid \omega) > 0$ for all $i \in \{1, \ldots, N\}$, $\omega \in \Omega$, and all $a_i \in A_i(\omega)$.

Lemma 4 *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 $\sigma_i(a_i \mid \omega)$, and for the $(i, a_i^{\omega}, \omega)$ column, where we write $-\sigma_i(a_i \mid \omega)$. The function f can be expressed as

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

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

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

where X_i equals

$$\left(\begin{array}{ccc|ccc|c|ccc} \multicolumn{3}{c|}{\sigma_i(\cdot \mid \omega^1)} & \multicolumn{3}{c|}{\sigma_i(\cdot \mid \omega^2)} & \ldots & \multicolumn{3}{c}{\sigma_i(\cdot \mid \omega^{|\Omega|})} \\ \hline 1 & \ldots & 1 & 0 & \ldots & 0 & \ldots & 0 & \ldots & 0 \\ 0 & \ldots & 0 & 1 & \ldots & 1 & \ldots & 0 & \ldots & 0 \\ \vdots & & \vdots & \vdots & & \vdots & & \vdots & & \vdots \\ 0 & \ldots & 0 & 0 & \ldots & 0 & \ldots & 1 & \ldots & 1 \end{array}\right)$$

and has rank $|\Omega|$, while $Z_i = M(\sigma, i)\Big(\mathbb{1} + \frac{1}{\rho}(P - Q)\mathcal{P}_{\sigma}\Big)^{-1}$. 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 $\sigma\in\Sigma$, $i\in\{1,\ldots,N\}$, 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

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

Using this notation, $\sigma\in\text{Equil}^0(<u,p,q,\rho>)$ if and only if $C_i(\sigma,\omega)\subseteq B_i(\sigma,\omega)$ for all $i\in\{1,\ldots,N\}$. We say that $\sigma\in\text{Equil}^0(<u,p,q,\rho>)$ is *quasi-strict* if $C_i(\sigma,\omega)=B_i(\sigma,\omega)$ for all $i\in\{1,\ldots,N\}$.

Lemma 5 *For almost all $u\in\mathbb{R}^{N\sum_{j=1}^{N}\sum_{\omega\in\Omega}|A_j(\omega)|}$, any $\sigma\in\text{Equil}^0(<u,p,q,\rho>)$ is quasi-strict.*

Proof. Given $i\in\{1,\ldots,N\}$, consider correspondences $B_i^*\colon\Omega\to\cup_{\omega\in\Omega}A_i(\omega)$ and $C_i^*\colon\Omega\to\cup_{\omega\in\Omega}A_i(\omega)$, 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 $\sigma\in\text{Equil}^0(<u,p,q,\rho>)$ with best replies $B^*=(B_i^*)_{i=1}^N$ and carriers $C^*=(C_i^*)_{i=1}^N$. Formally,

$$G(B^*,C^*)=\Big\{u\mid \text{ there exists } \sigma\in\text{Equil}^0(<u,p,q,\rho>) \text{ with } B_i(\sigma,\cdot)=B_i^* \text{ and } C_i(\sigma,\cdot)=C_i^* \text{ for all } i=1,\ldots,N\Big\}.$$

Consider first $\bar\sigma\in\text{Equil}^0(<\bar u,p,q,\rho>)$ such that $B_i(\sigma,\omega)=B_i^*(\omega)$ 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 $i\in\{1,\ldots,N\}$ and all $a_i\in B_i^*(\omega)$. For all $\omega\in\Omega$ and all $i\in\{1,\ldots,N\}$, 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

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

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

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

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 $i\in\{1,\ldots,N\}$ and some $\omega\in\Omega$, $C_i^*(\omega)\subsetneq B_i^*(\omega)$. Consider the set

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

where $A(C^*)=\{\sigma\in\Sigma \mid C_i(\cdot,\omega)=C_i^*(\omega) \text{ for all } i=1,\ldots,N\}$. 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 $\bar{\sigma}\in\text{Equil}^0(<u,p,q,\rho>)$ are regular.*

Proof. From Lemma 5, 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 $B_i^*:\Omega\to A_i$, 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 $i\in\{1,\ldots,N\}$ 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 4 and the transversality theorem (see the discussion in Section 7.1.2 in Doraszelski and Escobar 2010), 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 $\sigma^0\in\text{Equil}^0(<u,p,q,\rho>)$. Let $<\mathcal{J},\mathcal{P}>$ 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 $\Gamma =< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$, write the continuation value of player $i \in \{1, \ldots, N\}$ 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 $V_i^\Delta \colon \Omega \times \mathcal{J} \to \mathbb{R}$ 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 $\sigma \in \Sigma$, $V_i^\Delta(\cdot, J, \sigma) \to V_i^0(\cdot, \sigma)$ as $\Delta \to 0$.

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

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

with

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

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

Consider the profile $(a_i^\omega)_{i=1,\ldots,N,\omega\in\Omega}$ 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 $f \colon [0,1] \times \Sigma^\varepsilon \to \mathbb{R}^{\sum_{\omega\in\Omega} \sum_{i=1}^N |A_i(\omega)|}$ such that for all $a_i \neq a_i^\omega$

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

while

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

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 6 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 $< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$, consider first $a_i \in A_i(\omega)$ and $\omega \in \Omega$ such that $\sigma^0(a_i \mid \omega) = 0$. 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 $\sigma^\Delta(a_i \mid \omega) = 0$. Next consider $a_i \in A_i(\omega)$ and $\omega \in \Omega$ such that $\sigma^0(a_i \mid \omega) > 0$. We can use (i) to deduce that $\sigma^\Delta(a_i \mid \omega) > 0$ and, since $f(\Delta, \sigma^\Delta) = 0$, $\mathcal{U}_i^\Delta(\omega, a_i, \sigma^\Delta) = \mathcal{U}_i^\Delta(\omega, a_i^\omega, \sigma^\Delta)$. All of these observations prove that whenever $\sigma^\Delta(a_i, \omega) > 0$, a_i solves $\max_{\tilde{a}_i \in A_i(\omega)} \mathcal{U}_i^\Delta(\omega, \tilde{a}_i, \sigma^\Delta)$. Therefore, σ^Δ is a Markov perfect equilibrium of the separable dynamic game with noisy transitions $< \Delta, \mathcal{J}, \mathcal{P}, u, , q, \rho >$.

Since for almost all $u \in \mathbb{R}^{N\sum_{j=1}^{N}\sum_{\omega\in\Omega}|A_j(\omega)|}$, all $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$ 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 6 (Implicit Function Theorem) *Assume $f\colon [0,1] \times \Sigma^{\varepsilon} \to \mathbb{R}^{\sum_{\omega\in\Omega}\sum_{i=1}^{N}|A_i(\omega)|}$ 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 $\varphi(\sigma, \Delta) = \sigma - [D_\sigma f(0,\sigma^0)]^{-1} f(\Delta, \sigma)$ 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 $\varphi(\cdot, \Delta)$. Note that $D_\sigma \varphi(0, \sigma^0) = 0$ 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

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

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

Define now $m(\Delta) = \max_{\{\sigma \mid \|\sigma - \sigma^0\| \leq r\}} \|\varphi(\sigma, \Delta) - \varphi(\sigma, 0)\|$. 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 $\|\sigma - \sigma^0\| \leq r$ and $\Delta < \bar{\Delta}$

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

It follows that for all $\Delta < \bar{\Delta}$, the continuous function $\varphi(\cdot, \Delta)$ 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 $\varphi(\sigma^{\Delta}, \Delta) = \sigma^{\Delta}$. ■

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 $\sigma^{\Delta} \in \text{Equil}(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$, there exists $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$ 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

$\hat{\sigma}^{\Delta} \in \text{Equil}(< \Delta, \bar{\mathcal{J}}, \bar{\mathcal{P}}, u, p, q, \rho >)$ such that $\|\bar{\sigma}^{\Delta} - \sigma^0\| < \varepsilon/2$. Taking $\bar{\Delta} = \min\{\tilde{\Delta}, \hat{\Delta}\} > 0$, Theorem 1 follows from the triangle inequality. ■

A.2 Proof of Theorem 2

It is useful to begin by establishing the following lemma:

Lemma 7 *Let $\sigma \in \text{Equil}^0(< u, p, q, \rho >)$ with corresponding profile of value functions V. Then there exists $\tilde{\sigma} \in \text{Equil}^0(< u, p, q, \rho >)$ with corresponding profile of value functions $\tilde{V}$ such that for all $\omega \in \Omega$ and all $i \in \{1, \ldots, N\}$,*

a. $\tilde{V}_i(\omega) = V_i(\omega)$;

b. $\{a_i \in A_i(\omega) \mid \tilde{\sigma}_i(a_i \mid \omega) > 0\} \subseteq \{a_i \in A_i(\omega) \mid \sigma_i(a_i \mid \omega) > 0\}$;

c. $|\{a_i \in A_i(\omega) \mid \tilde{\sigma}_i(a_i \mid \omega) > 0\}| \leq N$.

In particular, if $\{a_i \in A_i(\omega) \mid \sigma_i(a_i \mid \omega) > 0\} \geq N + 1$ for some $i \in \{1, \ldots, N\}$ and some $\omega \in \Omega$, then $\tilde{\sigma}$ is not regular.

The proof of this result extends some of the ideas in the proof of Theorem 3.1 in Raghavan, Tijs, and Vrieze (1985).

Proof of Lemma 7. Write limit condition (3.2) as

$$\rho V_i(\omega) = \sum_{j=1}^{N} \sum_{a_j \in A_j(\omega)} \left(u_{i,j}(\omega, a_j) + \sum_{\omega' \in \Omega} (V_i(\omega') - V_i(\omega))\, \varphi_j(\omega' \mid \omega, a_j) \right) \sigma_j(a_j \mid \omega).$$

Fix player 1 and state $\omega \in \Omega$. Consider the finite set

$$C_1(\omega) = \Big\{ v = (v_2, \ldots, v_N) \in \mathbb{R}^{N-1} \mid \text{ there exists } a_1 \in A_1(\omega) \text{ such that}$$
$$v_i = u_{i,1}(\omega, a_1) + \sum_{\omega' \in \Omega} (V_i(\omega') - V_i(\omega))\, \varphi_1(\omega' \mid \omega, a_1) \text{ for all } i = 2, \ldots, N \Big\}.$$

Observe that the vector $\tilde{v} = (\tilde{v}_2, \ldots, \tilde{v}_N) \in \mathbb{R}^{N-1}$ defined by

$$\tilde{v}_i = \sum_{a_1 \in A_1(\omega)} \Big(u_{i,1}(\omega, a_1) + \sum_{\omega' \in \Omega} (V_i(\omega') - V_i(\omega))\, \varphi_j(\omega' \mid \omega, a_1) \Big) \sigma_1(a_1 \mid \omega)$$

belongs to the convex hull of $C_1(\omega)$. Caratheodory's theorem (Theorem 17.1 in Rockafellar 1970) implies that there exists a distribution $\tilde{\sigma}_1(\cdot \mid \omega) \in \mathtt{P}(A_1(\omega))$ with $\{a_1 \in A_1(\omega) \mid \tilde{\sigma}_1(a_1 \mid \omega) > 0\} \subseteq \{a_1 \in A_1(\omega) \mid \sigma_1(a_1 \mid \omega) > 0\}$ such that

$$\tilde{v} = \Big(\sum_{a_1 \in A_1(\omega)} u_{i,1}(\omega, a_1) + \sum_{\omega' \in \Omega} (V_i(\omega') - V_i(\omega))\, \varphi_j(\omega' \mid \omega, a_1) \mid i = 2, \ldots, N \Big) \tilde{\sigma}_1(a_1 \mid \omega) \in \mathbb{R}^{N-1}$$

and $|\{a_1 \in A_1(\omega) \mid \tilde{\sigma}_1(a_1 \mid \omega) > 0\}| \leq N$. Do this procedure for all players (not just player 1) and obtain a strategy profile $\tilde{\sigma} = (\tilde{\sigma}_i)_{i=1}^N$ that satisfies (b) and (c) for all $i = 1, \ldots, N$ and all $\omega \in \Omega$. Also note that, by construction, V_i is the value function for player i and $\sigma_i(\cdot, \omega)$ is putting positive weight only on actions $a_i \in A_i(\omega)$ that maximize

$$u_{i,i}(\omega, a_i) + \sum_{\omega' \neq \omega} (V_i(\omega') - V_i(\omega))\, \varphi_i(\omega' \mid \omega, a_i).$$

We therefore conclude that $(\tilde{\sigma}, V)$ is a solution to limit condition (3.2) and (3.3) having properties (a), (b), and (c). ■

Proof of Lemma 3. Proposition 5 shows that all solutions are regular for almost all flow payoffs $u \in \mathbb{R}^{N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$. Take such flow payoffs u and a solution σ. If for some $i = 1, \ldots, N$ and some $\omega \in \Omega$, $|\{a_i \in A_i(\omega) \mid \sigma_i(a_i \mid \omega)\}| \geq N + 1$, we could use Lemma 7 to find the non-regular solution $\tilde{\sigma}$. ■

Proof of Theorem 2. From the proof of Theorem 1, for almost all flow payoffs $u \in \mathbb{R}^{N \sum_{j=1}^N \sum_{\omega \in \Omega} |A_j(\omega)|}$, all $\varepsilon > 0$ and all protocols of moves $< \mathcal{J}, \mathcal{P} >$, we can find $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, the carrier of $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$ and $\sigma^\Delta \in \text{Equil}(< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$ coincide and $\|\sigma^\Delta - \sigma^0\| < \varepsilon$. Theorem 2 thus follows from Lemma 3. ■

A.3 Proofs for Section 5

Proof of Proposition 2. Given a limit protocol-invariant vector $v \in \mathbb{R}^N$, take the alternating-move protocol $< \mathcal{J}^{alt}, \mathcal{P}^{alt} >$ and a sequence $\Delta_n \to 0$ such that $w^n \in \text{Payoffs}^F(< \Delta_n, \mathcal{J}^{alt}, \mathcal{P}^{alt}, u, p, q, \rho >)$ and $\|v - w^n\| < 1/n$. In particular, there exists a sequence (σ^n) of strict subgame perfect equilibria in finite memory strategies with corresponding payoff profile $w^n \in \mathbb{R}^N$. From Bhaskar, Mailath, and Morris (2013), $\sigma^n \in \text{Equil}(< \Delta_n, \mathcal{J}^{alt}, \mathcal{P}^{alt}, u, p, q, \rho >)$. Because $\sigma^n \to \sigma^0$ (possibly through a subsequence), from Lemma 1, $\sigma^0 \in \text{Equil}^0(< u, p, q, \rho >)$ and therefore $v \in \text{Payoffs}^0(< u, p, q, \rho >)$. The converse follows from Theorem 1. ■

Proof Sketch of Proposition 3. 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. ■

Proof of Proposition 4. The results follow directly from the analysis in Section 3 in Lagunoff and Matsui (1997) and Theorem 1 in their working paper. ■

References

Abreu, D., and F. Gul (2000): "Bargaining and Reputation," *Econometrica*, 68(1), 85–117.

Acemoglu, D., and J. Robinson (2001): "A Theory of Political Transition," *American Economic Review*, 91(4), 938–963.

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," Discussion paper, Duke University.

Ambrus, A., and S. Lu (2015): "A Continuous-Time Model of Multilateral Bargaining," *American Economic Journal: Microeconomics*, 7, 208–249.

Arcidiacono, P., P. Bayer, J. Blevins, and P. Ellickson (2015): "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition," Discussion paper, Duke University.

Bergemann, D., and J. Välimäki (1996): "Learning and Strategic Pricing," *Econometrica*, pp. 1125–1149.

Bernheim, D., and D. Ray (1989): "Markov Perfect Equilibria in Altruistic Growth Economies with Production Uncertainty," *Journal of Economic Theory*, 47(1), 195–202.

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," *The 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," *The Review of Economic Studies*, 57, 531–553.

Cahuc, P., and H. Kempf (1997): "Alternative Time Patterns of Decisions and Dynamic Strategic Interactions," *The Economic Journal*, 107(445), 1728–1741.

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," *The 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.

Cournot, A. (1838): *Recherches sur les Principes Mathmatiques de la Thorie des Richesses*. Hachette, Paris.

Cyert, R., and M. DeGroot (1970): "Multiperiod 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 (2007): "Dynamic Stochastic Games with Sequential State-to-State Transitions," Working paper, Stanford University.

——— (2012): "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games," *Quantitative Economics*, 3(1), 53–93.

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. Satterthwaite (2010): "Computable Markov Perfect Industry Dynamics," *The 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.

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., 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," *The Review of Economic Studies*, 52(3), 383–401.

Gowrisankaran, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers," *RAND Journal of Economics*, 30(1), 56–83.

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.

——— (2010): "Markov Perfect Equilibria in Repeated Asynchronous Choice Games," *Journal of Mathematical Economics*, 46(6), 1103–1114.

Hanig, M. (1986): "Differential Gaming Models of Oligopoly," Ph.D. thesis, Massachusetts Institute of Technology.

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 $\delta \to 1$ and a Folk Theorem," *Econometrica*, 79(4), 1277–1318.

Iijima, R., and A. Kasahara (2015): "Gradual Adjustment and Equilibrium Uniqueness under Noisy Monitoring," Discussion paper, Harvard University.

Iskhakov, F., J. Rust, and B. Schjerning (2016): "Recursive Lexicographical Search: Finding All Markov Perfect Equilibria of Finite State Directional Dynamic Games," *The Review of Economic Studies*, 83(2), 658–703.

Kalai, E. (2004): "Large Robust Games," *Econometrica*, 72(6), 1631–1665.

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.

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.

McKelvey, R., and A. McLennan (1996): "Computation of Equilibria in Finite Games," in *Handbook of Computational Economics*, ed. by H. Amman, D. Kendrick, and J. Rust, pp. 87–142. North-Holland, Amsterdam.

Noel, M. (2008): "Edgeworth Cycles and Focal Prices: Computational Dynamic Markov Equilibria," *Journal of Economics and Management Strategy*, 17(2), 345–377.

Peski, M., and T. Wiseman (2015): "A Folk Theorem for Stochastic Games with Infrequent State Changes," *Theoretical Economics*, 10(1), 131–173.

Quint, D., and L. Einav (2005): "Efficient Entry," *Economics Letters*, 88(2), 278–283.

Raghavan, T., S. Tijs, and O. Vrieze (1985): "On Stochastic Games with Additive Reward and Transition Structure," *Journal of Optimization Theory and Applications*, 47(4), 451–464.

Reinganum, J. (1982): "A Dynamic Game of R and D: Patent Protection and Competitive Behavior," *Econometrica*, 50, 671–688.

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.

Rockafellar, R. T. (1970): *Convex Analysis*. Princeton University Press.

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.

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," *The 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.

Wen, Q. (2002): "A Folk Theorem for Repeated Sequential Games," *The Review of Economic Studies*, 69(2), 493–512.

Yoon, K. (2001): "A Folk Theorem for Asynchronously Repeated Games," *Econometrica*, 69(1), 191–200.

Online Appendix

This Appendix consists of three parts. Appendix OA.1 provides a counterexample complementing Example 4, Appendix OA.2 shows that the limit conditions (3.2) and (3.3) may not admit a solution in pure strategies, and Appendix OA.3 extends our notion of a protocol of moves.

OA.1 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 $p(1 \mid 0, (a_1, a_2)) = 1$. 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 4), the unique Markov perfect equilibrium is $\sigma_1^* = 1$, $\sigma_2^*(a_1) = a_1$, and the state eventually jumps to $\omega' = 1$.

OA.2 Non-Existence of Solution in Pure Strategies

Consider a separable dynamic game with noisy transitions, $N = 2$ players, $\Omega = \{1, 2\}$, $A_i(\omega) = \{1, 2\}$ if $\omega = i$, and $A_i(\omega) = \{1\}$ if $\omega \neq i$. This means that player i makes a nontrivial decision only when $\omega = i$. Flow payoffs are $u_{i,i}(\omega, a_i) = 0$ if $\omega = i$ for all $a_i \in \{1, 2\}$, while

$$u_{1,2}(a_2, 2) = \begin{cases} 10 & \text{if } a_2 = 1, \\ -10 & \text{if } a_2 = 2 \end{cases}$$

and

$$u_{2,1}(a_1, 1) = \begin{cases} -10 & \text{if } a_1 = 1, \\ 10 & \text{if } a_1 = 2. \end{cases}$$

Hence, the flow payoff of player i is 0 when $\omega = i$, but his decision determines whether the flow payoff of player $-i$ is 10 or -10. Transition probabilities are determined as in Example 3 with $\lambda = 1$ and

$$l_1(1 \mid 1, a_1) \begin{cases} 0 & \text{if } a_1 = 1, \\ 1 & \text{if } a_1 = 2 \end{cases}$$

and

$$l_2(1 \mid 2, a_2) = \begin{cases} 1 & \text{if } a_2 = 1, \\ 0 & \text{if } a_2 = 2, \end{cases}$$

while $l_i(i \mid i, a_{-i}) = 1$. This means that in state $\omega = i$, player i (and only player i) determines a probability distribution over the successor state ω'. (Note that given λ, l_1, and l_2, we can construct p and q as we did in Example 3.)

We show that the limit conditions (3.2) and (3.3) do not admit a solution in pure strategies. The intuition is similar to the non-existence of a Nash equilibrium in pure strategies in matching pennies. Consider a solution $\sigma^* \in \text{Equil}^0(< u, p, q, \rho >)$ in pure strategies. If $\sigma_1^*(\cdot \mid 1) = (1, 0)$, then it must be the case that player 2 is choosing $\sigma_2^*(1 \mid 2) = 1$ for otherwise player 1 makes a loss in state $\omega = 2$ while he can secure 0 by playing $a_1 = 2$. But this would mean that player 2 is willing to make a loss in state $\omega = 1$, while he can secure 0 by playing $a_2 = 2$. Similarly, it cannot be that $\sigma_1^*(\cdot \mid 1) = (0, 1)$. Thus, the limit conditions (3.2) and (3.3) do not admit a solution in pure strategies. They do, however, admit a solution in mixed strategies in which player i chooses $\sigma_i(\cdot \mid \omega) = (1/2, 1/2)$ when $\omega = i$.[19]

OA.3 Non-Partition Protocols of Moves

We relax Assumption 1 and assume that $\mathcal{J}$ is not a partition of the set of players but contains subsets $J \subseteq \{1, \ldots, N\}$ such that for all $i = 1, \ldots, N$, 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 $i = 1, \ldots, N$. As before, there is an irreducible Markov chain $\mathcal{P}$ defined on $\mathcal{J}$ that has a unique stationary distribution that is uniform on $\mathcal{J}$. Any such pair $< \mathcal{J}, \mathcal{P} >$ is an *extended protocol of moves*.

In the separable dynamic game with noisy transitions and extended protocol of moves, per-period payoffs are specified 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 $q_j \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathbb{R}^+ \cup \{0\}$ and $p_j \colon \{(\omega, a_j) \mid a_j \in A_j(\omega)\} \to \mathtt{P}(\Omega)$. The remaining aspects of the model are unchanged.

In this model, the identity of the players that move in conjunction with player i is a state variable. Thus, a Markovian strategy for player i is a function $\sigma_i \colon \Omega \times \{J \in \mathcal{J} \mid i \in J\} \to \cup_{\omega \in \Omega} \mathtt{P}(A_i(\omega))$. Overloading notation,

[19] This example can be easily adapted to show that so-called additive-reward, additive-transition dynamic stochastic games may not admit a Markov perfect equilibrium in pure strategies (Raghavan, Tijs, and Vrieze 1985). It is much simpler than other examples in the literature.

we denote Equil($< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$) the set of Markov perfect equilibria of the separable dynamic game with noisy transitions and extended protocol of moves $< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$. We say that a Markov perfect equilibrium σ of the game $< \Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >$ is *simple* if $\sigma_i(a_i \mid \omega, J) = \sigma_i(a_i \mid \omega, \tilde{J})$ for all $i = 1, \ldots, N$, $\omega \in \Omega$, $a_i \in A_i$, and all $J, \tilde{J} \in \mathcal{J}$. In this case, we write $\sigma_i(a_i \mid \omega)$.

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

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

In contrast to our main results, Proposition 6 assumes that Equil$^0(< u, p, q, \rho >)$ has only pure solutions. When mixed solutions are considered, the limit conditions could have a continuum of solutions because players could use the payoff-irrelevant realization of J to randomize over actions. As a result, our differential topology tools cannot be directly applied.

Proof. Consider a sequence $\left(\sigma^\Delta\right)$ with $\sigma^\Delta \in$ Equil$(\Delta, \mathcal{J}, \mathcal{P}, u, p, q, \rho >)$ and $\sigma^\Delta \to \sigma^0$ (possibly through a subsequence) as $\Delta \to 0$. 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

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

and

$$\sigma_i^0(a_i \mid \omega, J) > 0 \Rightarrow a_i \in \arg \max_{\tilde{a}_i \in A_i(\omega)} u_{i,i}(\omega, \tilde{a}_i) + \sum_{\omega' \neq \omega} \left(V_i^0(\omega') - V_i^0(\omega) \right) \varphi_i(\omega' \mid \omega, \tilde{a}_i). \tag{OA.0.1}$$

Define $\tilde{\sigma}_j(\cdot \mid \omega) = \frac{1}{\kappa} \sum_{J \in \mathcal{J} \mid j \in J} \sigma(\cdot \mid \omega, J)$ for all $i = 1, \ldots, N$ 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 all solutions to the limit conditions are in pure strategies, $\tilde{\sigma}$ must be a pure strategy solution. Since such a solution is strict, the maximization problem in equation (OA.0.1) has a unique solution and therefore $\sigma_i^0(a_i \mid \omega, J)$ does not depend on J and σ^0 is simple. In particular, $\sigma^0 \in$ Equil$^0(u, p, q, \rho)$ and therefore there exists $\bar{\Delta} > 0$ such that for all $\Delta < \bar{\Delta}$, $\sigma^\Delta \in$ Equil$^0(< u, p, q, \rho >)$. To see the converse, note that Equil$^0(< u, p, q, \rho >)$ 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 extended protocol of moves 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. ■