

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Bohren, J. Aislinn

Article Persistence in a dynamic moral hazard game

Theoretical Economics

Provided in Cooperation with: The Econometric Society

Suggested Citation: Bohren, J. Aislinn (2024) : Persistence in a dynamic moral hazard game, Theoretical Economics, ISSN 1555-7561, The Econometric Society, New Haven, CT, Vol. 19, Iss. 1, pp. 449-498, https://doi.org/10.3982/TE2680

This Version is available at: https://hdl.handle.net/10419/296464

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.



WWW.ECONSTOR.EU

https://creativecommons.org/licenses/by-nc/4.0/

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Persistence in a dynamic moral hazard game

J. AISLINN BOHREN Department of Economics, University of Pennsylvania

This paper explores how the persistence of past choices creates incentives in a continuous time stochastic game involving a large player (e.g., a firm) and a sequence of small players (e.g., customers). The large player faces moral hazard and her actions are distorted by a Brownian motion. Persistence refers to how actions impact a payoff-relevant state variable (e.g., product quality depends on past investment). I characterize actions and payoffs in Markov perfect equilibria (MPE) for a fixed discount rate, show that the perfect public equilibrium (PPE) payoff set is the convex hull of the MPE payoff set, and derive sufficient conditions for a MPE to be the *unique* PPE. Persistence can serve as an effective channel for intertemporal incentives in a setting where traditional channels fail. Applications to persistent product quality and policy targeting demonstrate the impact of persistence on equilibrium behavior.

Keywords. Continuous time games, stochastic games, moral hazard. JEL classification. C73, L1.

1. INTRODUCTION

This paper studies how the persistence of past choices can be used to create incentives in a continuous time stochastic game in which a large player interacts with a sequence of small players. Persistence refers to the impact that actions have on a payoff-relevant state variable, such as a worker's rating, a firm's product quality, or a government's key economic variables. It can capture exogenous features of the environment, such as how past investment influences current quality or how past policy choices map into the current level of an economic variable. It can also capture endogenous design choices, such as how a rating system aggregates past reviews and rewards a worker based on her rating. The large player faces moral hazard and her past actions are not perfectly observed by consumers: they are distorted by a Brownian motion. Incentives can depend on the noisy signal of action choices as well as on how persistence influences future payoffs through the impact that actions have on the state. The goal of this paper is to determine whether and how persistence strengthens incentives to overcome moral hazard.

J. Aislinn Bohren: abohren@sas.upenn.edu

Earlier versions of this paper were circulated under the titles "Stochastic games in continuous time: Persistent actions in long-run relationships" and "Using persistence to generate incentives in a dynamic moral hazard problem." I thank Simon Board, Matt Elliott, Jeff Ely, Ben Golub, Alex Imas, Bart Lipman, David Miller, George Mailath, Markus Mobius, Paul Niehaus, Andrew Postlewaite, Yuliy Sannikov, Andy Skrzypacz, Joel Sobel, Jeroen Swinkels, Joel Watson, Alex Wolitzky, and especially S. Nageeb Ali for useful comments. I also thank numerous seminar participants for helpful feedback.

^{© 2024} The Author. Licensed under the Creative Commons Attribution-NonCommercial License 4.0. Available at https://econtheory.org. https://doi.org/10.3982/TE2680

The framework captures many economic settings in which past choices shape key features of current and future interactions. For example, a worker's rating on a platform depends on the quality of service she has provided to previous customers. She may be rewarded for earning a good rating and punished for poor performance. This provides an incentive for her to earn and maintain a good rating. Similarly, a firm's ability to make a high quality product is a function not only of its effort today, but also its past investments in developing technology and training its workforce. Quality today is linked to a firm's future quality, in that customers experience similar quality across time due to the persistence of investment. When customers are willing to pay a higher price or buy a larger quantity of a high quality good, persistence provides an incentive for the firm to invest in developing a high quality product. Finally, a government's success in reaching the target level for an economic variable depends on both past and current policy choices. When past policy choices impact the future value of an economic variable, the government may be willing to undertake more costly actions today, since the benefit of such actions continue to accrue in future periods.

I study perfect public equilibria (PPE) in this framework; that is, equilibria in which strategies depend only on public information. I establish that the PPE payoff set is equal to the convex hull of the Markov perfect equilibrium (MPE) payoff set. In a MPE, equilibrium actions and payoffs depend only on the payoff-relevant components of the gamein this case, the observable state. Any paths of information that lead to the same current state prescribe the same continuation play. In contrast, a PPE can depend on past information in an arbitrary way. The intuition for this result stems from the type of incentives that are possible in games with small players and Brownian information. In a stochastic game, dynamic incentives can either be *informational*—signals are used to coordinate future equilibrium play—or *structural*—actions impact the structure of future interactions through their impact on the state, including both the state's direct impact on future feasible payoffs and its indirect impact through its effect on future equilibrium play. There are two main forms of informational incentives: burning value, where incentives are created by the threat of switching to an inefficient action profile, and transferring continuation payoffs tangent to the set of equilibrium payoffs. It is not possible to provide incentives with transfers when facing small players, and Brownian information is too noisy to create effective incentives via value-burning (Sannikov and Skrzypacz (2010)). Therefore, any nontrivial incentives in games with small players and Brownian information must be structural. This is precisely the channel for incentives in a MPE, as informational channels are precluded by definition.¹

In establishing this result, I characterize equilibrium payoffs and actions in MPE for a fixed discount rate. This characterization yields sharp insights. It shows that whether persistence allows the large player to overcome moral hazard depends on the marginal

¹In earlier related work, Faingold and Sannikov (2011) establish a similar result when small players have incomplete information about the large player's type and the state is the belief that the large player is committed to choosing a certain action. Stochastic games with multiple large players, Brownian information, and a failure of identifiability will likely have a similar equilibrium characterization to this paper, as such games face similar issues with informational incentives (Abreu, Milgrom, and Pearce (1991), Sannikov and Skrzypacz (2007)).

impact of its action on the state and how sensitive the continuation payoff is to changes in the state. In contrast to a folk theorem, it determines what type of equilibria one expects to emerge and what pattern of behavior generates a given payoff. It shows how the dynamics of behavior depend on observable outcomes (e.g., a restaurant rating or an economic variable) and how incentives and payoffs depend on key parameters of the model (e.g., the depreciation rate of investment). The characterization of the continuation payoff captures both the direct and equilibrium channels for structural incentives. For example, when consumers have observed a given level of quality in the recent past, their willingness to pay for a product depends on both the persistence of this quality the direct channel—as well as their belief about how quality influences the firm's current investment choice—the equilibrium channel. The interaction of these two channels can significantly strengthen or dampen incentives, depending on the structure of the game.

The second main result determines when a MPE emerges as the *unique* PPE. This result relies on determining when there is a unique MPE in the class of Markov equilibria. When there is a unique MPE, the result described above establishes that this will also be the unique PPE. Uniqueness depends on incentives as the state approaches the boundary of the state space. If boundary incentives are unique, e.g., it is possible to sustain a unique equilibrium action profile and payoff at the boundary, then from the MPE characterization, incentives must also be unique on the interior of the state space. I present sufficient conditions for uniqueness in two cases: (i) an unbounded state space and (ii) a bounded state space. In case (i), these conditions rule out complementarities between the direct and the equilibrium channels for incentives near the boundary, such as multiple optimal action profiles due to coordination motives. In case (ii), these conditions ensure that incentives collapse as the state approaches the boundary, which rules out the possibility of sustaining multiple equilibrium action profiles at the boundary.

Several applications illustrate how persistence can be used to create effective incentives. The first application modifies the canonical product choice setting to allow a firm's effort to have a persistent effect on the quality of its product. I show that persistence provides effective incentives for the firm to invest in building a high quality product. These incentives are present in the long run, in that the firm continues to choose a positive level of investment as the time period grows large. I also consider a variation of the product choice game in which the marginal return to quality is non-monotonic and show that this can lead to firms specializing in low or high quality. In the second application, constituents elect a board to implement a policy that targets an economic variable. The level of the variable depends on current and past decisions by the board. For example, the Federal Reserve targets an interest rate or a board of directors sets a growth target for a company. I show that the board's incentive to undertake costly intervention is strongest when the economic variable is an intermediate distance from its target; when it is far from its target, the benefit of intervention is significantly delayed, while when it is close to its target, the benefit of further intervention is small. In the final application, a government and a group of innovators invest in intellectual capital, and there is a strategic complementarity between their investments. This complementarity

gives rise to multiple Markov equilibria, including one in which neither party invests and several that sustain a positive level of investment. The equilibrium characterization in each application can be used to address important design questions. For example, a comparative static on how a firm's payoff varies with the persistence of its effort provides insight into the optimal durability for a production technology, while a comparative static on how a worker's effort varies with the persistence of its rating is useful for designing rating systems.

1.1 Literature

Recent results on repeated games between a long-run/large and short-run/small players show that the intersection of noise in monitoring and instantaneous adjustment of actions creates a genuine challenge in providing intertemporal incentives (Fudenberg and Levine (2007, 2009), Faingold and Sannikov (2011)).² In the analogue of this paper with no persistence, the large player cannot earn an equilibrium payoff above the best static Nash payoff.³ In contrast, the equilibrium characterization in this paper demonstrates that persistence can lead to effective intertemporal incentives and enable the large player to overcome moral hazard.

The literature on reputation with behavioral types is another important and well understood mechanism to overcome moral hazard in similar settings (Fudenberg and Levine (1989, 1992), Faingold and Sannikov (2011), Faingold (2020)). If consumers believe that there is a chance that the firm is committed to choosing high effort, then the firm will be able to charge a higher price for its product. Incomplete information about the firm's type creates a form of persistence, as consumers' beliefs depend on past effort choices. However, fixing a strategic firm's patience, such reputation effects vanish in the ex ante probability of behavioral types, and so the effectiveness of persistence via incomplete information requires a nontrivial fraction of behavioral types.⁴

The connection with the reputational literature motivates several key insights. First, when the firm is known to be strategic, this paper shows that other forms of persistence can overcome moral hazard.⁵ Second, in contrast to the temporary incentives in reputation models (Cripps, Mailath, and Samuelson (2004), Faingold and Sannikov (2011)), the incentives in a stochastic game persist in the long run.⁶ Finally, at a theoretical level, this

²Abreu, Milgrom, and Pearce (1991) first examined incentives in repeated games with imperfect monitoring and frequent actions. They established that shortening the period between actions has a crucial impact on the ability to structure effective incentives.

³Sannikov and Skrzypacz (2007) show that this is also the case in games between multiple long-run players in which deviations between individual players are indistinguishable.

⁴Kreps, Milgrom, Roberts, and Wilson (1982), Kreps and Wilson (1982), and Milgrom and Roberts (1982) first demonstrated that reputation, in the form of incomplete information about a player's type, has a dramatic effect on equilibrium behavior. Mailath and Samuelson (2001) show that reputational incentives can also come from a firm's desire to separate itself from an incompetent type.

⁵Along these lines, Dilmé (2019) shows that adjustment costs can help a firm overcome moral hazard by endogenously creating persistence.

⁶Long-run reputation effects are also possible in models with behavioral types when consumers cannot observe all past signals (Ekmekci (2011)) or the type of the firm is replaced over time.

paper explores the general properties of a stochastic game that has powerful intertemporal incentives. The reputational game can be viewed as a specific type of stochastic game. For instance, if instead of influencing the uncertainty about whether it is a behavioral type, a strategic firm makes a costly initial investment in a new production technology that benefits customers today and in the future, we observe similar intertemporal incentives in the resulting stochastic game.

This final point merits a closer comparison with Faingold and Sannikov (2011), who characterize the unique MPE in the stochastic game that corresponds to a continuous time reputation model. In their paper, payoffs and the evolution of the state take a specific form due to Bayesian updating. My characterization builds on the techniques in their paper to understand more generally what properties of stochastic games are needed for uniqueness of MPE and nondegenerate intertemporal incentives. I analyze a general class of stochastic games that places few restrictions on the process governing the evolution of the state and the structure of payoffs. The key technical advancement, relative to their paper, is for the case of an unbounded state space and payoff for the large player, as it requires significantly different techniques to complete the analysis.

Beyond reputation models with behavioral types, a rich literature analyzes dynamic games with a state variable in which effort is directly linked to future payoffs via the state. Ericson and Pakes (1995) were the first to analyze hidden investment and stochastic capital accumulation (the state) in a model that is similar in spirit to the quality example presented in Section 2. They study firm and industry dynamics, and establish equilibrium existence. Doraszelski and Satterthwaite (2010) modify Ericson and Pakes (1995) to guarantee the existence of a pure strategy MPE, which is computationally tractable. Neither paper establishes uniqueness, but instead focus on the dynamics associated with a particular MPE. More broadly, MPE is the workhorse solution concept across industrial organization and political economy. A comprehensive review of this literature is beyond the scope of this paper.

This paper also relates to a literature on stochastic games with an unobservable state. In these games, incentives stem from the large player's ability to manipulate the public belief about the state through her effort choice. Cisternas (2018) characterizes necessary conditions for the existence of Markov equilibria in a continuous time stochastic game with an unobservable state and sufficient conditions in two more restrictive classes of games. Hidden states significantly complicate the model, and it is not possible to establish uniqueness results or a full equilibrium characterization. Board and Meyer-ter-Vehn (2013) study a setting in which a firm's hidden quality depends on past effort and consumers learn about this quality from noisy signals. My paper differs in focus in that there is no adverse selection, there is strategic interaction between the large and small players, and it allows for a richer class of stage game payoffs.

Several folk theorems exist for discrete time stochastic games with observable states, beginning with a perfect monitoring setting in Dutta (1995) and extending to imperfect monitoring environments in Fudenberg and Yamamoto (2011) and Hörner, Takuo, Satoru, and Vieille (2011). My setting differs in that there is a single large player and information follows a diffusion process. It is already known that these two changes significantly alter incentives in standard repeated games (compare the folk theorem in Fudenberg, Levine, and Maskin (1994) to the equilibrium degeneracy in Fudenberg and

Levine (2007, 2009) and Faingold and Sannikov (2011)). The intuition is similar for the discrete time stochastic game folk theorems compared to the MPE uniqueness result in this paper.⁷

The organization of the paper proceeds as follows. Section 2 presents a product choice example to motivate the model. Section 3 sets up the model and characterizes the structure of PPE. Section 4 presents the three main results: existence of a Markov equilibrium, characterization of the PPE payoff set, and uniqueness of a Markov equilibrium in the class of all PPE. Section 5 presents structural results on the shape of equilibrium payoffs. Section 6 explores several applications. All proofs are provided in the Appendix.

2. Example 1: Product choice with persistent quality

Consider a variation of the canonical product choice setting in which a monopolist firm provides a product to consumers and the firm's effort has a persistent effect on the quality of the product. At each time *t*, the firm chooses an unobservable effort level $a_t \in [0, \overline{a}]$, where $\overline{a} > 0$. The quality of the firm's product at time *t* depends on both current and past effort, $q(a_t, X_t) = (1 - \lambda)a_t + \lambda X_t$, where past effort influences quality through the observable stock quality

$$X_t = \int_0^t e^{-\theta(t-s)} (a_s \, ds + dZ_s),$$

 $\theta > 0$ determines the decay rate of past effort, $(Z_t)_{t \ge 0}$ is a standard Brownian motion, and $\lambda \in [0, 1]$ captures the relative importance of past effort in determining current quality.⁸ Effort increases quality both today and in the future.

There is a continuum of identical consumers of unit mass. Consumers value quality: when they believe that the firm will choose effort level \tilde{a}_t at time t, they are willing to pay $q(\tilde{a}_t, X_t)$ for one unit of the product. Each consumer purchases the product for a price equal to her willingness to pay when it is positive and otherwise does not purchase. Therefore, the firm earns a flow revenue of $\bar{b}_t = q(\tilde{a}_t, X_t)$ when $q(\tilde{a}_t, X_t) > 0$ and $\bar{b}_t = 0$ otherwise. This exact form of revenue is chosen for simplicity; the important feature is that the flow revenue is increasing in quality and independent of the true current effort choice. Effort has flow cost $a_t^2/2$ and the firm discounts at rate r > 0. Therefore, the firm's average discounted payoff equals

$$r\int_0^\infty e^{-rt} \left(\overline{b}_t - a_t^2/2\right) dt.$$

In the unique PPE with no persistence, $\lambda = 0$, the firm exerts zero effort, quality is equal to zero, and the firm earns zero profit (this is a direct application of Theorem 3

⁷The paper also relates to an older literature on stochastic games and the existence of Markov equilibria in discrete time, including Shapley (1953), Dutta and Sundaram (1992), Nowak and Raghavan (1992), Duffie, John Geanakoplos, and McLennan (1994).

⁸In a slight abuse of notation, the Lebesgue integral and the stochastic integral are placed under the same integral sign.

from Faingold and Sannikov (2011)). Intertemporal incentives break down, despite the fact that the firm would earn higher profits if it could commit to higher effort.⁹

In this paper, I show that persistent quality incentivizes the firm to choose a positive level of effort and earn positive profits. Theorems 1 to 3 establish that there is a unique PPE, which is Markov in the stock quality X_t . The effort level and profit in this unique equilibrium are characterized as a function of the impact of past effort on current quality λ , the depreciation rate of quality θ , and the discount rate r. For any $\lambda > 0$, the firm chooses a positive level of effort and earns positive profits at positive and some (possibly all) negative levels of stock quality. Further, the firm has a long-run incentive to choose high effort. This contrasts with models in which the incentive to produce high quality is derived from consumers' uncertainty over the firm's payoffs and long-run effort converges to zero (Cripps, Mailath, and Samuelson (2004), Faingold and Sannikov (2011)).

Persistence increases the firm's payoffs through two complementary structural channels. First, the firm's effort increases the stock quality, which increases future revenue through its impact on future prices. This is the direct effect of persistence, as discussed in the Introduction. Second, persistence creates a link with future payoffs, which allows the firm to credibly choose a positive level of effort today, thereby increasing the current price, and hence, revenue. This second channel arises from the strategic interaction between the firm and consumers: it is the equilibrium effect discussed in the Introduction. When quality is high, the continuation value is approximately linear and it is possible to quantify the share of profit arising from each of these channels. The present value of the direct effect minus the cost of effort is approximately $\lambda^2/2(r+\theta)^2$, which is higher when past effort plays a larger role in determining current quality (higher λ), quality depreciates at a lower rate (lower θ), or the firm is more patient (lower r). The present value of the equilibrium effect is approximately $(1 - \lambda)\lambda/(r + \theta)$, which is also higher when quality depreciates at a lower rate or the firm is more patient. In contrast to the direct effect, the equilibrium effect is largest for intermediate values of λ . This is because the incentive to exert effort is increasing in λ while the impact of effort on the current price is increasing in $1 - \lambda$.

This example will be used throughout the paper to demonstrate the results. The product choice framework lends itself to other variations, several of which are discussed in Section 6.1.

3. Model

3.1 Model setup

States and actions A large player and a continuum $I \equiv [0, 1]$ of identical small players, indexed by *i*, play a continuous time stochastic game with imperfect monitoring. At each instant of time $t \in [0, \infty)$, a publicly observable state variable X_t in nonempty closed interval $\mathcal{X} \subset \mathbb{R}$ determines the action set and feasible flow payoffs. If \mathcal{X} is

⁹In contrast to Abreu, Milgrom, and Pearce (1991) and Sannikov and Skrzypacz (2007), this breakdown of incentives takes place despite there being no failure of identifiability.

bounded, denote the upper and lower boundary states by $\overline{X} \equiv \sup \mathcal{X}$ and $\underline{X} \equiv \inf \mathcal{X}$, respectively, and assume $X_0 \in (\underline{X}, \overline{X})$. Large and small players simultaneously choose actions a_t from A and b_t^i from $B(X_t)$, respectively, where A is a nonempty compact subset of a Euclidean space and B(X) is a nonempty compact subset of a closed Euclidean space B with continuous correspondence $X \mapsto B(X)$. Denote the set of feasible pairs of small player actions and states as $E \equiv \{(b, X) \in B \times \mathcal{X} | b \in B(X)\}$. Assume that the boundary of the feasible set of actions for small players grows at most linearly with the state; that is, there exists a K_b , $c_b > 0$ such that for all $(b, X) \in E$, $|b| \le K_b |X| + c_b$.¹⁰ Individual actions are privately observed. Players observe the aggregate distribution of small players' actions, $\overline{b}_t \in \Delta B(X_t)$, and do not observe the large player's action.

Given initial state X_0 , the state evolves stochastically according to

$$dX_t = \mu(a_t, \overline{b}_t, X_t) dt + \sigma(\overline{b}_t, X_t) dZ_t, \tag{1}$$

where $(Z_t)_{t\geq 0}$ is a one-dimensional Brownian motion, and the drift and volatility are determined by Lipschitz continuous functions $\mu : A \times E \to \mathbb{R}$ and $\sigma : E \to \mathbb{R}$, which are linearly extended to $A \times \{(\overline{b}, X) \in \Delta B \times \mathcal{X} | \operatorname{supp} \overline{b} \subset B(X) \}$ and $\{(\overline{b}, X) \in \Delta B \times \mathcal{X} | \operatorname{supp} \overline{b} \subset B(X) \}$, respectively.¹¹ The drift depends on the large player's action, the aggregate action of the small players, and the state. Volatility is independent of the large player's action to maintain the assumption that it is not perfectly observed. If the state space is bounded, then to prevent the state from escaping its boundary and maintain imperfect monitoring at the boundary, the volatility must be zero at the upper and lower bounds, $\sigma(b, \overline{X}) = 0$ for all $b \in B(\overline{X})$ and $\sigma(b, \underline{X}) = 0$ for all $b \in B(\underline{X})$, and the drift must be weakly negative at the upper bound, weakly positive at the lower bound, and independent of (a, b) at both bounds, $\mu(a, b, \overline{X}) = \overline{m} \leq 0$ for all $(a, b) \in A \times B(\overline{X})$ and $\mu(a, b, \underline{X}) = \underline{m} \geq 0$ for all $(a, b) \in A \times B(\underline{X})$. To ensure that the future path of the state is stochastic, except at boundary states, assume that its volatility is positive at all interior states.

ASSUMPTION 1 (Positive Volatility). When $\mathcal{X} = \mathbb{R}$, $\inf_E \sigma(b, X) > 0$. When \mathcal{X} is compact, there exists a C > 0 such that $\sigma(b, X) \ge C(\overline{X} - X)(X - \underline{X})$ for all $(b, X) \in E$.

This assumption rules out interior absorbing states, where state *X* is *absorbing* if the drift and volatility are equal to zero, $\mu(a, \overline{b}, X) = 0$ and $\sigma(\overline{b}, X) = 0$ for all $(a, \overline{b}) \in A \times \Delta B(X)$.

The path of the state provides a public signal of the large player's action. There are no additional public signals. This is without loss of generality, as additional public signals have no effect on the equilibrium characterization (see the discussion in Section 3.3). Let $(F_t)_{t\geq 0}$ represent the filtration generated by the public information $(X_t)_{t\geq 0}$. Small players observe no information about the large player's action beyond what is contained in $(F_t)_{t\geq 0}$.

 $^{^{10}}$ I use $|\cdot|$ to denote the Euclidean norm for vectors..

¹¹Functions μ and σ are extended to distributions as $\int_{B(X)} \mu(a, b, X) d\overline{b}(b)$ and $\int_{B(X)} \sigma(b, X)^2 d\overline{b}(b)$.

Payoffs The payoff of the large player depends on her action, the distribution of small players' actions, and the state. She seeks to maximize the expected value of her discounted payoff

$$r\int_0^\infty e^{-rt}g(a_t,\overline{b}_t,X_t)\,dt$$

where r > 0 is the discount rate and $g : A \times E \to \mathbb{R}$ is a Lipschitz continuous function representing the flow payoff, which is linearly extended to $A \times \{(\overline{b}, X) \in \Delta B \times \mathcal{X} | \operatorname{supp} \overline{b} \subset B(X)\}$. Small players have identical preferences. The payoff of player $i \in I$ depends on her action, the distribution of small players' actions, the large player's action, and the state,

$$r\int_0^\infty e^{-rt}h(a_t,b_t^i,\overline{b}_t,X_t),$$

where $h: A \times \{(b, b', X) \in B^2 \times \mathcal{X} | b, b' \in B(X)\} \to \mathbb{R}$ is a continuous function, which is linearly extended to $A \times \{(b, \overline{b}, X) \in B \times \Delta B \times \mathcal{X} | b \in B(X), \operatorname{supp} \overline{b} \subset B(X)\}$. As is standard, assume that small players do not learn any information about the long-run player's action from observing their flow payoffs beyond that conveyed in the public information. The dependence of payoffs on the state creates a form of action persistence, since the state depends on prior actions.

To ensure that the expected discounted payoff of the large player is well behaved requires a restriction on either the flow payoff of the large player or the growth rate of the state. Assumption 2 states that either the flow payoff is bounded or the drift of the state grows at a linear rate less than the discount rate.

ASSUMPTION 2 (Bounded Payoff or Growth of Drift). At least one of the following conditions holds: (i) the flow payoff g is bounded; (ii) the drift μ has linear growth at a rate less than r: there exists a $K_{\mu} \in [0, r)$ and $c_{\mu} > 0$ such that for all $(a, b, X) \in A \times E$, if $X \ge 0$, then $\mu(a, b, X) \le K_{\mu}X + c_{\mu}$, and if $X \le 0$, then $\mu(a, b, X) \ge K_{\mu}X - c_{\mu}$.

No lower bound is necessary on the slope of the drift when X > 0, since a negatively sloped drift pulls the state toward zero; similarly, no upper bound is necessary when X < 0. This assumption is trivially satisfied when the state space is bounded.

Strategies and equilibrium A public pure strategy for the large player is a stochastic process $(a_t)_{t\geq 0}$ with $a_t \in A$ that is progressively measurable with respect to $(F_t)_{t\geq 0}$. Likewise, a public pure strategy for small player $i \in I$ is a stochastic process $(b_t^i)_{t\geq 0}$ with $b_t^i \in B(X_t)$ that is progressively measurable with respect to $(F_t)_{t\geq 0}$. Given that small players have identical preferences, it is without loss of generality to work with aggregate strategy $(\bar{b}_t)_{t\geq 0}$. The large player's expected discounted payoff at time t under strategy $S = (a_t, \bar{b}_t)_{t\geq 0}$ is given by

$$V_t(S) \equiv E_t \left[r \int_0^\infty e^{-rs} g(a_s, \overline{b}_s, X_s) \, ds \right]. \tag{2}$$

I restrict attention to *pure strategy perfect public equilibria* (PPE). In any PPE, small players' strategies must be myopically optimal because their individual behavior is not observed and does not influence the course of equilibrium play. The following definition modifies the definition in Sannikov (2007) to allow for small players. In a slight abuse of notation, I directly incorporate the myopic incentive constraint for small players into this definition.

DEFINITION 1 (PPE). A public strategy profile $S = (a_t, \overline{b}_t)_{t\geq 0}$ is a *perfect public equilibrium* if, after all public histories, the strategy of the large player maximizes her expected payoff, $V_t(S) \ge V_t(S')$ a.s. for all public strategies $S' = (a'_t, \overline{b}'_t)_{t\geq 0}$ with $(\overline{b}'_t)_{t\geq 0} = (\overline{b}_t)_{t\geq 0}$ almost everywhere, and the strategy of each small player maximizes his expected payoff,

$$b \in \underset{b' \in B(X_t)}{\operatorname{arg\,max}} h(a_t, b', \overline{b}_t, X_t) \quad \forall b \in \operatorname{supp} \overline{b}_t.$$

Timing At each instant *t*, players observe the current state X_t and choose actions. Then nature stochastically determines payoffs and the next state, given the current state and the chosen action profile.

3.2 PPE structure

This section extends the recursive characterization of PPE for continuous time repeated gamed (Sannikov (2007), Faingold and Sannikov (2011)) to continuous time stochastic games.¹² Given strategy profile $S = (a_t, \overline{b}_t)_{t\geq 0}$, define the large player's continuation value as the expected value of the future discounted payoff at time *t*,

$$W_t(S) \equiv E_t \left[r \int_t^\infty e^{-r(s-t)} g(a_s, \overline{b}_s, X_s) \, ds \right]. \tag{3}$$

The expected average discounted payoff at time t can be represented as

$$V_t(S) = r \int_0^t e^{-rs} g(a_s, \overline{b}_s, X_s) \, ds + e^{-rt} W_t(S). \tag{4}$$

Lemma 1 characterizes the evolution of the large player's continuation value and the incentive constraint in PPE. It is the analogue of Theorem 2 in Faingold and Sannikov (2011) without uncertainty over types, allowing for an unbounded state space and flow payoff. Two challenges in extending the PPE characterization to allow for an unbounded flow payoff are showing that $E|V_t(S)| < \infty$ for all $t \ge 0$ and showing that $(W_t(S))_{t\ge 0}$ has linear growth with respect to the state. Establishing these properties requires Assumption 2(ii) to ensure that the state grows at a slow enough rate relative to the discount rate.¹³

¹²Sannikov (2007)'s characterization is for games with two long-run players, while Faingold and Sannikov (2011) focus on public sequential equilibria in games with a single long-run player and incomplete information. Both characterizations build on the recursive methods developed in Abreu, Pearce, and Stacchetti (1990) for discrete time games of imperfect monitoring.

¹³Given Assumption 2(ii), the linear growth of $(W_t)_{t\geq 0}$ is a transversality condition. This result is similar in spirit to Lemma 1 in Strulovici and Szydlowski (2015), which shows that the value function of an optimal control problem is finite and satisfies a linear growth condition with respect to the state.

LEMMA 1 (PPE Characterization). Assume Assumptions 1 and 2. A public strategy profile $S = (a_t, \overline{b}_t)_{t\geq 0}$ is a PPE with continuation values $(W_t)_{t\geq 0}$ if and only if, for some (F_t) -measurable process $(\beta_t)_{t>0}$ in \mathcal{L} ,¹⁴ the following statements hold:

(i) The continuation value $(W_t)_{t\geq 0}$ satisfies

$$dW_t = r(W_t - g(a_t, \overline{b}_t, X_t)) dt + r\beta_t (dX_t - \mu(a_t, \overline{b}_t, X_t) dt)$$
(5)

and there exists a K, $M \ge 0$ such that $|W_t| \le M + K|X_t|$ for all $t \ge 0$, with K = 0 if g is bounded.

(ii) The strategy profile S satisfies sequential rationality, $(a_t, \overline{b}_t)_{t\geq 0} \in S^*(X_t, r\beta_t)$ for almost all $t \geq 0$, where

$$S^{*}(X, z) \equiv \left\{ (a, \overline{b}) : \frac{a \in \arg\max_{a' \in A} g(a', \overline{b}, X) + \frac{z}{r} \mu(a', \overline{b}, X)}{b \in \arg\max_{b' \in B(X)} h(a, b', \overline{b}, X) \forall b \in \operatorname{supp} \overline{b}} \right\}$$
(6)

for $(X, z) \in \mathcal{X} \times \mathbb{R}$.

The first part of the lemma establishes that the continuation value is a stochastic process that is measurable with respect to public information, $(F_t)_{t\geq 0}$. Its drift, W - g(a, b, X), is the difference between the current continuation value and the flow payoff, which captures the expected change in the continuation value. Its volatility $r\beta_t$ determines the sensitivity of the continuation value to information; future payoffs are more sensitive when the volatility is larger. The second part of the lemma shows that the incentive constraint for the large player depends on the trade-off between her action's impact on her flow payoff today and her action's expected impact on her future payoff via the drift of the state, weighted by the incentive weight β_t . It is analogous to the one-shot deviation principle in discrete time. The continuation value and incentive constraint are linear with respect to $(\beta_t)_{t\geq 0}$. This key property of continuous time games with Brownian information is due to the martingale representation theorem and lends significant tractability to the model.

Multiple PPE of the form characterized in Lemma 1 may arise for several reasons. First, at a state X and incentive weight β that are on the equilibrium path, there may be multiple sequentially rational action profiles $(a, \overline{b}) \in S^*(X, r\beta)$. In this case, it is clear that there will be multiple PPE.¹⁵ Second, even if each state and incentive weight prescribe a unique sequentially rational action profile, there may be multiple equilibrium paths of incentive weights $(\beta_t)_{t\geq 0}$ that satisfy Lemma 1 and, hence, multiple PPE. This paper focuses on the latter class of games, in which there is a unique sequentially rational action profile at each X and β , but potentially multiple equilibrium paths of incentive weights. The paper also focuses on a class of games in which the oscillation

¹⁴The notation \mathcal{L} denotes the space of progressively measurable processes $(\beta)_{t\geq 0}$ that are squareintegrable, $E[\int_0^T \beta_t^2 dt] < \infty$ for all $T < \infty$.

¹⁵For example, suppose there are two action profiles in $S^*(X, r\beta)$ for some X and β , denoted by (a_1, \overline{b}_1) and (a_2, \overline{b}_2) . Then when $X_t = X$ and $\beta_t = \beta$, there is a PPE with $(a_t, \overline{b}_t) = (a_1, \overline{b}_1)$ and a PPE with $(a_t, \overline{b}_t) = (a_2, \overline{b}_2)$, where the change in the continuation value dW_t is determined by (5) evaluated at the respective action profile.

of the payoff, drift, and volatility functions is limited. Lipschitz continuity guarantees this when the state space is compact, and, analogously, a monotonicity assumption for large and small states guarantees this when $\mathcal{X} = \mathbb{R}$. The following assumption formalizes these conditions.

ASSUMPTION 3 (Sequentially Rational Action Profile). For all $(X, z) \in \mathcal{X} \times \mathbb{R}$, S^* is nonempty, is single-valued, and returns $\overline{b} = \delta_b$ for some $b \in B(X)$, where δ_b is the Dirac measure on action b. When \mathcal{X} is compact, S^* is Lipschitz continuous on every bounded subset of $\mathcal{X} \times \mathbb{R}$. When $\mathcal{X} = \mathbb{R}$, S^* is Lipschitz continuous on $\mathcal{X} \times \mathbb{R}$ and there exists a $\delta > 0$ such that for all $|X| > \delta$ and $z \in \mathbb{R}$, the rate of change of $g(S^*(X, z), X) + z\mu(S^*(X, z), X)/r$ with respect to X is monotone in X and $\sigma(S^*(X, z), X)$ is monotone in X and constant in z.¹⁶

As we illustrate in Section 6, the assumption is straightforward to verify from the primitives of the model.¹⁷

Importantly, Assumption 3 does not preclude the existence of multiple PPE. As illustrated in Section 6.3, there can be multiple equilibria when it is satisfied.¹⁸ In order to establish that there is a unique PPE, it is also necessary to show that there is a unique path of incentive weights $(\beta_t)_{t\geq 0}$ that satisfies the conditions in Lemma 1. When multiple paths of incentive weights satisfy Lemma 1, each path corresponds to a different PPE with a different strength incentive scheme. For example, there may be a "low" incentive path of $\beta_t = 0$ for all *t* along which players do not invest and a "high" incentive path of $\beta_t > 0$ for all *t* along which players choose a positive level of investment.

One implication of Assumption 3 is that the stage game at any state must have a unique static Nash equilibrium, as the static Nash equilibrium profile corresponds to $S^*(X, 0)$. This rules out coordination games and some games with strategic complementarities. It allows for a broad class of games, including games in which actions are strategic substitutes, games with strategic complementarities that have a unique fixed point, and games with one-sided complementarities between actions. Another implication is that the distribution of the small players' equilibrium actions has a trivial support: all small players play the same action at a given time *t*.

Under Assumption 3, let $(a(X, z), b(X, z)) \equiv S^*(X, z)$ denote the unique sequentially rational action profile at state *X* and incentive weight z/r. Let $g^*(X, z) \equiv$ $g(S^*(X, z), X), \ \mu^*(X, z) \equiv \mu(S^*(X, z), X)$, and $\sigma^*(X, z) \equiv \sigma(b(X, z), X)$ denote the flow payoff, drift, and volatility of the state, respectively, evaluated at this unique sequentially rational action profile. The Lipschitz continuity of S^* implies that the same Lipschitz properties extend to $g^*, \ \mu^*$, and σ^* .

¹⁶In the case of $\mathcal{X} = \mathbb{R}$ and *g* bounded, the weaker assumption $g^*(S^*(X, z), X) + z\mu^*(S^*(X, z), X)/r$ monotone in *X* for large *X* suffices.

 $^{^{17}}$ When S^* is not single-valued, it may not be lower hemicontinuous. Different techniques are necessary to characterize Markov equilibrium payoffs. Similar to Faingold and Sannikov (2011), differential inclusions can be used to characterize the "greatest" and "least" Markov equilibrium payoffs as a function of the state and to show that the PPE payoff set is bounded by these payoffs.

¹⁸Similarly, many discrete time games that satisfy an analogous assumption have multiple nontrivial equilibria. This analogous assumption is more complex since the incentive weights are functions rather than scalars.

EXAMPLE 1 (Product Choice, cont.). To demonstrate Assumptions 1 to 3, return to the example introduced in Section 2. The boundary of the feasible action set for small players is linear in X, $B(X) = [0, \overline{a} + X]$. The volatility of the state is constant, $\sigma(\overline{b}, X) = 1$ (Assumption 1). The drift is $\mu(a, \overline{b}, X) = a - \theta X$, which is negative when X is high and positive when X is low (Assumption 2(ii)). The large player's payoff is $g(a, \overline{b}, X) = \overline{b} - a^2/2$. Given (X, z), sequential rationality for the firm requires

$$a \in \arg \max_{a \in [0,\overline{a}]} \overline{b} - \frac{1}{2}a^2 + \frac{z}{r}(a - \theta X),$$

which yields a(X, z) = z/r for $z \in [0, r\overline{a}]$, a(X, z) = 0 for z < 0, and $a(X, z) = \overline{a}$ for $z > r\overline{a}$. In equilibrium, consumers' beliefs about the effort choice of the firm are correct. Therefore, consumers are willing to pay $b(X, z) = \max\{0, q(a(X, z), X)\}$. This sequentially rational action profile is unique and Lipschitz continuous in (X, z) (Assumption 3). Given a(X, z) and b(X, z), the flow payoff is $g^*(X, z) = b(X, z) - a(X, z)^2/2$, the drift of the state is $\mu^*(X, z) = a(X, z) - \theta X$, and the volatility of the state is $\sigma^*(X, z) = 1$, which is trivially monotone in X and constant in z for large |X|. Further, $\frac{d}{dX}(g^*(X, z) + z\mu^*(X, z)/r)$ is $\lambda - z\theta/r$ for large X and $-z\theta/r$ for small X, which is constant and, hence, monotone. \Diamond

3.3 Discussion of model

Equilibrium actions In many applications—including rational expectations equilibrium models and learning models, where the state is a belief—the transition of the state and/or the large player's flow payoff depend on the large player's chosen action and her equilibrium action (i.e., the action other players expect). The framework in this paper indirectly allows for such dependences, since the best response of the small player depends on the expected action of the large player, which is correct in equilibrium. One could also model such dependences directly by adding the equilibrium action, denoted \tilde{a} , as an argument to the drift and the volatility, i.e., $\mu'(a, \tilde{a}, b, X)$ and $\sigma'(\tilde{a}, b, X)$.¹⁹ The analysis is unchanged, provided μ' and σ' satisfy Assumptions 1 to 3.²⁰ For example, suppose the drift is $\mu(a, b, X) = \theta_1 b + \theta_2 a$ and the small player's payoff is $ab - b^2/2$. When the small player believes that the large player will choose \tilde{a} , her best response is $b = \tilde{a}$. This is isomorphic to a model in which the drift directly depends on the equilibrium action, $\mu'(a, \tilde{a}, b, X) = \theta_1 \tilde{a} + \theta_2 a$.

The framework presented here does rule out some classes of stochastic games. In particular, consider Bayesian learning games with a binary outcome space, and let the state $X \in [0, 1]$ denote the belief that the outcome is high. Assumption 1 rules out games in which there exists an action profile that shuts down learning at interior beliefs, i.e., there exists an $(\tilde{a}, b) \in A \times B$ such that $\sigma'(\tilde{a}, b, X) = 0$ at some $X \in (0, 1)$. This contrasts with Faingold and Sannikov (2011), in which it is feasible for the normal player to shut

¹⁹Imperfect monitoring is maintained when the volatility depends on the equilibrium action, in contrast to the chosen action. This is because the expected action does not reveal the chosen action.

²⁰This highlights the distinction from a single-agent decision problem, as the present framework is a fixed-point problem.

down learning by perfectly mimicking the behavioral type. Therefore, their setup does not satisfy Assumption 1 and their model requires an alternative approach to establish that the volatility of the state is bounded away from zero in equilibrium.

Incentives in continuous time stochastic games As discussed in the Introduction, in a stochastic game, incentives can either be *informational*—past signals are used to coordinate future equilibrium play—or *structural*—past actions impact the structure of future interactions through their impact on the state. This latter channel includes both the state's direct impact on future feasible payoffs and its indirect impact through its effect on future equilibrium play. The process $(\beta_t)_{t\geq 0}$ characterized in Lemma 1 captures all of these channels for intertemporal incentives.

The linear structure of the continuation value with respect to $(\beta_t)_{t\geq 0}$ (as characterized in (5)) plays a key role in determining incentives. In a repeated game with small players, this linearity precludes effective intertemporal incentives (Faingold and Sannikov (2011)). This is because when the continuation value is at its maximum, a linear transfer with a nontrivial incentive weight $\beta_t > 0$ results in the continuation value exceeding its maximum unless the transfer is tangential to the boundary of the equilibrium payoff set.²¹ But tangential transfers are not possible, since small players are myopic. Therefore, it must be that $\beta_t = 0$ and the large player acts myopically, yielding a static Nash payoff. However, in a stochastic game, β_t can depend on the state. Therefore, it may be possible to have a linear transfer with nontrivial incentive weight $\beta_t > 0$ at states that do not yield the maximum continuation value, while setting $\beta_t = 0$ at states that do to ensure that the continuation value does not exceed its boundary. The remainder of the paper explores whether and when it is possible to create effective incentives in this manner.²²

Additional public signals The path of the state both serves as a public signal of the large player's action and directly impacts payoffs. The results and analysis are unchanged if there are additional payoff-irrelevant Brownian public signals. Markov equilibria ignore such signals, so the characterization of Markov equilibria remains the same. Further, it is not possible to effectively use such signals to coordinate additional equilibria, due to reasoning similar to Faingold and Sannikov (2011). An older working paper version of the current paper allows for an arbitrary finite number of public signals (Bohren (2016)).

4. Equilibrium analysis

This section presents the main results of the paper. I establish the existence of Markov equilibria, characterize the correspondence of PPE payoffs of the large player, and derive conditions under which there is a unique PPE, which is Markov.

²¹This intuition is formalized in Sannikov and Skrzypacz (2010), who show that the only effective ways to use Brownian information are linearly (i.e., the continuation value depends linearly on $(\beta_t)_{t\geq 0}$) and to transfer value tangentially along the boundary of the equilibrium payoff set.

²²Nonlinear incentive structures, such as value-burning, are ineffective in both repeated and stochastic games with Brownian information, because the expected loss from false punishment exceeds the expected gain from cooperating (Fudenberg and Levine (2007, 2009), Sannikov and Skrzypacz (2007, 2010)).

4.1 Existence of Markov equilibria

In a Markov equilibrium, the continuation value and actions depend solely on the current value of the state; they are independent of the past path of the state. Since the path of the state provides a signal of the large player's action, using it to punish or reward the large player could give rise to PPE in which different paths of the state specify different continuation payoffs and equilibrium actions, even when these paths map to the same current state. In a Markov equilibrium, this is not allowed.

Theorem 1 establishes existence of a Markov equilibrium and characterizes equilibrium behavior and payoffs in Markov equilibria. The continuation value is characterized as the solution(s) $U : \mathcal{X} \to \mathbb{R}$ to an ordinary differential equation that maps each state to a payoff. If there are multiple solutions, then each solution characterizes a Markov equilibrium (see Section 6.3 for an illustration of a setting with multiple Markov equilibria). Given a solution U, the corresponding Markov equilibrium action profile is the sequentially rational action profile at state X and incentive weight U'(X)/r. The large player has nondegenerate incentives at any state with $U'(X) \neq 0$.

THEOREM 1. Assume Assumptions 1 to 3. Given initial state X_0 , if U is a solution to the optimality equation

$$rU(X) = rg^*(X, U'(X)) + U'(X)\mu^*(X, U'(X)) + \frac{1}{2}U''(X)\sigma^*(X, U'(X))^2$$
(7)

on \mathcal{X} (on $(\underline{X}, \overline{X})$ if \mathcal{X} is compact) and U has linear growth (is bounded if g is bounded), then U characterizes a Markov equilibrium with the following payoffs and actions:

- (i) Equilibrium payoff $U(X_0)$.
- (ii) Continuation values $(W_t)_{t\geq 0} = (U(X_t))_{t\geq 0}$.
- (iii) Equilibrium actions $(a_t, \overline{b}_t)_{t\geq 0} = (S^*(X_t, U'(X_t)))_{t\geq 0}$, where $S^*(X, U'(X))$ is the unique solution to (6) at state X and incentive weight U'(X)/r.

The optimality equation has at least one twice continuously differentiable solution that lies in the range of feasible payoffs for the large player and has linear growth (is bounded if g is bounded). Thus, there exists at least one Markov equilibrium.

From the optimality equation, the continuation value U(X) is equal to the sum of the equilibrium flow payoff $g^*(X, U'(X))$ and the expected change in the continuation value. This expected change has two components: (i) the interaction between the slope of the continuation value and the drift of the state, $U'(X)\mu^*(X, U'(X))/r$, and (ii) the interaction between the concavity of the continuation value and the volatility of the state, $U''(X)\sigma^*(X, U'(X))^2/2r$.

In relation to the discussion in Section 3.3, a nontrivial incentive weight is possible at some states without the continuation value escaping the payoff set. Theorem 1 shows that the volatility of the continuation value in a Markov equilibrium is equal to its slope, $r\beta_t = U'(X_t)$. At any interior state X^* that yields the maximum continuation value

across all states, $U'(X^*) = 0$. Therefore, when $X_t = X^*$, the volatility of the continuation value is zero, $r\beta_t = 0$, which ensures that the continuation value does not escape the payoff set. In these periods, the large player acts myopically and earns the static Nash payoff in state X^* . At other states, the continuation value can be sensitive to changes in the state, $U'(X) \neq 0$, generating nontrivial incentives.

Outline of proof In a Markov equilibrium, continuation values take the form of $W_t = U(X_t)$ for some function U. Assuming that U is twice continuously differentiable, by Ito's formula the continuation value must follow the law of motion,

$$dU(X_t) = U'(X_t)\mu(a_t^*, \overline{b}_t^*, X_t) dt + \frac{1}{2}U''(X_t)\sigma(\overline{b}_t^*, X_t)^2 dt + U'(X_t)\sigma(\overline{b}_t^*, X_t) dZ_t$$

By Lemma 1, the continuation value must also follow the law of motion in (5). Matching the drifts of these two laws of motion yields the optimality equation, while matching the volatilities yields the equilibrium volatility of the continuation value, $r\beta_t = U'(X_t)$. Showing that the optimality equation has at least one twice continuously differentiable solution that lies in the range of feasible payoffs for the large player establishes existence.

Faingold and Sannikov (2011) follow similar steps to derive a Markov equilibrium in a game of incomplete information. Relative to their derivation, the innovative part of my proof lies in establishing existence of a solution to the optimality equation when the state space is unbounded, particularly when *g* is also unbounded. I show by construction that there exist lower and upper solutions to the optimality equation, $\underline{\alpha} : \mathcal{X} \to \mathbb{R}$ and $\overline{\alpha} : \mathcal{X} \to \mathbb{R}$, that have linear growth. This is only possible when the maximum drift of the state has linear growth at rate less than *r* (Assumption 2). The lower and upper solutions characterize bounds on the solution to the optimality equation, $\underline{\alpha}(X) \leq U(X) \leq \overline{\alpha}(X)$ for all *X*. Next I show that the bound on the optimality equation grows linearly with respect to U'(X) and, therefore, the optimality equation does not grow too quickly (technically speaking, it satisfies a growth condition on any compact subset of the state space). These conditions establish that the optimality equation has a twice continuously differentiable solution with linear growth. When *g* is bounded, the lower and upper solutions are constant, which establishes existence of a bounded solution.

The final step is to show that the continuation value and actions characterized above constitute a Markov equilibrium. Given a solution U(X) and an action profile uniquely specified at state X_t by $(a_t^*, \overline{b}_t^*) = S^*(X_t, U'(X_t))$ (where uniqueness follows from Assumption 3), the state variable evolves uniquely according to (1), the continuation value $(U(X_t))_{t\geq 0}$ satisfies the law of motion (5), and the action profile satisfies the conditions for sequential rationality (6). Therefore, $(a_t^*, \overline{b}_t^*, U(X_t))$ constitute a PPE.

For a given solution U(X), the state evolves uniquely and actions are uniquely specified as a function of the state. Therefore, *each* solution to the optimality equation characterizes a unique Markov equilibrium. If there are multiple solutions, then there will be multiple Markov equilibria.

EXAMPLE 1 (Product Choice, cont.). Given a(X, z) and b(X, z) characterized in Section 3.2, any solution to

$$rU(X) = rb(X, U'(X)) - \frac{r}{2}a(X, U'(X))^{2} + U'(X)(a(X, U'(X)) - \theta X) + \frac{1}{2}U''(X)$$
(8)

with linear growth as $X \to \infty$ and bounded as $X \to -\infty$ characterizes a Markov equilibrium with equilibrium actions a(X, U'(X)) and b(X, U'(X)).

4.2 The PPE payoff set

Let $\xi : \mathcal{X} \rightrightarrows \mathbb{R}$ denote the correspondence that maps each state onto the corresponding set of PPE payoffs for the large player, and let $Y : \mathcal{X} \rightrightarrows \mathbb{R}$ denote the analogous correspondence for the Markov equilibrium payoffs characterized by the optimality equation in Theorem 1. Theorem 2 shows that in any PPE, the large player cannot achieve a payoff above the highest or below the lowest Markov equilibrium payoff in Y.

THEOREM 2. Assume Assumptions 1 to 3. Then for any state $X \in \mathcal{X}$ (state $X \in (\underline{X}, \overline{X})$ if \mathcal{X} is compact), the set of PPE payoffs of the large player at state X is equal to the convex hull of the set of Markov equilibrium payoffs at state X, $\xi(X) = co(Y(X))$.

The impossibility of the large player achieving a PPE payoff above the highest Markov payoff in Y yields insight into the type of incentives generated by persistence. As discussed in the Introduction and Section 3.3, incentives can be either informational or structural. When a Markov equilibrium yields the highest equilibrium payoff, it precludes the existence of equilibria that achieve higher payoffs using informational incentives. Therefore, any nontrivial incentives arising from persistence are structural.

Outline of proof The key argument in the proof shows that any PPE with an initial payoff above the highest Markov equilibrium payoff in Y will eventually yield a continuation value that lies outside the set of feasible payoffs for the large player, which is a contradiction. Suppose that a PPE with continuation values $(W_t)_{t\geq 0}$ yields a payoff higher than the maximum Markov equilibrium payoff in Y at state X_0 . Let $D_t \equiv W_t - U(X_t)$ be the difference between the continuation values in these two equilibria at time t. I show that whenever $D_0 > 0$, D_t will grow arbitrarily large with positive probability, independent of X_t . By Lemma 1, $|W_t(S)|$ is bounded with respect to X_t . Thus, D_t can only grow arbitrarily large when X_t grows arbitrarily large, so it cannot be that $D_0 > 0$.

This escape argument is similar to other papers in the literature, in particular Faingold and Sannikov (2011). Their proof relies on the compactness of the state space to show that the volatility of D_t is bounded away from zero and relies on the boundedness of the flow payoff to reach a contradiction when D_t grows arbitrarily large. Therefore, their proofs do not trivially extend to an unbounded state space or an unbounded flow payoff. The innovative parts of this proof are to establish that the volatility of D_t is bounded state space and to show that when D_t grows arbitrarily large, it can jump outside of the feasible payoff set (a contradiction) provided the state does not grow too quickly.

Equilibrium degeneracy without persistent actions If the state evolves independently of the large player's action, then there is no link between the current action and the continuation value. It is not possible to generate effective intertemporal incentives and the large player acts myopically. In the unique PPE, both players play the static Nash equilibrium action profile $S^*(X, 0)$ at all states X.

COROLLARY 1. Assume Assumptions 1 to 3 and suppose μ is independent of a for all X. Then in the unique PPE, $(a_t, \overline{b}_t) = S^*(X_t, 0)$ for all $t \ge 0$ and the continuation value is characterized by the unique solution to the optimality equation (7).

This is the stochastic game analogue of the equilibrium degeneracy result in repeated games with a long-run player and short-run/small players (Fudenberg and Levine (2007, 2009), Faingold and Sannikov (2011)).

4.3 Equilibrium uniqueness

This section establishes sufficient conditions for there to be a unique PPE, which is Markov. The main step is to determine when the optimality equation has a unique feasible solution. When this is the case, Theorem 2 establishes that PPE payoffs are uniquely specified as the payoffs in this unique Markov equilibrium. The behavior of the optimality equation as the state approaches its boundary plays a key role in establishing when it has a unique solution. Any two feasible solutions that satisfy the same boundary conditions cannot differ on the interior of the state space: they must be equivalent (see Lemma 7 in Appendix A.4). Therefore, establishing that all feasible solutions satisfy the same boundary conditions is necessary and sufficient to establish a unique solution. I outline a set of sufficient conditions to guarantee this when $\mathcal{X} = \mathbb{R}$; the case of a compact state space requires no additional conditions. The application in Section 6.3 illustrates how multiple Markov equilibria can arise when this condition fails.

4.3.1 Unbounded state space $(\mathcal{X} = \mathbb{R})$ Assumption 4 (below) outlines a set of sufficient conditions for a unique Markov equilibrium when $\mathcal{X} = \mathbb{R}$. The first condition requires the large player's equilibrium flow payoff and the equilibrium drift to be additively separable in the state X and incentive weight z as X approaches ∞ and $-\infty$. This rules out complementarities between the direct and equilibrium channels for incentives near the boundary, which prevents multiple equilibrium incentive weights—and hence, equilibrium action profiles—at a given state. It is used to establish that the slope of the continuation value converges to the same limit in all Markov equilibria. The second condition relates to the volatility: it is a technical condition that helps establish that two distinct solutions to the optimality equation cannot have the same limit slope. The third condition applies to a growth model where the drift of the state approaches infinity as $X \to \infty$ (or approaches negative infinity as $X \to -\infty$); it ensures that the volatility does not also grow arbitrarily large. It is also used to pin down a unique boundary continuation value.

Assumption 4.

- (i) Additive Separability Near Boundary. There exists $a \delta > 0$ and continuously differentiable functions $g_1, \mu_1 : \mathcal{X} \to \mathbb{R}$ and $g_2, \mu_2 : \mathbb{R} \to \mathbb{R}$ with μ_1 monotone such that for $|X| > \delta$, $g^*(X, z) = g_1(X) + g_2(z)$ and $\mu^*(X, z) = \mu_1(X) + \mu_2(z)$.
- (ii) Volatility. The function $\sigma^*(X, z)^2$ is Lipschitz continuous.

(iii) Growth Case. When $\lim_{X\to\infty} \mu_1(X) = \infty$, then there exists an ε , $\delta > 0$ such that for $X > \delta$ and $z \in \mathbb{R}$, $|\mu_1(X)|/\sigma^*(X, z)^2 > \varepsilon$, and similarly when $\lim_{X\to-\infty} \mu_1(X) = -\infty$.²³

Given Assumption 4(i), select $g_1(X)$ and $g_2(z)$ such that $g_2(z)$ contains any constant term in $g^*(X, z)$ to uniquely pin down each function, and similarly for $\mu_1(X)$ and $\mu_2(z)$. When g is bounded, it is possible to establish uniqueness without additive separability; Assumption 5 in Supplemental Appendix D.3 (available at http://econtheory.org/supp/ 2680/supplement.pdf) presents an alternative condition.

Theorem 3 establishes uniqueness and characterizes the limit of the continuation value and its slope as the state grows large.

THEOREM 3. Suppose $\mathcal{X} = \mathbb{R}$ and assume Assumptions 1 to 4. For each initial state $X_0 \in \mathcal{X}$, there exists a unique PPE that is Markov and characterized by the unique solution U of (7) on \mathcal{X} with linear growth (bounded when g is bounded). The slope of the continuation value converges to a constant,

$$\lim_{X \to x} U'(X) = z_x \quad \text{where } z_x \equiv \lim_{X \to x} rg_1(X) / (rX - \mu_1(X)), \tag{9}$$

and the continuation value converges to

$$\lim_{X \to x} U(X) - y(X) = g_2(z_x) + z_x \mu_2(z_x)/r$$
(10)

for $x \in \{-\infty, \infty\}$, where $y(X) \equiv -\phi(X) \int (rg_1(X)/\phi(X)\mu_1(X)) dX$ and $\phi(X) \equiv \exp(\int (r/\mu_1(X)) dX)$ when $\lim_{X \to x} \mu_1(x) \neq 0$, and $y(X) \equiv g_1(X)$ when $\lim_{X \to x} \mu_1(X) = 0$. When g is bounded, this implies the continuation value converges to the limit static Nash equilibrium payoff and the slope of the continuation value converges to zero: for $x \in \{-\infty, \infty\}$,

$$\lim_{X \to x} (U(X) - g^*(X, 0)) = 0 \quad and \quad \lim_{X \to x} U'(X) = 0.$$
(11)

Theorem 3 establishes that the slope of the continuation value converges to a unique limit slope, which is equal to the ratio of the growth rate of the flow payoff to the growth rate of the drift with respect to the state. Given this slope, the boundary condition (10) highlights the impact of structural incentives on the continuation payoff. Repeated play of the static Nash equilibrium profile yields a payoff U^{NE} that satisfies $\lim_{X\to x} U^{\text{NE}}(X) - y(X) = g_2(0) + z_x \mu_2(0)/r$. Therefore, from (10), the continuation value approaches the sum of this repeated static Nash payoff and a constant $g_2(z_x) - g_2(0) + z_x(\mu_2(z_x) - \mu_2(0))/r$. This constant determines the extent to which structural incentives persist at the boundary of the state space. The first term, $g_2(z_x) - g_2(0)$, captures the equilibrium effect of persistence. It is the portion of the equilibrium flow payoff that arises from future strategic interaction; it captures the effect of the large player's action on the small players' actions, net of the cost of *a*. The

²³Part (ii) is unnecessary when g is bounded. Note that part (iii) holds trivially when $\sigma(b, X)$ is bounded.

second term, $z_x(\mu_2(z_x) - \mu_2(0))/r$, captures the direct effect of persistence on future feasible payoffs, measured by how the continuation value changes with respect to the state and how the state changes with respect to the large player's equilibrium action relative to the static Nash action. If this constant is positive, then as the state becomes large, structural incentives provide the large player with a payoff that is strictly higher than the payoff from playing the static Nash profile at each state.

When the asymptotic slope z_x is nonzero, it is possible to sustain nontrivial intertemporal incentives as the state grows large. This is an important and novel insight of this paper. If it is possible to sustain nontrivial incentives at the boundary of the state space, then incentives are permanent in the sense that they do not dissipate with time, regardless of the asymptotic behavior of the state with respect to time. In the case of a bounded flow payoff, $z_x = 0$. Therefore, incentives collapse at the boundary and the continuation value converges to the limit of the static Nash payoff. However, this does not preclude the existence of long-run incentives: even when incentives collapse at the boundary, the state does not necessarily converge to a boundary state as $t \to \infty$. Therefore, it can be possible to sustain nontrivial incentives in the long run.

The continuation of Example 1 below illustrates how to verify Assumption 4 and derive the boundary conditions in Theorem 3 when the flow payoff is unbounded, while Section 6.1 illustrates how to do so for a bounded flow payoff. Section 6.3 shows that there can be multiple MPE in an application in which Assumption 4 (specifically, additive separability) fails.

Outline of proof I first show that all solutions to the optimality equation have the same boundary conditions. Faingold and Sannikov (2011) also characterize boundary conditions as a step towards establishing that the optimality equation in their paper has a unique solution. Relative to their result, the innovative part of my proof is in establishing boundary conditions for an unbounded flow payoff and state space, as I next describe. Let $\psi(X, z) \equiv g^*(X, z) + z\mu^*(X, z)/r$ be the sum of the large player's flow payoff and return on effort at the sequentially rational action profile (a(X, z), b(X, z)), and let U(X)be a solution to the optimality equation. Suppose that U'(X) does not converge as $X \rightarrow$ ∞ . Then for any slope *z* such that the continuation value has slope *z* infinitely often at large X, U(X) will alternate between being convex and concave at slope z. From the optimality equation, $\psi(X, z)$ will lie above U(X) when it is concave at slope z and will lie below U(X) when it is convex at slope z. Therefore, the oscillation of $\psi'(X, U'(X))$ is at least as large as the oscillation of U'(X). This violates the monotonicity of ψ' , so it must be that U'(X) has a limit $z_{\infty} \in \mathbb{R}$. Since U(X) has linear growth (by Theorem 1), this limit must be finite. Moreover, it is equal to $\lim_{X\to\infty} U(X)/X$. Given additive separability, as well as the Lipschitz continuity and monotonicity of μ_1 and g_1 , the limits of $\psi(X, z)/X$ and $\psi'(X, z)$ exist and are equal as $X \to \infty$. Denote these limits by $\psi_{\infty}(z)$. We use these properties and the optimality equation to show that $\lim_{X\to\infty} \sigma^*(X, U'(X))^2 U''(X)/X =$ 0 and, therefore, $\lim_{X\to\infty} U(X)/X - \psi(X, U'(X))/X = 0$. This establishes that the limit slope z_{∞} is a fixed point of $\psi_{\infty}(z)$. The additively separable assumption on g^* and μ^* is sufficient to ensure that $\psi_{\infty}(z)$ has a unique fixed point, which is equal to $z_{\infty} = \lim_{X \to \infty} rg_1(X)/(rX - \mu_1(X))$. This guarantees that all solutions to the optimality equation have the same limit slope.

Using the characterization of the limit slope, it can be shown that any solution U(X) to the optimality equation satisfies $\lim_{X\to\infty} U(X) - U'(X)\mu_1(X)/r - g_1(X) = g_2(z_{\infty}) + z_{\infty}\mu_2(z_{\infty})/r$. Consider the linear first-order differential equation (FODE) $y(X) - y'(X)\mu_1(X)/r - g_1(X) = 0$. Establishing that any solution U(X) satisfies $\lim_{X\to\infty} U(X) - y(X) = g_2(z_{\infty}) + z_{\infty}\mu_2(z_{\infty})/r$ for any linear growth solution y to this FODE yields the boundary condition for U(X) (i.e., (10)). Therefore, all solutions to the optimality equation approach the same value and slope as the state grows large or small.

Finally, I show that any two such solutions U and V cannot differ on the interior of the state space. Similar to Faingold and Sannikov (2011), if there exists an X such that U(X) - V(X) > 0, the structure of the optimality equation prevents these solutions from satisfying the same boundary conditions for at least one boundary.

EXAMPLE 1 (Product Choice, cont.). This example satisfies Assumption 4. From the characterization in Section 3.2, the sequentially rational effort a(X, z) is independent of X and the consumers' willingness to pay b(X, z) is additively separable in (X, z). Therefore, the flow payoff $g^*(X, z) = b(X, z) - a(X, z)^2/2$ and the drift $\mu^*(X, z) = a(X, z) - \theta X$ are additively separable in (X, z). From these expressions, $g_1(X) = \lambda X$ for X > 0 and $g_1(X) = 0$ for X < 0, while $\mu_1(X) = -\theta X$ for all X, which is monotone. Finally, $\sigma^*(X, z)^2 = 1$ trivially satisfies Lipschitz continuity, and the growth condition is not relevant since $\lim_{X \to \infty} \mu_1(X) = -\infty$ and similarly for $X \to -\infty$.

From Theorem 3, the limit slopes are $z_{\infty} = r\lambda/(r + \theta)$ and $z_{-\infty} = 0$. Therefore, equilibrium effort approaches $a(X, z_{\infty}) = \lambda/(r + \theta)$ as *X* grows large, which is strictly positive. As discussed in Section 2, this contrasts with settings in which effort does not have a persistent effect on quality and long-run effort converges to zero (Cripps, Mailath, and Samuelson (2004), Faingold and Sannikov (2011)). From (10), for large *X* the continuation value approximates

$$U(X) \approx \frac{r\lambda}{r+\theta} X + \frac{(1-\lambda)\lambda}{r+\theta} + \frac{\lambda^2}{2(r+\theta)^2},$$

where the first term is the payoff from repeated play of the static Nash equilibrium profile, and the second and third terms capture the impact of structural incentives on the equilibrium payoff: the equilibrium effect of persistence stemming from future strategic interaction between the firm and consumers, and the direct effect of persistence on future payoffs via the stock quality, respectively.²⁴ In contrast, as *X* approaches $-\infty$, equilibrium effort approaches zero and the continuation value converges to zero, $\lim_{X\to -\infty} U(X) = 0$. Therefore, at large negative values of the state, incentives collapse.

²⁴Given the expression for a(X, z) above, $g_2(z) = (1 - \lambda)a(X, z) - a(X, z)^2/2$ and $\mu_2(z) = a(X, z)$, the constant on the right hand side of (10) is $(1 - \lambda)\lambda/(r + \theta) + \lambda^2/2(r + \theta)^2$. The payoff from repeated play of the static Nash equilibrium profile, $y(X) = r\lambda X/(r + \theta)$, is calculated from the expression for y(X) in Theorem 3, using the expressions for $g_1(X)$ and $\mu_1(X)$ above and $\phi(X) = \exp(-\int (r/\theta X) dX) = X^{-r/\theta}$.

4.3.2 Bounded state space (\mathcal{X} compact) When \mathcal{X} is compact, uniqueness follows from Assumptions 1 to 3. No additional conditions are needed as in Theorem 3, as Lipschitz continuity together with the conditions on the drift and volatility that prevent the state from escaping its boundary (i.e., positive drift and zero volatility at \underline{X} , and analogously for \overline{X}) establish that the large player plays a unique action at the boundary and pin down a unique boundary continuation value. Theorem 4 establishes uniqueness when the state space is compact, and characterizes the limit of the continuation value and the large player's incentive constraint.²⁵

THEOREM 4. Suppose \mathcal{X} is compact and assume Assumptions 1 to 3. For each initial state $X_0 \in \mathcal{X}$, there exists a unique PPE that is Markov and characterized by the unique bounded solution U of (7) on $(\underline{X}, \overline{X})$. When the boundary states are absorbing, the continuation value converges to the static Nash equilibrium payoff and intertemporal incentives collapse at the boundary,

$$\lim_{X \to x} (U(X) - g^*(X, 0)) = 0 \quad and \quad \lim_{X \to x} \mu^* (X, U'(X)) U'(X) = 0$$
(12)

for $x \in \{\underline{X}, \overline{X}\}$. When the boundary states are not absorbing,

$$\lim_{X \to \underline{X}} U(X) = g^*(\underline{X}, 0) + \underline{mu'}/r \quad and \quad \lim_{X \to \overline{X}} U(X) = g^*(\overline{X}, 0) + \overline{mu'}/r$$
(13)

given unique finite limit slopes $\underline{u}' \equiv \lim_{X \to \underline{X}} U'(X)$ and $\overline{u}' \equiv \lim_{X \to \overline{X}} U'(X)$.

The continuation value at a boundary state depends on whether the boundary state is absorbing or reflecting. When the boundary is absorbing, the state remains at the boundary once it is reached and, therefore, the continuation value converges to the static Nash payoff. When the boundary is reflecting, the limit of the continuation value also depends on its (unique) limit slope and the boundary drift, which captures how quickly the state moves away from the boundary and how the continuation value changes as the state changes. In either case, the impact of the long-run player's action on the drift of the state converges to zero at the boundary. Therefore, incentives collapse and the equilibrium action profile converges to the static Nash action profile. This rules out the possibility of sustaining multiple equilibrium action profiles at the boundary, a key step in establishing uniqueness. An important difference from Theorem 3 is that incentives collapse even if the slope of the continuation value does not converge to zero. This stems from the requirement that the boundary drift is independent of the

²⁵Theorems 1, 2, and 4 also hold for an alternative version of Assumption 3 when the static Nash payoff $g^*(X, 0)$ is increasing in X. Specifically, assume that the restriction of S^* to $\mathcal{X} \times [0, \infty)$ is nonempty, is single-valued, and returns $\overline{b} = \delta_b$ for some $b \in B(X)$, where δ_b is the Dirac measure on action b, S^* is Lipschitz continuous on every bounded subset of $\mathcal{X} \times [0, \infty)$ when \mathcal{X} is compact and on $\mathcal{X} \times [0, \infty)$ when $\mathcal{X} = \mathbb{R}$, and when $\mathcal{X} = \mathbb{R}$, there exists a $\delta > 0$ such that for all $|X| > \delta$ and $z \in [0, \infty)$, the rate of change of $g(S^*(X, z), X) + z\mu(S^*(X, z), X)/r$ with respect to X is monotone in X and $\sigma(S^*(X, z), X)$ is monotone in X and constant in z. Change the definition of S^* to set $S^*(X, z) = S^*(X, 0)$ for z < 0. By Proposition 2, the solution U(X) is increasing, and the values for z < 0 are irrelevant. An analogous restriction to $(-\infty, 0]$ is possible when $g^*(X, 0)$ is decreasing in X.

long-run player's action—in order to maintain imperfect monitoring when the volatility is zero—combined with continuity as the drift approaches its boundary. As in Theorem 3, when incentives collapse at the boundary, this does not preclude the existence of long-run incentives, as the state does not necessarily converge to a boundary state as $t \rightarrow \infty$. Section 6.2 provides an illustration of Theorem 4.

5. PROPERTIES OF EQUILIBRIUM PAYOFFS

The optimality equation yields rich insights into how the correspondence of PPE payoffs is tied to the underlying structure of the game. Propositions 1 and 2 show that the shape of the static Nash equilibrium payoff $g^*(X, 0)$ is a key determinant of the shape of the Markov equilibrium continuation value. Note that $g^*(X, 0)$ is straightforward to derive from the primitives of the game.

Proposition 1 relates the number and type of extrema for U(X) to the shape of $g^*(X, 0)$. Given a solution U(X) to the optimality equation, define an *interval minimum* of U on a closed proper interval $I \subset \mathcal{X}$ as $[X_a, X_b] \subset \text{int } I$ such that U'(X) = 0for all $X \in [X_a, X_b]$, and there exists an $\varepsilon > 0$ such that $U(X_a) < U(X)$ for all $X \in (X_a - \varepsilon, X_a) \cup (X_b, X_b + \varepsilon)$, with an analogous definition for *interval maximum*.²⁶

PROPOSITION 1. Assume Assumptions 1 to 3. Let $I \subset X$ denote a closed proper interval of states and let U(X) denote a linear growth or bounded (when g bounded) solution to (7).

- (i) If $g^*(X, 0)$ is constant on I, then U(X) has at most one interval extremum on I.
- (ii) If $g^*(X, 0)$ is strictly monotone on I, then U(X) has at most two interval extrema on I and is not constant on I. If $g^*(X, 0)$ is strictly increasing (decreasing) on I, and U(X) has an interval minimum $[X_{1a}, X_{1b}]$ and maximum $[X_{2a}, X_{2b}]$, then $X_{1b} < X_{2a}$ $(X_{2b} < X_{1a})$.
- (iii) If $g^*(X, 0)$ has n interval extrema on I, then U(X) has at most n + 2 interval extrema on I.

The intuition for Proposition 1 stems from the behavior of the continuation value at interior extrema. Given solution U(X), if there is an extremum at state X, then U'(X) = 0 and the optimality equation simplifies to $U(X) = g^*(X, 0) + U''(X)\sigma^*(X, 0)^2/2r$. If the extremum is a minimum, $U''(X) \ge 0$, and, therefore, $U(X) \ge g^*(X, 0)$. Similarly, at a maximum, $U''(X) \le 0$, and, therefore, $U(X) \le g^*(X, 0)$. Hence, the oscillation of U(X) is bounded by the oscillation of $g^*(X, 0)$.

When the continuation value converges to the static Nash payoff at boundary states, then it is possible to characterize additional results on the shape of payoffs across the entire state space. Proposition 2 relates the monotonicity or single-peakedness of U(X) to the monotonicity or single-peakedness of $g^*(X, 0)$.

²⁶Note that since *U* is twice continuously differentiable, if U'(X) = 0 for all *X* in some open interval (X_a, X_b) , then $U'(X_a) = U'(X_b) = 0$ and, therefore, U'(X) = 0 on closed interval $[X_a, X_b]$. In the case of $X_a = X_b$, this definition corresponds to a strict extremum point.

PROPOSITION 2. Assume Assumptions 1 to 3 and g bounded. When $\mathcal{X} = \mathbb{R}$, assume Assumption 4, and when \mathcal{X} is compact, assume the boundary states $\{\underline{X}, \overline{X}\}$ are absorbing. Let U(X) denote the unique bounded solution to (7).

- (i) The term $g^*(X, 0)$ is constant on \mathcal{X} if and only if U(X) is constant on \mathcal{X} .
- (ii) If $g^*(X, 0)$ is monotonically increasing (decreasing) on \mathcal{X} , then U(X) is monotonically increasing (decreasing) on \mathcal{X} .
- (iii) If $g^*(X, 0)$ is single-peaked with a unique interval maximum (minimum) and $g^*(\underline{X}, 0) = g^*(\overline{X}, 0)$ (or, in the case of \mathcal{X} , unbounded, $\lim_{X\to\infty} g^*(X, 0) = \lim_{X\to-\infty} g^*(X, 0)$), then U(X) is single-peaked with a unique interval maximum (minimum).
- (iv) If $g^*(X, 0)$ has N interval extrema on \mathcal{X} , then U(X) has at most N interval extrema on \mathcal{X} .

Applying Propositions 1 and 2 to specific applications will yield structural empirical predictions about how equilibrium behavior and payoffs change with the state. This is illustrated in Section 6.1 when the static Nash payoff is monotonic and in Section 6.2 when the static Nash payoff is single-peaked.

Proposition 3 establishes a bound on the PPE payoff across all states when the continuation value converges to the static Nash payoff at the boundary states. Let $\overline{W} \equiv \sup_{X \in \mathcal{X}} U(X)$ and $\underline{W} \equiv \inf_{X \in \mathcal{X}} U(X)$ be the least upper bound and greatest lower bound of the large player's PPE payoff across all states, and let \mathcal{X}_H and \mathcal{X}_L denote the sets of states that yield these payoffs (where, in a slight abuse of notation, I say $\infty \in \mathcal{X}_H$ if $\lim_{X \to \infty} U(X) = \overline{W}$ and similarly for $-\infty$ and the case of \mathcal{X}_L). The following result shows that the smallest static Nash payoff in \mathcal{X}_H bounds the PPE payoff from above and, similarly, the largest static Nash payoff in \mathcal{X}_L bounds the PPE payoff from below.²⁷

PROPOSITION 3. Assume Assumptions 1 to 3 and g bounded. When $\mathcal{X} = \mathbb{R}$, assume Assumption 4, and when \mathcal{X} is compact, assume the boundary states $\{\underline{X}, \overline{X}\}$ are absorbing. Then the PPE payoff is bounded above (below) by the least (greatest) static Nash payoff at the states that yield the highest (lowest) PPE payoff,

$$\sup_{X\in\mathcal{X}_L}g^*(X,0)\leq\underline{W}\leq\overline{W}\leq\inf_{X\in\mathcal{X}_H}g^*(X,0),$$

where, in a slight abuse of notation, if $X \in \{-\infty, \infty\}$, then $g^*(X, 0)$ corresponds to $\lim_{x\to X} g^*(x, 0)$.

These bounds follow directly from the optimality equation. To see this, consider the case in which there is an interior state X_H such that $\overline{W} = U(X_H)$. Then $U'(X_H) =$

²⁷In general, it may be difficult to characterize \mathcal{X}_H from the primitives of the game, as \mathcal{X}_H does not necessarily correspond to the set of states that maximizes the static Nash payoff. A weaker bound that can be easily characterized is the highest static Nash payoff across all states, $\overline{W} \leq \sup_{\mathcal{X}} g^*(\mathcal{X}, 0)$, and similarly, $\underline{W} \geq \inf_{X \in \mathcal{X}} g^*(\mathcal{X}, 0)$.

0 and $U''(X_H) \le 0$, which from (7) implies $U(X_H) \le g^*(X_H, 0)$. This yields the upper bound. If the continuation value is sufficiently flat around X_H (i.e., $U''(X_H) = 0$), then $\overline{W} = g^*(X_H, 0)$. Otherwise, the continuation value or the state changes too quickly to maintain $g^*(X_H, 0)$ and $\overline{W} < g^*(X_H, 0)$.

This bound illustrates an important feature of the structural incentives generated by persistence. When selecting an action, the large player actively manages her tradeoff between short-run and long-run gains. At state(s) that maximize her long-run payoff, she rests on her laurels and focuses on short-run gains. In essence, incentives collapse at the top and she acts myopically. Therefore, she cannot earn a payoff higher than her best myopic payoff at this state (i.e., the static Nash payoff). This intuition for "shirking at the top" is similar in spirit to the reputation dynamics in Mailath and Samuelson (2001).

6. Applications

This section develops several applications to illustrate the breadth of the model. Section 6.1 presents two variations on the product choice setting introduced in Section 2, Section 6.2 presents an application in which a government selects a policy to target a persistent economic variable, and Section 6.3 presents an application in which the strategic complementarity between the investments of a government and a group of innovators leads to multiple MPE.

6.1 Variations of persistent quality

Consumer budget constraint This application modifies Example 1 to illustrate a setting with a bounded flow payoff. In Example 1, consumers' willingness to pay increases linearly with quality q(a, X). Now suppose that the marginal value of quality is decreasing. In particular, each consumer has a budget constraint and is willing to pay up to \overline{B} for the product, $B = [0, \overline{B}]$. The best response is the same as in Section 2, except now $b_i = \overline{B}$ if $q(\tilde{a}, X) \ge \overline{B}$. Payoffs and the drift of the stock quality are as defined in Section 2.

This modified version continues to satisfy Assumption 4. The boundary conditions at $-\infty$ are as before, but the boundary conditions at ∞ differ. Now $g_1(X) = 0$ for sufficiently large X and so $z_{\infty} = 0$. Therefore, equilibrium effort approaches zero and the continuation value converges to the limit of the static Nash payoff, \overline{B} , as X approaches ∞ . Given g is bounded, Proposition 2 can be used to characterize the shape of the continuation value. The static Nash payoff is $g^*(X, 0) = \max\{0, \lambda X\}$, which is monotonically increasing when $\lambda > 0$. Therefore, the continuation value is monotonically increasing on \mathcal{X} .

Figure 1(a) plots equilibrium effort, as derived from the equilibrium characterization in Theorem 1. The firm has the strongest incentive to invest at intermediate quality stock levels, as revenue rapidly increases with quality. This reputation building phase is characterized by high effort and rising quality. When the firm has high stock quality, the consumers' budget constraints prevent the firm from continuing to benefit from building its quality. The firm rides its good reputation by enjoying high payoffs today at the expense of allowing the quality to drift down. Very negative shocks lead to periods

Theoretical Economics 19 (2024)



FIGURE 1. Parameters: $\lambda = 0.8$, $\overline{B} = 20$, $\theta = 0.2$.

of reputation recovery where consumers stop purchasing for a time. The firm chooses low effort and allows the negative shock to dissipate before beginning to rebuild. Quality is stable when effort exactly offsets decay or, mathematically, when the drift is zero ($a = \theta X$). As the firm becomes more patient, this switch from reputation building to reputation riding occurs at a higher level of quality.

Figure 1(b) illustrates the equilibrium continuation value for several discount rates. It is convex at low levels of the stock quality and concave at high levels. When the stock quality is high, consumers are purchasing near their maximum level. The firm is risk averse in quality, in that negative quality shocks reduce revenue more than positive quality shocks increase revenue. On the other hand, when the stock quality is low, the firm faces the potential for substantial gains if quality rises, but the risk of loss from a negative quality shock is small. The continuation payoff has an interesting non-monotonicity with respect to the discount rate. As the firm becomes more patient, it places greater weight on the future, which gives it a stronger incentive to choose high effort and build its quality. On the other hand, as it becomes more patient, it values transitory positive quality shocks less. When stock quality is low, the first effect dominates and low discount rates yield higher payoffs; this relationship flips at high levels of quality.

Quality specialization This application adapts Example 1 to a setting with quality indivisibilities, which are modeled as intervals of quality in which the return to quality is flat. For example, the marginal value of an upgrade to a new software version may be larger for some versions and smaller for others, while the cost of developing an upgrade is constant. This illustrates how the characterization in Theorem 1 can be used to study equilibrium dynamics in an environment that differs from the standard setting with constant or decreasing returns to quality, as typically studied in the reputation and dynamic games literature.

To model this, we use a simple parameterization of quality that captures the idea that the marginal return to quality is non-monotonic in the state, while maintaining the property that the overall return to quality is increasing. Each consumer's expected

value for the product is $q(\tilde{a}, X) + \sin q(\tilde{a}, X)$, where $q(\tilde{a}, X) = (1 - \lambda)\tilde{a} + \lambda X$ as before.²⁸ The expected value of quality is increasing in *X*, so higher stock quality is always more valuable, but the marginal value of an increase varies between being relatively high and relatively low. As in the previous application, the consumer faces a budget constraint, $B = [0, \overline{B}]$. The remainder of the model is as defined in Section 2.

Using Theorem 1 to characterize the Markov equilibrium in this setting shows that these quality indivisibilities lead to novel dynamics related to the firm's incentive to invest. A key feature is that a firm may specialize in providing intermediate or low quality, rather than always striving to provide high quality. This stems from the firm's nonmonotonic incentive to exert effort, which leads to multiple regions of the state space in which the firm cycles between building and dissipating quality. When the stock quality is such that the marginal return is high, the firm invests in building its quality. Once the firm reaches a region where the marginal return is flat, it slacks off and chooses lower effort. With positive probability, quality drifts back down to a level at which the firm has an incentive to invest again. But also with positive probability, quality continues to rise and the firm reaches a new level of quality with a high marginal return. The firm then begins investing to maintain this new, higher level of quality. This leads the firm to specialize in different levels of quality. A low quality firm may be better off remaining a low quality firm, rather than trying and failing to move up the market. But if a firm has a positive shock and reaches a high quality level, it will then have the incentive to invest in maintaining this higher quality. This variation has the same properties as the consumer budget constraint variation for large and small X. Therefore, Assumption 4 is satisfied and Theorem 3 guarantees that this MPE is unique.

6.2 Targeting an Economic Variable

This application illustrates a setting in which the state space is compact and the equilibrium payoff is single-peaked. Suppose constituents elect a board to implement a policy to target an economic variable. Elected officials and governing bodies often play a role in formulating and implementing such policies. For example, the Federal Reserve targets an interest rate, a board of directors sets growth and return targets for its company, and the housing authority targets home ownership rates. Achieving such targets requires costly effort on behalf of officials, and often the level of the variable will depend on both current and past policy efforts. Moral hazard issues arise when the officials' preferences are not aligned with the population they serve.

To model this, consider a setting in which the state is an economic variable that takes on values in $\mathcal{X} = [0, 2]$. Constituents want to target X = 1, but in the absence of intervention, the state drifts toward its natural level $d \in [0, 2]$. The board can undertake costly intervention $a \in [-1, 1]$ to alter the state. The state has drift $\mu(a, b, X) = X(2-X)(a + \theta(d - X))$, where $\theta > 0$ captures the persistence of past interventions, and

²⁸One could also use a parameterization of quality with kinks, where the marginal return to quality is zero on some intervals of stock quality and positive on others, provided the parameterization is Lipschitz continuous. Such a parameterization would yield qualitatively similar equilibrium dynamics.

volatility $\sigma(b, X) = X(2 - X)$; it is most volatile at intermediate levels. A negative intervention decreases *X*, while a positive intervention increases *X*. Constituents choose an action each period that represents their campaign contributions or support for the board. When constituents believe the board chooses intervention \tilde{a} and the economic variable is equal to *X*, they are willing to contribute $\lambda \tilde{a}^2 + 1 - (1 - X)^2$, where $\lambda > 0$ captures the weight placed on an intervention. This contribution is a reduced form representation of the constituents' preferences: they pledge higher support when *X* is closer to their preferred target and when the board undertakes a stronger intervention. The board has no direct preference over the economic variable. Its flow payoff is increasing in the support it receives from constituents and decreasing in the cost of intervention, $g(a, \overline{b}, X) = \overline{b} - ca^2$.

The state space is compact and the boundary states are absorbing, so uniqueness of a MPE follows from Theorem 4. Using Theorem 1 to characterize the Markov equilibrium and Proposition 2(iii) to characterize the shape of the continuation value establishes that the board will intervene to increase the economic variable at low states and intervene to decrease the variable at high states, and that the continuation value is single-peaked with a maximum and is not constant on any interval of states. The point at which the board switches from a positive to a negative intervention depends on the constituents' target and the natural drift d. If the drift lies above the target, d > 1, then when the state is low, it will naturally move toward the target. This benefits the board. At very low levels, the board chooses a positive intervention to increase the rate at which the state moves toward the target. It switches to a negative intervention when the state is slightly below the target in order to prevent the state from overshooting its target. The opposite holds when d < 1. The board has the strongest incentive to intervene when the economic variable is an intermediate distance from its target, in which case the state is sensitive to an intervention and the benefit from intervening is high. When the economic variable is far from its target, an intervention has a small impact on the state and the board has a low incentive to intervene. When the economic variable is close to its target, the continuation value is flat so the benefit from an intervention is low and the board again has a low incentive to intervene. Figure 2 plots the equilibrium intervention and continuation value as a function of the state for two levels of d. See Supplemental Appendix E.1 for derivations not contained in the text.

6.3 Complementary investment and multiple equilibria

This example illustrates how multiple Markov equilibria can arise in a setting with an unbounded flow payoff. In this example, $g^*(X, z)$ and $\mu^*(X, z)$ are not additively separable, which violates Assumption 4(i). Complementarities between the direct and equilibrium channels for incentives create coordination motives that give rise to multiple equilibrium incentive weights, each associated with a different optimal action profile.

Suppose a government and a sequence of small innovators can invest to generate intellectual capital. The state *X* represents the current level of intellectual capital in the economy. The government chooses an investment level $a \in [0, \overline{a}]$, where $\overline{a} > 0$ is the maximum feasible investment for the government. Each innovator chooses investment

Dynamic moral hazard game 477



FIGURE 2. Parameters: $\lambda = 0.5$, $\theta = 0.4$, c = 1, r = 0.1.

 $b_i \in [-\gamma|X|, \gamma|X|]$, where $\gamma > 0$ and the bound on feasible investment is proportional to available intellectual capital. Both government and innovator investment contribute to the growth of intellectual capital, with returns $\theta_1 > 0$ and $\theta_2 > 0$, respectively. Intellectual capital depreciates at rate $\theta_3 > 0$. Therefore, the expected change in intellectual capital is $\mu(a, b, X) = \theta_1 b + \theta_2 a - \theta_3 X$. Assume the volatility of intellectual capital is constant, $\sigma(b, X) = 1$, and $\gamma < (r + \theta_3)/\theta_1$ to satisfy Assumption 2.²⁹

For an innovator, government investment is a strategic complement with her own investment and the current level of intellectual capital. For example, when an innovator invests in a new project, her return depends on both the stock of intellectual capital in the economy and the investment from the government to make this intellectual capital accessible. This is captured by payoff $h(a, b, \overline{b}, X) = abX - cb^2/2$, where the first term captures the property that intellectual capital is only valuable to an innovator if both the innovator and the government invest, and the second term captures the cost of investment for some c > 0. The government receives a return of $\alpha > 0$ on each innovator's investment. Therefore, even though it does not directly value intellectual capital, it values it indirectly through its impact on future investment. For example, α is the tax rate on investment; the government is willing to invest today if this increases future tax returns. This is captured by payoff $g(a, b, X) = \alpha b - a^2/2$, where the second term captures the cost of investment. For technical reasons, assume $\overline{a} \leq \gamma c.^{30}$

From Lemma 1, the sequentially rational investment level for the government is

$$a(X, z) = \begin{cases} \theta_2 z/r & \text{if } z/r \in [0, \overline{a}/\theta_2] \\ \overline{a} & \text{if } z/r > \overline{a}/\theta_2 \\ 0 & \text{if } z < 0. \end{cases}$$

²⁹Note that innovator investment can be unboundedly negative. While unrealistic, this specification yields a closed-form solution for the continuation value, which makes it straightforward to illustrate the existence of multiple equilibria. In the more realistic case that the lower bound on innovator investment is zero, the equilibrium characterization is qualitatively similar.

³⁰This guarantees that an interior solution is always feasible for the innovator.

Investment is increasing in the impact that it has on the growth of intellectual capital θ_2 and the incentive weight z/r. When an innovator believes that the government will choose investment level \tilde{a} and the current stock of intellectual capital is X, the innovator's best response is to select investment $\tilde{a}X/c$. The innovator's investment is increasing in the investment of the government and the current stock of intellectual capital, reflecting the complementarity of these two inputs. If $z/r \in [0, \overline{a}/\theta_2]$, then the government chooses an interior level of investment, yielding

$$g^*(X,z) = \left(\frac{\alpha\theta_2 z}{cr}\right) X - \frac{\theta_2^2 z^2}{2r^2}$$
(14)

and

$$\mu^*(X,z) = \left(\frac{\theta_1 \theta_2 z}{cr}\right) X + \frac{\theta_2^2 z}{r} - \theta_3 X.$$
(15)

Neither expression is additively separable in (X, z), so Assumption 4 does not hold.

We first use Theorem 1 to show that there is an equilibrium in which neither the government nor the innovators invest, a(X) = b(X) = 0 for all X, and the government's equilibrium payoff is U(X) = 0. Due to the strategic complementarity, if the government does not invest, then neither will the innovators, yielding a payoff of zero for all players.

We next use Theorem 1 to show that there can also be nontrivial equilibria that sustain positive investment. When $\gamma \approx (r + \theta_3)/\theta_1$, $\overline{a} = \gamma c$, and $c\theta_3 > \alpha\theta_2$, there exists an equilibrium that has nonzero equilibrium investment levels, $a(X) = (cr - \alpha\theta_2 + c\theta_3)/\theta_1$ and b(X) = a(X)X/c, and continuation value

$$U(X) = r\left(\frac{cr - \alpha\theta_2 + c\theta_3}{\theta_1\theta_2}\right)X + \frac{(cr - \alpha\theta_2 + c\theta_3)^2}{2\theta_1^2}.$$

The slope captures the government's net present value of the current stock of intellectual capital, while the constant term captures the equilibrium effect stemming from the value of future strategic interaction between the government and the innovators. This latter effect is positive, given $c\theta_3 > \alpha\theta_2$. As the government becomes arbitrarily patient, $r \rightarrow 0$, the net present value of the current stock of intellectual capital approaches zero and $U(X) \rightarrow (c\theta_3 - \alpha\theta_2)^2/2\theta_1^2$. Intuitively, the government cares more about the long-run return from the strategic interaction rather than the short-run return from the current stock of intellectual capital interpretation: it is the equilibrium flow payoff for the patient government when the stock of intellectual capital is at its long-run average, which depends on equilibrium investment and the rate of depreciation. Additionally, when $\alpha > (cr - \theta_1 \overline{a} + c\theta_3)/\theta_2$, there is an equilibrium with investment levels $a(X) = \overline{a}$ and $b(X) = \overline{a}X/c$. See Supplemental Appendix E.2 for both derivations.

This example illustrates that both trivial and nontrivial Markov equilibria can exist when there are complementarities between the players' actions. Even when uniqueness does not hold, Theorem 1 can be used to characterize these Markov equilibria and Theorem 2 can be used to characterize the PPE payoff set.

7. Conclusion

This paper shows that persistence provides an important channel for intertemporal incentives and develops a tractable method to characterize Markov equilibrium behavior and payoffs. The tools developed in this paper will yield insights into equilibrium behavior in a broad range of settings, from industrial organization to political economy to macroeconomics. Once functional forms are specified for payoffs and the evolution of the state, it is straightforward to use Theorem 1 to construct Markov equilibria. This in turn can be used to derive empirically testable comparative statics and predictions about the dynamics of equilibrium behavior based on observable features of the environment. Future research can use this framework to address design questions in specific applications, such as determining the optimal structure of persistence in a rating mechanism.

Furthermore, the equilibrium characterization can be used for structural estimation. Markov equilibria have an intuitive appeal in empirical work due to their simplicity and dependence on payoff-relevant variables to structure incentives. Players do not need to condition on past behavior in a complex way, as actions and payoffs are fully determined by the current value of the state. Establishing that a Markov equilibrium exists and is unique provides a strong justification for focusing on this equilibrium concept, while the equilibrium characterization yields expressions for payoffs and actions that can calibrated and estimated.

Appendix A: Proofs

A.1 Proof of Lemma 1

I first show that $(V_t(S))_{t\geq 0}$ is a martingale and $(W_t(S))_{t\geq 0}$ is bounded with respect to $(X_t)_{t\geq 0}$.

CLAIM 1. Under Assumption 2, for any public strategy profile $S = (a_t, \overline{b}_t)_{t\geq 0}$, initial state X_0 , and path of the state variable $(X_t)_{t\geq 0}$ that evolves according to (1) given S, $V_t(S)$ is a martingale and there exists a $K_W > 0$ such that $|W_t(S)| \leq K_W(1 + |X_t|)$ for all $t \geq 0$.

Suppose *g* is unbounded. By Assumption 2, there exists a $k \in [0, r)$ and c > 0 such that for all $(a, b, X) \in A \times E$, if $X \ge 0$, then $\mu(a, b, X) \le kX + c$, and if $X \le 0$, then $\mu(a, b, X) \ge kX - c$. Lipschitz continuous functions have linear growth. Therefore, by Lipschitz continuity of *g* and σ , the compactness of *A*, and the assumption that $|b| \le K_b |X| + c_b$ for all $(b, X) \in E$, there exists a $K_g, K_\sigma, c > 0$ such that for all $(a, b, X) \in A \times E$, $|g(a, b, X)| \le K_g (\frac{c}{k} + |X|)$ and $|\sigma(b, X)| \le K_\sigma (1 + |X|)$.

I first derive a bound on $E_{\tau}|g(a_t, \overline{b}_t, X_t)|$, the expected flow payoff at time *t* conditional on available information at time $\tau \leq t$. This bound will be independent of the

strategy profile. Define $f : \mathcal{X} \to \mathbb{R}$ as

$$f(X) \equiv \begin{cases} K_g \left(\frac{c}{k} - X\right) & \text{if } X \leq -1 \\ -\frac{1}{8} K_g X^4 + \frac{3}{4} K_g X^2 + \frac{3}{8} K_g + K_g \frac{c}{k} & \text{if } X \in (-1, 1) \\ K_g \left(\frac{c}{k} + X\right) & \text{if } X \geq 1. \end{cases}$$

Note that $f \in C^2$, $f \ge 0$, $|f'| \le K_g$, and

$$f''(X) = \begin{cases} 0 & \text{if } |X| \ge 1\\ \frac{3}{2}K_g(1 - X^2) & \text{if } |X| < 1. \end{cases}$$

Ito's lemma holds for any C^2 function. Given a strategy profile $S = (a_t, \overline{b}_t)_{t \ge 0}$, initial state $X_\tau < \infty$, and path of the state variable $(X_t)_{t \ge \tau}$ that evolves according to (1),

$$f(X_t) = f(X_\tau) + \int_{\tau}^{t} \left(f'(X_s)\mu(a_s, \overline{b}_s, X_s) + \frac{1}{2}f''(X_s)\sigma(\overline{b}_s, X_s)^2 \right) ds$$

+
$$\int_{\tau}^{t} f'(X_s)\sigma(\overline{b}_s, X_s) dZ_s$$

$$\leq f(X_\tau) + \int_{\tau}^{t} \left(K_g(k|X_s| + c) + 3K_gK_{\sigma}^2 \right) ds + K_gK_{\sigma} \int_{\tau}^{t} \left(1 + |X_s| \right) dZ_s$$

$$\leq f(X_\tau) + k \int_{\tau}^{t} f(X_s) ds + 3K_gK_{\sigma}^2(t - \tau) + K_gK_{\sigma} \int_{\tau}^{t} \left(1 + |X_s| \right) dZ_s$$

for all $t \ge \tau$, where the first inequality follows from $f'(X)\mu(a, b, X) \le K_g(k|X| + c)$, $\frac{1}{2}f''(X)\sigma(b, X)^2 \le 3K_gK_{\sigma}^2$, and $f'(X)\sigma(b, X)z \le K_gK_{\sigma}(1 + |X|)z$ for all $z \in \mathbb{R}$ and for all $(a, b, X) \in A \times E$, and the second inequality follows from the definition of f. The addition of the absolute value sign in $f'(X)\mu(a, b, X) \le K_g(k|X| + c)$ follows from the sign of f', and the bound on $\frac{1}{2}f''(X)\sigma(b, X)^2$ follows from $f''(X)\sigma(b, X)^2 = 0$ if $|X| \ge 1$ and

$$f''(X)\sigma(b,X)^2 = \frac{3}{2}K_g(1-X^2)\sigma(b,X)^2 \le \frac{3}{2}K_g(1-X^2)K_\sigma^2(1+|X|)^2 \le 6K_gK_\sigma^2$$

if |X| < 1. Taking expectations and noting that $(1 + |X_s|)$ is square-integrable on $[\tau, t]$, so the expectation of the stochastic integral is zero,

$$E_{\tau}[f(X_t)] \leq f(X_{\tau}) + 3K_g K_{\sigma}^2(t-\tau) + k \int_{\tau}^t E_{\tau}[f(X_s)] ds$$
$$\leq (f(X_{\tau}) + 3K_g K_{\sigma}^2(t-\tau)) e^{k(t-\tau)},$$

where the last line follows from Gronwall's inequality. Note that $|g(a, b, X)| \le f(X)$ for all $(a, b, X) \in A \times E$. Therefore,

$$e^{-r(t-\tau)}E_{\tau}|g(a_t, \overline{b}_t, X_t)| \le e^{-r(t-\tau)}E_{\tau}[f(X_t)] \le (f(X_{\tau}) + 3K_gK_{\sigma}^2(t-\tau))e^{-(r-k)(t-\tau)}.$$

I next show that if $X_t < \infty$, then $W_t(S) < \infty$,

$$|W_t(S)| = \left| E_t \left[r \int_t^\infty e^{-r(s-t)} g(a_s, \overline{b}_s, X_s) \, ds \right] \right|$$

$$\leq r \int_t^\infty e^{-r(s-t)} E_t \left| g(a_s, \overline{b}_s, X_s) \right| \, ds$$

$$\leq r \int_t^\infty \left(f(X_t) + 3K_g K_\sigma^2(s-t) \right) e^{-(r-k)(s-t)} \, ds$$

$$= \left(\frac{r}{r-k} \right) f(X_t) + \frac{3rK_g K_\sigma^2}{(r-k)^2},$$

which is finite for any $X_t < \infty$ and k < r. Also, given that f has linear growth, there exists a $K_W > 0$ such that $|W_t(S)| \le K_W(1 + |X_t|)$. By similar reasoning, $E|V_t(S)| < \infty$ for any $X_0 < \infty$ since

$$E|V_t(S)| = E\left|E_t\left[r\int_0^\infty e^{-rs}g(a_s,\overline{b}_s,X_s)\,ds\right]\right| \le E\left[r\int_0^\infty e^{-rs}\left|g(a_s,\overline{b}_s,X_s)\right|\,ds\right]$$

is finite for any $X_0 < \infty$ and k < r.

Finally,

$$E_t[V_{t+k}(S)] = E_t\left[r\int_0^{t+k} e^{-rs}g(a_s,\overline{b}_s,X_s)\,ds + e^{-r(t+k)}W_{t+k}(S)\right]$$
$$= r\int_0^t e^{-rs}g(a_s,\overline{b}_s,X_s)\,ds$$
$$+ E_t\left[r\int_t^{t+k} e^{-rs}g(a_s,\overline{b}_s,X_s)\,ds\right.$$
$$+ e^{-r(t+k)}E_{t+k}\left[r\int_{t+k}^{\infty} e^{-r(s-(t+k))}g(a_s,\overline{b}_s,X_s)\,ds\right]\right]$$
$$= r\int_0^t e^{-rs}g(a_s,\overline{b}_s,X_s)\,ds + e^{-rt}W_t(S) = V_t(S).$$

Taken together, this implies $V_t(S)$ is a martingale and establishes Claim 1 for the case of g unbounded. If g is bounded, then trivially, $W_t(S) < \infty$ and $E|V_t(S)| < \infty$ for all $t \ge 0$ and $X_0 \in \mathcal{X}$, and only the final step is needed to establish the claim.

I next derive the evolution of the continuation value. This part of the proof follows from almost identical reasoning to the proof of Proposition 2 in Faingold and Sannikov (2011). The derivative of $V_t(S)$ with respect to t is:

$$dV_t(S) = re^{-rt}g(a_t, \overline{b}_t, X_t) dt - re^{-rt}W_t(S) dt + e^{-rt} dW_t(S).$$

By the martingale representation theorem (Karatzas and Shreve (1991)), there exists a progressively measurable process $(\beta_t)_{t\geq 0}$ such that V_t can be represented as $dV_t(S) =$

 $re^{-rt}\beta_t\sigma(\overline{b}_t, X_t) dZ_t$. Combining these two expressions for $dV_t(S)$ yields the law of motion for the continuation value,

$$dW_t(S) = r(W_t(S) - g(a_t, b_t, X_t)) dt + r\beta_t \sigma(b_t, X_t) dZ_t$$

= $r(W_t(S) - g(a_t, \overline{b}_t, X_t)) dt + r\beta_t (dX_t - \mu(a_t, \overline{b}_t, X_t) dt),$

where β_t captures the sensitivity of the continuation value to the state variable. As shown above, any continuation value has linear growth with respect to X_t and is bounded when *g* is bounded.

Finally, I establish sequential rationality. This part of the proof follows from almost identical reasoning to the proof of Proposition 3 in Faingold and Sannikov (2011). Consider strategy profile $(a_t, \overline{b}_t)_{t\geq 0}$ played from period τ onward and alternative strategy $(\widetilde{a}_t, \overline{b}_t)_{t\geq 0}$ played up to time τ . Recall that all values of X_t are possible under both strategies, but that each strategy induces a different measure over sample paths $(X_t)_{t\geq 0}$. At time τ , the state variable is equal to X_{τ} . Action a_{τ} will induce $dX_{\tau} = \mu(a_{\tau}, \overline{b}_{\tau}, X_{\tau}) dt + \sigma(\overline{b}_{\tau}, X_{\tau}) dZ_{\tau}$, whereas action \widetilde{a}_{τ} will induce $dX_{\tau} = \mu(\widetilde{a}_{\tau}, \overline{b}_{\tau}, X_{\tau}) dt + \sigma(\overline{b}_{\tau}, X_{\tau}) dZ_{\tau}$. Let \widetilde{V}_{τ} be the expected average payoff conditional on information at time τ when the large player follows \widetilde{a} up to τ and a afterward, and let W_{τ} be the continuation value when the large player follows strategy $(a_t)_{t\geq 0}$ starting at time τ :

$$\widetilde{V}_{\tau} = r \int_0^{\tau} e^{-rs} g(\widetilde{a}_s, \overline{b}_s, X_s) \, ds + e^{-r\tau} W_{\tau}.$$

Consider changing τ so that the large player plays strategy (\tilde{a}_t, \bar{b}_t) for another instant: $d\tilde{V}_{\tau}$ is the change in average expected payoffs when the large player switches to $(a_t)_{t\geq 0}$ at $\tau + d\tau$ instead of τ . When the large player switches strategies at time τ ,

$$\begin{split} d\widetilde{V}_{\tau} &= re^{-r\tau} \big(g(\widetilde{a}_{\tau}, \overline{b}_{\tau}, X_{\tau}) - W_{\tau} \big) \, d\tau + e^{-r\tau} \, dW_{\tau} \\ &= re^{-r\tau} \big(g(\widetilde{a}_{\tau}, \overline{b}_{\tau}, X_{\tau}) - g(a_{\tau}, \overline{b}_{\tau}, X_{\tau}) \big) \, d\tau + re^{-r\tau} \beta_{\tau} \big(dX_{\tau} - \mu(a_{\tau}, \overline{b}_{\tau}, X_{\tau}) \, d\tau \big) \\ &= re^{-r\tau} \big(g(\widetilde{a}_{\tau}, \overline{b}_{\tau}, X_{\tau}) - g(a_{\tau}, \overline{b}_{\tau}, X_{\tau}) + \beta_{\tau} \mu(\widetilde{a}_{\tau}, \overline{b}_{\tau}, X_{\tau}) - \beta_{\tau} \mu(a_{\tau}, \overline{b}_{\tau}, X_{\tau}) \big) \, d\tau \\ &+ re^{-r\tau} \beta_{\tau} \sigma(\overline{b}_{\tau}, X_{\tau}) \, dZ_{\tau}. \end{split}$$

There are two components to this strategy change: how it affects the immediate flow payoff and how it affects the future state X_t , which impacts the continuation value. The profile $(\tilde{a}_t, \bar{b}_t)_{t\geq 0}$ yields the large player a payoff of

$$\begin{aligned} \widetilde{W}_0 &= E_0[\widetilde{V}_\infty] = E_0 \bigg[\widetilde{V}_0 + \int_0^\infty d\widetilde{V}_t \bigg] \\ &= W_0 + E_0 \bigg[r \int_0^\infty e^{-rt} \big(g(\widetilde{a}_t, \overline{b}_t, X_t) + \beta_t \mu(\widetilde{a}_t, \overline{b}_t, X_t) - g(a_t, \overline{b}_t, X_t) \big] \\ &- \beta_t \mu(a_t, \overline{b}_t, X_t) \big) dt \bigg]. \end{aligned}$$

If

$$g(a_t, \overline{b}_t, X_t) + \beta_t \mu(a_t, \overline{b}_t, X_t) \ge g(\widetilde{a}_t, \overline{b}_t, X_t) + \beta_t \mu(\widetilde{a}_t, \overline{b}_t, X_t)$$

holds for all $t \ge 0$, then $W_0 \ge \widetilde{W}_0$ and deviating to $S = (\widetilde{a}_t, \overline{b}_t)$ is not a profitable deviation. A strategy $(a_t)_{t\ge 0}$ is sequentially rational for the large player if, given $(\beta_t)_{t\ge 0}$, for all t,

$$a_t \in \arg\max_{a \in A} g(a, \overline{b}_t, X_t) + \beta_t \mu(a, \overline{b}_t, X_t)$$

A.2 Proof of Theorem 1

In a Markov equilibrium, the continuation value and equilibrium actions are characterized as a function of the state variable as $W_t = U(X_t)$, $a_t^* = a(X_t)$, and $\overline{b}_t^* = b(X_t)$. By Ito's formula, if a Markov equilibrium with a twice continuously differentiable continuation value exists, the continuation value will evolve according to

$$dU(X_{t}) = U'(X_{t}) dX_{t} + \frac{1}{2}U''(X_{t})\sigma(\overline{b}_{t}^{*}, X_{t})^{2} dt$$

$$= U'(X_{t})\mu(a_{t}^{*}, \overline{b}_{t}^{*}, X_{t}) dt + \frac{1}{2}U''(X_{t})\sigma(\overline{b}_{t}^{*}, X_{t})^{2} dt$$

$$+ U'(X_{t})\sigma(\overline{b}_{t}^{*}, X_{t}) dZ_{t}.$$
 (16)

Similar to the derivation of a Markov equilibrium in Faingold and Sannikov (2011), matching the drift of (16) with the drift of the continuation value characterized in (5) yields the optimality equation

$$U''(X) = \frac{2r(U(X) - g(a(X), b(X), X))}{\sigma(b(X), X)^2} - \frac{2\mu(a(X), b(X), X)U'(X)}{\sigma(b(X), X)^2},$$
(17)

which is a second-order nonhomogenous differential equation, and matching the volatilities characterizes the process governing incentives, $r\beta_t = U'(X_t)$. Substituting this expression into the condition for sequential rationality characterized in (6) yields the Markovian action profile $(a(X), b(X)) = S^*(X, U'(X))$ (by Assumption 3, S^* is single-valued.) Plugging this into (17) yields (7).

I first establish that (7) has at least one solution $U \in C^2$ that takes on values in the interval of feasible payoffs for the large player. In the case of an unbounded state space, Theorem 5.6 from De Coster and Habets (2006) gives sufficient conditions for the existence of a solution to a second-order differential equation defined on \mathbb{R}^3 . I construct upper and lower solutions to (17) at action profile $S^*(X, U'(X))$ to show that these conditions are satisfied. This leads to the following lemma, which is the innovative part of this proof and is proven in Supplemental Appendix B.

LEMMA 2. If $\mathcal{X} = \mathbb{R}$, then (7) has at least one solution $U \in C^2$ on \mathcal{X} that lies in the range of feasible payoffs for the large player.

In the case of a bounded state space, I use an extension of a standard existence result from De Coster and Habets (2006), which was developed in Faingold and Sannikov (2011). The extension is necessary because (7) is undefined at the boundary of the state space, { X, \overline{X} }. This leads to the following lemma, which is proven in Supplemental Appendix B.

LEMMA 3. If X is compact, then (7) has at least one solution $U \in C^2$ on $(\underline{X}, \overline{X})$ that lies in the range of feasible payoffs for the large player.

Finally, I construct a Markov equilibrium that yields payoff $U(X_0)$, where U is a solution to (7). The function $X \mapsto S^*(X, U'(X))$ is Lipschitz continuous, as are $X \mapsto \mu^*(X, U'(X))$ and $X \mapsto \sigma^*(X, U'(X))$. Therefore, the state variable starts at X_0 and evolves according to the unique strong solution $(X_t)_{t\geq 0}$ to the stochastic differential equation

$$dX_t = \mu^* \big(X_t, U'(X_t) \big) dt + \sigma^* \big(X_t, U'(X_t) \big) dZ_t.$$

Moreover,

$$dU(X_t) = U'(X_t)\mu^* (X_t, U'(X_t)) dt + \frac{1}{2}U''(X_t)\sigma^* (X_t, U'(X_t))^2 dt + U'(X_t)\sigma^* (X_t, U'(X_t)) dZ_t = r (U(X_t) - g^* (X_t, U'(X_t))) dt + U'(X_t)\sigma^* (X_t, U'(X_t)) dZ_t$$

and, therefore, the process of continuation values $W_t = U(X_t)$ satisfies (5) with process of incentive weights $\beta_t = U'(X_t)/r$. Finally, the strategy profile $(a_t^*, \overline{b}_t^*)_{t\geq 0}$ satisfies (6) given $(\beta_t)_{t\geq 0}$ with $\beta_t = U'(X_t)/r$. Therefore, $(a_t^*, \overline{b}_t^*)_{t\geq 0}$ is a PPE yielding equilibrium payoff $U(X_0)$.

A.3 Proof of Theorem 2

Let *U* be the linear growth (when *g* is unbounded) or bounded (when *g* is bounded) solution to (7) that yields the highest MPE payoff at X_0 . Suppose there exists an initial state $X_0 \in \mathcal{X}$ and a PPE strategy profile $S = (a_t, \overline{b}_t)_{t \ge 0}$ that yields an equilibrium payoff $W_0 > U(X_0)$. In such a PPE, the state $(X_t)_{t \ge 0}$ evolves according to (1) given $S = (a_t, \overline{b}_t)_{t \ge 0}$ and, by Lemma 1, the continuation value evolves according to

$$dW_t(S) = r(W_t(S) - g(a_t, \overline{b}_t, X_t)) dt + r\beta_t(dX_t - \mu(a_t, \overline{b}_t, X_t) dt)$$
(18)

for some process $(\beta_t)_{t\geq 0}$. By Assumption 3, a unique action profile satisfies (6) at each $(X, r\beta)$. Therefore, by Lemma 1, equilibrium actions satisfy $(a_t, \overline{b}_t) = S^*(X_t, r\beta_t)$. By Ito's formula, the process $(U(X_t))_{t\geq 0}$ evolves according to

$$dU(X_{t}) = U'(X_{t})\mu^{*}(X_{t}, r\beta_{t}) dt + \frac{1}{2}U''(X_{t})\sigma^{*}(X_{t}, r\beta_{t})^{2} dt + U'(X_{t})\sigma^{*}(X_{t}, r\beta_{t}) dZ_{t}.$$
(19)

Define a process $D_t \equiv W_t(S) - U(X_t)$ with initial condition $D_0 = W_0(S) - U(X_0) > 0$. Then D_t evolves according to $dD_t = dW_t(S) - dU(X_t)$. Plugging in Eqs. (18) and (19), the process has volatility $f(X_t, \beta_t)$, where $f(X, \beta) \equiv (r\beta - U'(X))\sigma^*(X, r\beta)$, and has

drift $rD_t + d(X_t, \beta_t)$, where

$$d(X,\beta) \equiv r(U(X) - g^*(X,r\beta)) - U'(X)\mu^*(X,r\beta) - U''(X)\sigma^*(X,r\beta)^2/2$$

= $r(g^*(X,U'(X)) - g^*(X,r\beta)) + U'(X)(\mu^*(X,U'(X)) - \mu^*(X,r\beta))$
+ $U''(X)(\sigma^*(X,U'(X))^2 - \sigma^*(X,r\beta)^2)/2,$

and the second line follows from substituting the right hand side of (7) for U(X).

LEMMA 4. If $f(X, \beta) = 0$ and $\sigma(X, r\beta) > 0$, then $d(X, \beta) = 0$.

PROOF. Suppose $f(X, \beta) = 0$ for some (X, β) and $\sigma(X, r\beta) > 0$. Then $r\beta = U'(X)$. The action profile associated with $S^*(X, U'(X))$ corresponds to the actions played in the Markov equilibrium with continuation value U(X) at state X. Therefore, $d(X, \beta) = 0$.

LEMMA 5. For every $\varepsilon > 0$, there exists a $\eta > 0$ such that either $d(X, \beta) > -\varepsilon$ or $|f(X, \beta)| > \eta$.

PROOF. Suppose the state space is unbounded, $\mathcal{X} = \mathbb{R}$. Note that in this case, $\sigma^*(X, r\beta)$ is bounded away from 0 by Assumption 1, so, by Lemma 4, if $f(X, \beta) = 0$, then $d(X, \beta) =$ 0. First show that there exists an M > 0 such that this is true for $(X, \beta) \in \Omega_a \equiv \{\mathcal{X} \times \mathbb{R} :$ $|\beta| > M$. Since U' is bounded by Lemma 9 in the case of g unbounded or Lemma 26 in Supplemental Appendix D.2 in the case of g bounded (note that neither lemma requires Assumption 4), and since $\sigma^*(X, r\beta)$ is bounded away from 0, there exists an M > 0 and $\eta_1 > 0$ such that $|f(X, \beta)| > \eta_1$ for all $|\beta| > M$ and $X \in \mathcal{X}$, regardless of d. Next show that there exists an $\delta > 0$ such that this is true for $(X, \beta) \in \Omega_b \equiv \{\mathcal{X} \times \mathbb{R} : |\beta| \le M, |X| > \delta\}.$ Consider the set $\Phi_b \subset \Omega_b$ with $d(X, \beta) \leq -\varepsilon$. It must be that β is bounded away from U'(X)/r on Φ_b . Suppose not. Then either (i) there exists some $(X, \beta) \in \Phi_b$ with $\beta = U'(X)/r$, which implies $f(X, \beta) = 0$ and, therefore, $d(X, \beta) = 0$ —a contradiction or (ii) as *X* becomes large, the boundary of the set Φ_b approaches $\beta = U'(X)/r$. The latter implies that for any $\delta_1 > 0$, there exists an $(X, \beta) \in \Phi_h$ with $r\beta - U'(X) < \delta_1$. Choose δ_1 so that $|g^*(X, U'(X)) - g^*(X, r\beta)| < \varepsilon/4r, |U'(X)||\mu^*(X, U'(X)) - \mu^*(X, r\beta)| < \varepsilon/4,$ and $|U''(X)||\sigma^*(X, U'(X))^2 - \sigma^*(X, r\beta)^2| = 0$, which is possible given that g^* and μ^* are Lipschitz, U' is bounded, and σ^* is independent of z for large X. Then $|d(X,\beta)| < 1$ $\varepsilon/4 + \varepsilon/4 + \varepsilon/4 = 3\varepsilon/4$, which is a contradiction. Therefore, there exists a η_2 such that $|f(X,\beta)| > \eta_2$ on Φ_b . Then on the set Ω_b , if $d(X,\beta) \leq -\varepsilon$, then $|f(X,\beta)| > \eta_2$. Finally show this is true for $(X, \beta) \in \Omega_c \equiv \{\mathcal{X} \times \mathbb{R} : |\beta| \le M \text{ and } |X| \le \delta\}$. Consider the set $\Phi_c \subset \Omega_c$ where $d(X, \beta) \leq -\varepsilon$. The function d is continuous and Ω_c is compact, so Φ_c is compact. The function |f| is continuous and, therefore, achieves a minimum η_3 on Φ_c . If $\eta_3 = 0$, then d = 0 by Lemma 4—a contradiction. Therefore, $\eta_3 > 0$ and $|f(X, \beta)| > \eta_3$ for all $(X, \beta) \in \Phi_c$. Take $\eta \equiv \min\{\eta_1, \eta_2, \eta_3\}$. Then when $d(X, \beta) \leq -\varepsilon$, $|f(X, \beta)| > \eta$. The proof for a bounded state space is analogous (see Supplemental Appendix C).

LEMMA 6. Given X_0 , any PPE payoff W_0 is such that $W_0 \leq U(X_0)$.

PROOF. Choose $\varepsilon = rD_0/4$ and suppose $D_t \ge D_0/2$. Then, by Lemma 5, there exists a $\eta > 0$ such that whenever the drift of D_t is less than $rD_t - \varepsilon > rD_0/2 - rD_0/4 = rD_0/4 > 0$, $|f(X_t, \beta_t)| > \eta$. Thus, as long as $D_t \ge D_0/2 > 0$, it has either positive drift or positive volatility. This implies it grows arbitrarily large with positive probability, irrespective of X_t . This is a contradiction, since in the case that g is unbounded, by Lemma 1, D_t is the difference of two processes that are bounded with respect to X_t , and in the case that g is bounded, D_t is the difference of two bounded processes. Thus, it cannot be that $D_0 > 0$ and it must be the case that $W_0 \le U(X_0)$.

Letting *U* be the linear growth (when *g* is unbounded) or bounded (when *g* is bounded) solution to (7) that yields the lowest MPE payoff at X_0 , by analogous reasoning it is not possible to have $D_0 = W_0(S) - U(X_0) < 0$, implying $W_0 \ge U(X_0)$. The proof of Theorem 2 immediately follows from Lemma 6, the analogue for $W_0 \ge U(X_0)$, and the fact that at any state $X \in \mathcal{X}$, it is possible for the large player to achieve any payoff in the convex hull of the set of Markov equilibrium payoffs at state X by randomization at time zero.

Proof of Corollary 1 The existence of a Markov equilibrium follows from Theorem 1. When μ is independent of *a*, the sequential rationality condition (6) in a Markov equilibrium collapses to maximizing the static flow payoff, and the large player plays the unique static Nash action profile $S^*(X, 0)$ in each state. Therefore, any solution to (7) must satisfy

$$U(X_t) = E_t \bigg[r \int_t^\infty e^{-rs} g^*(X_s, 0) \, dt \bigg],$$
(20)

where the measure over the state is independent of the solution U since equilibrium actions are independent of U. Given that the right hand side of (20) is independent of U, (7) must have a unique solution and there is a unique Markov equilibrium. By Theorem 2, this is also the unique PPE. The solution to (7) evaluated at state X_t analytically characterizes (20).

A.4 Proof of Theorems 3 and 4

I prove Theorems 3 and 4 simultaneously. The proof proceeds in three steps:

Step 1. Any solution to the optimality equation has the same boundary conditions.

- Step 2. If all solutions have the same boundary conditions, then there is a unique linear growth (bounded) solution.
- Step 3. When there is a unique solution, then there is a unique PPE.

Let $\psi(X, z) \equiv g^*(X, z) + z\mu^*(X, z)/r$ be the value of the large player's incentive constraint at the sequentially rational action profile for incentive weight z/r. All intermediate theorems and lemmas maintain Assumptions 1 to 3 and, as stated, Assumption 4. As a reminder, $|\cdot|$ denotes the Euclidean norm for vectors. I first present an intermediate result that will be used in Steps 1 and 2.

LEMMA 7. Suppose U and V are both linear growth (bounded) solutions to (7), with U(X) < V(X) for some interior state $X \in \mathcal{X}$. Then V - U does not have an interior maximum and is monotone for large |X|.

PROOF. First suppose \mathcal{X} is compact. It follows from identical reasoning to Lemma C.7 in Faingold and Sannikov (2011) that if U and V are two linear growth (bounded) solutions of (7) such that $U(X_0) \leq V(X_0)$ and $U'(X_0) \leq V'(X_0)$, with at least one strict inequality, then U(X) < V(X) and U'(X) < V'(X) for all $X \in (X_0, \overline{X})$.³¹ Similarly if $U(X_0) \leq V(X_0)$ and $U'(X_0) \geq V'(X_0)$, with at least one strict inequality, then U(X) < V(X) and U'(X) > V'(X) for all $X \in (X, X_0)$.

Suppose *U* and *V* are both bounded solutions to (7), with U(X) < V(X) for some $X \in (\underline{X}, \overline{X})$. Suppose V - U has an interior maximum at some $X^* \in (\underline{X}, \overline{X})$. Then by continuity, this implies that $U'(X^*) = V'(X^*)$. If $U(X^*) < V(X^*)$, then by the above statement, U'(X) < V'(X) for all $X > X^*$ and, therefore, V(X) - U(X) is strictly increasing for $X > X^*$. This contradicts that X^* is an interior maximum. If $U(X^*) > V(X^*)$, then by the above statement, U(X) > V(X) and U'(X) > V(X) for all $X > X^*$, and U(X) > V(X) and U'(X) < V'(X) for all $X < X^*$. Therefore, X^* is a global maximum. This contradicts U(X) < V'(X) for some $X \in (\underline{X}, \overline{X})$. Therefore, V - U does not have an interior maximum. Given this, V - U has at most one interior minimum. Therefore, there exists a $\delta > 0$ such that V - U is monotone for $|\overline{X} - X| < \delta$ and $|\underline{X} - X| < \delta$. The proof for the case of $\mathcal{X} = \mathbb{R}$ is analogous, replacing \overline{X} and \underline{X} with ∞ and $-\infty$, respectively.

Step 1: Boundary conditions Lemmas 8 to 19 as well as Lemmas 26 and 27 in Supplemental Appendix D.2 establish the following boundary conditions for the case of $\mathcal{X} = \mathbb{R}$. When *g* is unbounded, any solution *U* of (7) with linear growth satisfies $\lim_{X\to p} U(X) - y^L(X) = g_2(z_p) + z_p\mu_2(z_p)/r$, $\lim_{X\to p} U'(X) = z_p$, and $\lim_{X\to p} \sigma(X, U'(X))^2 U''(X) = 0$ for $p \in \{-\infty, \infty\}$, where $z_p \equiv r\overline{g}_p/(r - \overline{\mu}_p)$ given $\overline{\mu}_p \equiv \lim_{X\to p} \mu^*(X, z)/X$ and $\overline{g}_p \equiv \lim_{X\to p} g^*(X, z)/X$, which exist and are finite, and $y^L(x) \equiv -f(x)\int rg_1(x)/f(x)\mu_1(x) dx$ with integrating factor $f(x) \equiv \exp(\int r/\mu_1(x) dx)$ when $\lim_{x\to p} \mu_1(x) \neq 0$ and $y^L(x) \equiv g_1(x)$ when $\lim_{X\to p} g^*(X, 0)$, and $\lim_{X\to p} U'(X) = 0$. Supplemental Appendix D.1 establishes analogous boundary conditions for the case of \mathcal{X} compact, and Supplemental Appendix D.3 establishes the same boundary conditions for the case of $\mathcal{X} = \mathbb{R}$ and *g* bounded under an alternative to Assumption 4.

Define $\overline{\psi}(X, z) \equiv \psi(X, z)/X$ and $\overline{U}(X) \equiv U(X)/X$. Let ψ' and $\overline{\psi'}$ denote the partial derivatives of ψ and $\overline{\psi}$ with respect to X. Let $\delta_0 > 0$ denote the lower bound above which the large |X| properties of Assumptions 3 and 4 hold. Several lemmas use the property that $g^*(X, z)$, $\mu^*(X, z)$, and $\sigma^*(X, z)$ are bounded in z, which follows from the compactness of A and B(X). The Lipschitz continuity of g_1 , μ_1 , g_2 , and μ_2 is also used, which follows from the Lipschitz continuity of $g^*(X, z)$ and $\mu^*(X, z)$. The following series of lemmas are stated for an unbounded flow payoff g; to apply them to a bounded

³¹Analogous to the definition of ϕ_1 in their result, set $X_1 \equiv \inf\{X \in [X_0, \overline{X}) : U'(X) \ge V'(X)\}$ and apply the same reasoning.

flow payoff, simply substitute "bounded solution to (7)" for "linear growth solution to (7)" throughout.

LEMMA 8. Suppose $\mathcal{X} = \mathbb{R}$. Given $p \in \{-\infty, \infty\}$, $\overline{\mu}_p \equiv \lim_{X \to p} \mu^*(X, z)/X$ and $\overline{g}_p \equiv \lim_{X \to p} g^*(X, z)/X$ exist and are finite. Moreover, $\lim_{X \to p} \overline{\psi}(X, z) = \lim_{X \to p} \psi'(X, z) = \psi_p(z)$ for all $z \in \mathbb{R}$, where $\psi_p(z) \equiv \overline{g}_p + z\overline{\mu}_p/r$.

PROOF. Let $p = \infty$ and fix $z \in \mathbb{R}$. Given Assumption 4(i), $\psi'(X, z) = g'_1(X) + z\mu'_1(X)/r$ for $X > \delta_0$. By the Lipschitz continuity of g_1 and μ_1 , g'_1 and μ'_1 are bounded, and, therefore, $\psi'(\cdot, z)$ is bounded for any $z \in \mathbb{R}$. By Assumption 3, $\psi'(\cdot, z)$ and g'_1 are monotone for large X (the latter follows from the assumption holding at z = 0). Therefore, by the monotone convergence theorem, $\psi_{\infty}(z) \equiv \lim_{X \to \infty} \psi'(X, z)$ and $\overline{g}_{\infty} \equiv \lim_{X \to \infty} g'_1(X)$ exist and are finite. Given that ψ' and g'_1 have well defined limits and $\psi'(X, z) = g'_1(X) + z\mu'_1(X)/r$ for large $X, \ \overline{\mu}_{\infty} \equiv \lim_{X \to \infty} \mu'_1(X)$ exists and is finite. Moreover, $\psi_{\infty}(z) = \overline{g}_{\infty} + z\overline{\mu}_{\infty}/r$. When g_1 and μ_1 are unbounded, then by l'Hopital's rule, $\lim_{X \to \infty} \overline{\psi}(X, z) = \psi_{\infty}(z)$, $\lim_{X \to \infty} g_1(X)/X = \overline{g}_{\infty}$, and $\lim_{X \to \infty} \mu_1(X)/X = \overline{\mu}_{\infty}$. In the case where g_1 or μ_1 is bounded, this immediately follows from $\overline{g}_{\infty} = 0$ or $\overline{\mu}_{\infty} = 0$. Given that $g_2(z)$ and $\mu_2(z)$ are independent of X, $\lim_{X \to \infty} g_2(z)/X = 0$ and $\lim_{X \to \infty} \mu_2(z)/X = 0$. This implies $\lim_{X \to \infty} g^*(X, z)/X = \overline{g}_{\infty}$ and $\lim_{X \to \infty} \mu^*(X, z)/X = \overline{\mu}_{\infty}$. Note that $\overline{\mu}_{\infty} < r$ by Assumption 2. The proof for $p = -\infty$ is analogous.

LEMMA 9. Suppose $\mathcal{X} = \mathbb{R}$ and U is a solution of (7) with linear growth. Then for $p \in \{-\infty, \infty\}$, there exists a finite $U'_p \in \mathbb{R}$ such that $\lim_{X \to p} \overline{U}(X) = \lim_{X \to p} U'(X) = U'_p$.

PROOF. Let $p = \infty$ and let U be a solution of (7) with linear growth. Suppose $\liminf_{X\to\infty} U'(X) \neq \limsup_{X\to\infty} U'(X)$. Then for all $\delta > 0$, by the continuity of U', there exists a z and an increasing sequence $(X_n)_{n\in\mathbb{N}}$ of alternating consecutive X such that $X_1 > \delta$, $U'(X_n) = z$, and $U''(X_n) \leq 0$ for n odd, and $U'(X_n) = z$ and $U''(X_n) \geq 0$ for n even, with one inequality for U'' strict. From (7), this implies $U(X_n) \leq \psi(X_n, z)$ for n odd and $\psi(X_n, z) \leq U(X_n)$ for n even. Thus, the oscillation of $\psi'(X, z)$ is at least as large as the oscillation of U'. But by Assumption 3, $\psi'(X, z)$ is monotone for $X > \delta_0$. Therefore, it must be that $\liminf_{X\to\infty} U'(X) = \limsup_{X\to\infty} U'(X)$. Let U'_{∞} denote this limit. Given U has linear growth, $|U'_{\infty}| < \infty$ and, by l'Hopital's rule, $\lim_{X\to\infty} \overline{U}(X) = U'_{\infty}$. In the case of g bounded, U bounded implies $U'_{\infty} = 0$ and $\lim_{X\to\infty} \overline{U}(X) = 0$. The proof for $p = -\infty$ is analogous.

LEMMA 10. Suppose $\mathcal{X} = \mathbb{R}$ and U is a solution of (7) with linear growth. Then $\lim_{X \to p} \overline{\psi}(X, U'(X)) = \psi_p(U'_p)$ for $p \in \{-\infty, \infty\}$, where $U'_p \equiv \lim_{X \to p} U'(X)$.

PROOF. Let $p \in \{-\infty, \infty\}$ and let *U* be a solution of (7) with linear growth. Given μ^* and g^* are Lipschitz continuous and additively separable in (X, z) for $|X| > \delta_0$, there exists a $M_1, M_2, M_3, c > 0$ and $\delta > \delta_0$ such that for $|X| > \delta$,

$$|\psi(X, z_1) - \psi(X, z_2)| \le M_1 |z_1 - z_2| + M_2 |z_1| |z_1 - z_2| + M_3 |z_1 - z_2| (|X| + |z_2|).$$

From Lemma 9, $U'_p \equiv \lim_{X \to p} U'(X)$ exists and is finite. Therefore,

$$\begin{split} &\lim_{X \to p} \left| \overline{\psi} \big(X, U'(X) \big) - \overline{\psi} \big(X, U'_p \big) \right| \\ &= \lim_{X \to p} \left| \psi \big(X, U'(X) \big) - \psi \big(X, U'_p \big) \big| / |X| \\ &\leq \lim_{X \to p} \big(M_1 \big| U'(X) - U'_p \big| + M_2 \big| U'(X) \big| \big| U'(X) - U'_p \big| \\ &+ M_3 \big| U'(X) - U'_p \big| \big(|X| + \big| U'_p \big| \big) \big) / |X| = 0. \end{split}$$

From Lemma 8, $\lim_{X \to p} \overline{\psi}(X, U'_p) = \psi_p(U'_p)$. Therefore, $\lim_{X \to p} \overline{\psi}(X, U'(X)) = \psi_p(U'_p)$.

LEMMA 11. Suppose $\mathcal{X} = \mathbb{R}$ and $f : \mathbb{R} \to \mathbb{R}$ has linear growth. Then any solution U of (7) with linear growth satisfies $\liminf_{X \to p} |f(X)| U''(X) \le 0 \le \limsup_{X \to p} |f(X)| U''(X)$ for $p \in \{-\infty, \infty\}$.

PROOF. Let $p = \infty$ and let U be a solution of (7) with linear growth. Suppose f has linear growth and $\liminf_{X\to\infty} |f(X)|U''(X) > 0$. There exists an $\delta_1, M > 0$ such that when $X > \delta_1, |f(X)| \le MX$. Given $\liminf_{X\to\infty} |f(X)|U''(X) > 0$, there exists a $\delta_2, \varepsilon > 0$ such that when $X > \delta_2, |f(X)|U''(X) > \varepsilon$. Take $\delta = \max\{\delta_1, \delta_2\}$. Then for $X > \delta$, $U''(X) > \varepsilon/|f(X)| \ge \varepsilon/MX$. The antiderivative of ε/MX is $(\varepsilon/M) \ln X$, which converges to ∞ as $X \to \infty$. Therefore, U' must grow unboundedly large as $X \to \infty$, which violates the linear growth of U. Therefore, $\liminf_{X\to\infty} |f(X)|U''(X) \le 0$. The proof is analogous for $\limsup_{X\to\infty} |f(X)|U''(X) \ge 0$ as well as the case of $p = -\infty$.

LEMMA 12. Suppose $\mathcal{X} = \mathbb{R}$. Any solution U of (7) with linear growth satisfies $U'_p = \psi_p(U'_p)$ and $\lim_{X \to p} U''(X) \sigma^*(X, U'(X))^2 / X = 0$ for $p \in \{-\infty, \infty\}$.

PROOF. Let $p = \infty$ and let *U* be a solution of (7) with linear growth. From the optimality equation,

$$\lim_{X \to \infty} \frac{|U''(X)| \sigma^* (X, U'(X))^2}{X} = \lim_{X \to \infty} \frac{2r |U(X) - g^* (X, U'(X)) - U'(X) \mu^* (X, U'(X)) / r|}{X}$$
$$= \lim_{X \to \infty} 2r |\overline{U}(X) - \overline{\psi} (X, U'(X))| = 2r |U'_{\infty} - \psi_{\infty} (U'_{\infty})|,$$

where the second equality follows from Lemmas 9 and 10. Let $c \equiv |U'_{\infty} - \psi_{\infty}(U'_{\infty})|$ and suppose c > 0. Then for any $\varepsilon > 0$, there exists a $\delta_1 > \delta_0$ such that for $X > \delta_1$, $|U''(X)|\sigma^*(X, U'(X))^2/X > c - \varepsilon$. This implies $|U''(X)| > (c - \varepsilon)X/\sigma^*(X, U'(X))^2$. Given that σ^* is Lipschitz continuous and bounded in z, there exists an M and a $\delta_2 > \delta_1$ such that for $X > \delta_2$, $\sigma^*(X, U'(X)) \le MX$. Therefore, for $X > \delta_2$, |U''(X)| > $(c - \varepsilon)X/M^2X^2 = (c - \varepsilon)/M^2X$. The antiderivative of $(c - \varepsilon)/M^2X$ is $((c - \varepsilon)/M^2) \ln X$, which converges to ∞ as $X \to \infty$. Therefore, U' must grow unboundedly large as $X \to \infty$, which violates the linear growth of U. Therefore, c = 0, which implies $U'_{\infty} =$ $\psi_{\infty}(U'_{\infty})$ and $\lim_{X\to\infty} U''(X)\sigma^*(X, U'(X))^2/X = 0$. The proof is analogous for the case of $p = -\infty$.

LEMMA 13. Suppose $\mathcal{X} = \mathbb{R}$ and that U is a solution of (7) with linear growth. Then for $p \in \{-\infty, \infty\}, U'_p = r\overline{g}_p/(r - \overline{\mu}_p)$.

PROOF. Let $p \in \{-\infty, \infty\}$ and let *U* be a solution of (7) with linear growth. From Lemma 12, $U'_p = \psi_p(U'_p)$, so U'_p is a fixed point of ψ_p . From Lemma 8, $\psi_p(z) = \overline{g}_p + z\overline{\mu}_p/r$. The unique fixed point is $z_p = r\overline{g}_p/(r - \overline{\mu}_p)$. Therefore, $U'_p = r\overline{g}_p/(r - \overline{\mu}_p)$.

LEMMA 14. Suppose $\mathcal{X} = \mathbb{R}$, and suppose U and V are solutions of (7) with linear growth and $\lim_{X \to p} \mu_1(X) \neq p$ for $p \in \{-\infty, \infty\}$. Then U = V.

PROOF. Let *U* and *V* be solutions to (7) with linear growth and suppose $\lim_{X\to p} \mu_1(X) \neq p$ for $p \in \{-\infty, \infty\}$. Suppose $U \neq V$. Without loss of generality suppose $U(X_k) < V(X_k)$ for some $X_k \in (-\infty, \infty)$. Define $D \equiv V - U$, with D' = V' - U' and D'' = V'' - U''. Given the continuity of D, $D(X_k) > 0$ implies that there exists an $\varepsilon_1 > 0$ such that $D(X_k) > \varepsilon_1$. By Lemma 7, there exists a $\delta_1 > \delta_0$ such that D is monotone for $|X| > \delta_1$ and D does not have an interior maximum. Choose δ_1 large enough such that either (a) $D(X) > \varepsilon_1$ and $D'(X) \ge 0$ for $X > \delta_1$ or (b) $D(X) > \varepsilon_1$ and $D'(X) \le 0$ for $X < -\delta_1$. By Assumption 3, $\sigma^*(X, z)$ is independent of z for $X > \delta_0$. In a slight abuse of notation, write $\sigma^*(X)$ to simplify notation throughout the proof. From (7) and Assumption 4(i), for $|X| > \delta_1$,

$$\frac{1}{2}\sigma^{*}(X)^{2}D''(X) = rD(X) - D'(X)\mu_{1}(X) - r(g_{2}(V'(X)) - g_{2}(U'(X))) - (V'(X)\mu_{2}(V'(X)) - U'(X)\mu_{2}(U'(X))).$$
(21)

First consider case (a) where $D(X) > \varepsilon_1$ and $D'(X) \ge 0$ for $X > \delta_1$. From the Lipschitz continuity of g_2 and μ_2 and Lemma 13, for any $\varepsilon_2 > 0$, there exists a $\delta_2 > \delta_1$ such that for $X > \delta_2$, $|rg_2(V'(X)) - rg_2(U'(X)) + V'(X)\mu_2(V'(X)) - U'(X)\mu_2(U'(X))| < \varepsilon_2$. By Assumption 4(i), μ_1 is monotone for $X > \delta_1$. Together with $\lim_{X\to\infty} \mu_1(X) \ne \infty$, this implies that either μ_1 is bounded for $X > \delta_1$ or $\lim_{X\to\infty} \mu_1(X) = -\infty$. If $\mu_1(X)$ is bounded for $X > \delta_1$, then $\lim_{X\to\infty} D'(X)\mu_1(X) = 0$, and if $\lim_{X\to\infty} \mu_1(X) = -\infty$, then $\mu_1(X) < 0$ for large X, and, therefore, $D'(X)\mu_1(X) \le 0$ for large X. Therefore, for any $\varepsilon_3 > 0$, there exists a $\delta_3 > \delta_1$ such that for $X > \delta_3$, $D'(X)\mu_1(X) < \varepsilon_3$. Choosing $\varepsilon_2 = \varepsilon_3 = r\varepsilon_1/4$, there exists a δ_4 such that for $X > \delta_4$,

$$\frac{1}{2}\sigma^*(X)^2 D''(X) > r\varepsilon_1 - r\varepsilon_1/4 - r\varepsilon_1/4 = r\varepsilon_1/2 > 0.$$
(22)

By Assumption 4(ii), $\sigma^*(X)^2$ is Lipschitz continuous and, therefore, has linear growth, and *D* has linear growth since *U* and *V* have linear growth. By similar reasoning to Lemma 11, it must be that $\liminf_{X\to p} \sigma^*(X)^2 D''(X) \leq 0$. This is a contradiction.³² Therefore, it cannot be that $U \neq V$. The reasoning for case (b) is analogous.

LEMMA 15. Suppose $\mathcal{X} = \mathbb{R}$ and that U is a solution of (7) with linear growth. For $p \in \{-\infty, \infty\}$, if $\lim_{X \to p} \sigma^*(X, z) = \infty$, suppose $|\mu_1(X)|/\sigma^*(X, z)^2$ is bounded away from zero near p. Then $\lim_{X \to p} \sigma^*(X, U'(X))^2 U''(X) = 0$.

³²When g is bounded, a contradiction is reached from $\sigma^*(X)$ Lipschitz by similar reasoning to Lemma 27 in Supplemental Appendix D.2.

PROOF. Let $p = \infty$ and let U be a solution of (7) with linear growth. By Assumption 3, $\sigma^*(X, z)$ is independent of z for $X > \delta_0$. In a slight abuse of notation, write $\sigma^*(X)$ to simplify notation throughout the proof. In the case where $\lim_{X\to\infty} \sigma^*(X) = \infty$, suppose $|\mu_1(X)|/\sigma^*(X)^2$ is bounded away from zero for sufficiently large X. Suppose $\lim_{X\to\infty} \sigma^*(X)^2 U''(X) \neq \limsup_{X\to\infty} \sigma^*(X)^2 U''(X)$. Then for all $\delta > \delta_0$, by the continuity of U'', there exists a $z \neq 0$ and an increasing sequence $(X_n)_{n\in\mathbb{N}}$ such that $X_1 > \delta$, $\sigma^*(X_n)^2 U''(X_n) = z$, and $2\sigma^*(X_n)\sigma^{*'}(X_n)U''(X_n) + \sigma^*(X_n)^2 U'''(X_n) < 0$ for n odd, and $2\sigma^*(X_n)\sigma^{*'}(X_n)U''(X_n) + \sigma^*(X_n)^2 U'''(X_n) < 0$ for n odd, and $2\sigma^*(X_n)\sigma^{*'}(X_n)U''(X_n) > 0$ for n even. From differentiating (7), for $X > \delta_0$,

$$\sigma^{*}(X)\sigma^{*'}(X)U''(X) + \frac{1}{2}\sigma^{*}(X)^{2}U'''(X)$$

= $r(U'(X) - \psi'(X, U'(X)))$
 $- U''(X)[\mu_{1}(X) + rg'_{2}(U'(X)) + \mu_{2}(U'(X)) + U'(X)\mu'_{2}(U'(X))]),$ (23)

where $\psi'(X, U'(X)) = rg'_1(X) + U'(X)\mu'_1(X)$. The right hand side of (23) is strictly negative at X_n for *n* odd and strictly positive at X_n for *n* even. Using $\sigma^*(X_n)^2 U''(X_n) = z$ and rearranging terms, this implies

$$r(U'(X_n) - \psi'(X_n, U'(X_n))/z - (rg'_2(U'(X_n)) + \mu_2(U'(X_n)) + U'(X_n)\mu'_2(U'(X_n)))/\sigma^*(X_n)^2 < \mu_1(X_n)/\sigma^*(X_n)^2$$
(24)

for *n* odd and the inequality reversed for *n* even when z > 0, with the opposite when z < 0. By Lemma 13, $U'(X) - \psi'(X, U'(X)) \rightarrow 0$ and U'(X) converges. By Assumption 1, $\sigma^*(X)^2$ is bounded away from zero, and by Assumption 3, $\sigma^*(X)^2$ is monotone. Therefore, $1/\sigma^*(X)^2$ converges to a finite limit. Together this implies that the left hand side of (24) converges to some finite K, with K = 0 when $\sigma^*(X)^2 \to \infty$. Therefore, for any $\varepsilon > 0$, there exists a $\delta_1 > \delta_0$ such that for $X > \delta_1$, $K - \varepsilon < \mu_1(X_n) / \sigma^*(X_n)^2$ for *n* odd and $K + \varepsilon > \mu_1(X_n) / \sigma^*(X_n)^2$ for *n* even. Note that μ_1 is monotone for $|X| > \delta_0$ by Assumption 4(i) and, therefore, $\mu_1(X)/\sigma^*(X)^2$ either converges to a finite limit or approaches infinity. In the case where $\sigma^*(X)$ is bounded as $X \to \infty$, this leads to a contradiction provided $\mu_1(X)$ does not converge to $rg'_2(U'_{\infty}) + \mu_2(U'_{\infty}) + U'_{\infty}\mu'_2(U'_{\infty})$. In particular, this is a contradiction when $|\mu_1(X)| \to \infty$. In the case where $\lim_{X\to\infty} \sigma^*(X) = \infty$, K = 0. This leads to a contradiction given $|\mu_1(X)|/\sigma^*(X)^2$ is bounded away from zero. Therefore, it must be that $\liminf_{X\to\infty} \sigma^*(X)^2 U''(X) = \limsup_{X\to\infty} \sigma^*(X)^2 U''(X)$. By Assumption 4(ii), $\sigma^*(X)^2$ is Lipschitz continuous and, therefore, has linear growth. Given that $\lim_{X\to\infty} \sigma^*(X)^2 U''(X)$ exists, by Lemma 11, it must be that $\lim_{X\to\infty} \sigma^*(X)^2 U''(X) =$ 0.³³ The proof for $p = -\infty$ is analogous.

LEMMA 16. Suppose $\lim_{x\to p} \mu_1(x) \neq 0$ for $p \in \{-\infty, \infty\}$ and y is a solution to the ordinary differential equation (ODE)

$$y'(x) - (r/\mu_1(x))y(x) = 0$$
(25)

on \mathbb{R} with linear growth. Then $\lim_{x\to p} y(x) = 0$ for $p \in \{-\infty, \infty\}$.

³³When g is bounded, this follows from $\sigma^*(X)$ Lipschitz by Lemma 27 in Supplemental Appendix B.2.

PROOF. The general solution to (25) is $y(x) = c \exp(\int r/\mu_1(x) dx)$, where $c \in \mathbb{R}$ is a constant. Trivially, there always exists a bounded (and, therefore, linear growth) solution because y(x) = 0 is a solution. Consider $p = \infty$. By Assumption 2, μ_1 has linear growth with rate slower than r. Given that $\lim_{x\to\infty} \mu_1(x) \neq 0$ and μ_1 is monotone for large x by Assumption 4(i), there exists a $\delta > 0$ such that for $x > \delta$, either (i) there exists a $k \in (0, r)$ such that $\mu_1(x) \in (0, kx]$ or (ii) there exists a k > 0 such that $\mu_1(x) \in [-kx, 0)$. First suppose there exists a $k \in (0, r)$ and $\delta > 0$ such that for $x > \delta$, $\mu_1(x) \in (0, kx]$. Then $1/\mu_1(x) \ge 1/kx$. But $\exp(\int r/kx \, dx) = \exp((r/k) \ln x) = x^{r/k}$ is not in O(x) since r/k > 1. Therefore, $\exp(\int r/\mu_1(x) dx)$ is not in O(x). Therefore, any solution that has linear growth must have c = 0. The unique solution with linear growth is y(x) = 0, which trivially satisfies $\lim_{x\to\infty} y(x) = 0$. Next suppose there exists a k, $\delta > 0$ such that for $x > \delta$, $\mu_1(x) \in [-kx, 0)$. Then $1/\mu_1(x) \le -1/kx$. But $\exp(\int -r/kx \, dx) = \exp(-r \ln x/k) = x^{-r/k}$ and $\lim_{x\to\infty} x^{-r/k} \to 0$. Therefore, $\lim_{x\to\infty} \exp(\int r/\mu_1(x) dx) = 0$. Therefore, for all c, $\lim_{x\to\infty} y(x) = 0$ and any solution satisfies this property. The case for $p = -\infty$ is analogous.

LEMMA 17. Suppose $\mathcal{X} = \mathbb{R}$, and that U and V are solutions of (7) with linear growth. Then $\lim_{X \to p} V(X) - U(X) = 0$ for $p \in \{-\infty, \infty\}$.

PROOF. Let *U* and *V* be solutions of (7) with linear growth. Lemma 14 established that when $\lim_{X\to p} \mu_1(X) \neq p$ for $p \in \{-\infty, \infty\}$, then U = V, which trivially implies the result. Therefore, consider the case where either $\lim_{X\to\infty} \mu_1(X) = \infty$ or $\lim_{X\to-\infty} \mu_1(X) = -\infty$. First suppose $\lim_{X\to\infty} \mu_1(X) = \infty$. By Assumption 4(iii), $|\mu_1(X)|/\sigma^*(X, z)^2$ is bounded away from zero near ∞ . Define D = V - U. Then D' = V' - U', *D* has linear growth since *U* and *V* have linear growth, and $\lim_{X\to\infty} D'(X) = 0$ by Lemma 13. By Lemma 15, $\lim_{X\to\infty} \sigma^*(X, U'(X))^2 D''(X) = 0$. Combined with (7) and the Lipschitz continuity of g_2 and μ_2 , this implies $\lim_{X\to\infty} D(X) - \mu_1(X)D'(X)/r = 0$. In the case where $\lim_{X\to\infty} \mu_1(X) = 0$, then $\lim_{X\to\infty} D(X) = 0$ follows from $\lim_{X\to\infty} \mu_1(X)D'(X)/r = 0$ for any solution *y* with linear growth. Therefore, $\lim_{X\to\infty} D(X) = 0$, which implies $\lim_{X\to\infty} V(X) - U(X) = 0$. The proof establishing $\lim_{X\to-\infty} V(X) - U(X) = 0$ when $\lim_{X\to-\infty} \mu_1(X) = -\infty$ is analogous.

When $\lim_{X\to p} \mu_1(X) = p$ for both $p \in \{-\infty, \infty\}$, this yields the result. When this only holds for $p = \infty$, then $\lim_{X\to\infty} D(X) = 0$ implies that only case (b) is relevant in Lemma 14 since D(X) cannot be bounded above ε_1 for large X, and similarly, only case (a) is relevant when $\lim_{X\to p} \mu_1(X) = p$ only holds for $p = -\infty$. Therefore, by identical reasoning to the proof of Lemma 14, U = V. This trivially establishes the result.

LEMMA 18. Suppose $\lim_{x\to p} \mu_1(x) \neq 0$ for $p \in \{-\infty, \infty\}$ and y is a solution to the ODE

$$y(x) - g_1(x) - \mu_1(x)y'(x)/r = 0$$
(26)

Dynamic moral hazard game 493

on \mathbb{R} with linear growth. Then for $p \in \{-\infty, \infty\}$, $\lim_{x \to p} y(x) - y^L(x) = 0$, where

$$y^{L}(x) \equiv -\phi(x) \int \left(\frac{1}{\phi(x)}\right) \frac{rg_{1}(x)}{\mu_{1}(x)} dx$$
(27)

and $\phi(x) \equiv \exp(\int r/\mu_1(x) dx)$.

PROOF. The general solution to (26) is $y(x) = -\phi(x) \int (rg_1(x)/\phi(x)\mu_1(x)) dx - c\phi(x))$, where ϕ is as defined above and $c \in \mathbb{R}$ is a constant. Consider $p = \infty$. By Lemma 16, $\lim_{x\to\infty} c\phi(x) = 0$ for any solution with linear growth. Therefore, $\lim_{x\to\infty} y(x) - y^L(x) = 0$ for any solution *y*.

LEMMA 19. Suppose $\mathcal{X} = \mathbb{R}$ and that U is a solution of (7) with linear growth, and for $p \in \{-\infty, \infty\}$, if $\lim_{X \to p} \sigma^*(X, z) = \infty$, suppose $|\mu_1(X)|/\sigma^*(X, z)^2$ is bounded away from zero near p. Then for $p \in \{-\infty, \infty\}$, when g is unbounded, $\lim_{X \to p} U(X) - y^L(X) = g_2(z_p) + z_p \mu_2(z_p)/r$, where y_L is defined by (27) when $\lim_{x \to p} \mu_1(x) \neq 0$ and $y^L(x) \equiv g_1(x)$ when $\lim_{x \to p} \mu_1(x) = 0$, while when g is bounded, $\lim_{X \to p} U(X) = g_p$, where $g_\infty \equiv \lim_{X \to \infty} g^*(X, 0)$.

PROOF. Let $p \in \{-\infty, \infty\}$ and let *U* be a solution of (7) with linear growth. Then from Assumption 4, Lemmas 13 and 15, and the Lipschitz continuity of g_2 and μ_2 ,

$$\lim_{X \to p} U(X) - g_1(X) - U'(X)\mu_1(X)/r = g_2(z_p) + z_p\mu_2(z_p)/r.$$

Therefore, when $\lim_{x\to p} \mu_1(x) \neq 0$, there exists a solution *y* to (26) with linear growth such that $\lim_{X\to p} U(X) - y(X) = g_2(z_p) + z_p\mu_2(z_p)/r$. By Lemma 18, $\lim_{X\to p} y(X) - y^L(X) = 0$. Therefore, $\lim_{X\to p} U(X) - y^L(X) = g_2(z_p) + z_p\mu_2(z_p)/r$, which establishes the boundary condition for *U*. When $\lim_{x\to p} \mu_1(x) = 0$, by Lemma 13, $\lim_{X\to p} U'(X) \times \mu_1(X)/r = 0$. It immediately follows that $\lim_{X\to p} U(X) - g_1(X) = g_2(z_p) + z_p\mu_2(z_p)/r$.

The boundary conditions for *g* bounded use two lemmas from Supplemental Appendix D. Consider $p = \infty$. In the case where *g* is bounded, $g_{\infty} \equiv \lim_{X\to\infty} g^*(X, 0)$ exists and is finite given that g^* is monotone for large |X|. Moreover, by Lemma 26, $U_{\infty} \equiv \lim_{X\to\infty} U(X)$ exists and is finite, and $\lim_{X\to\infty} U'(X) = 0$ (the latter also follows from Lemma 13). Then from (7),

$$\lim_{X \to \infty} \sigma^* (X, U'(X))^2 U''(X) = \lim_{X \to \infty} 2r (U(X) - g^* (X, U'(X))) - 2\mu^* (X, U'(X)) U'(X)$$

$$\Rightarrow \quad 0 = 2r (U_\infty - g_\infty) - \lim_{X \to \infty} 2\mu^* (X, U'(X)) U'(X), \tag{28}$$

where the left hand side of the second line follows from Lemma 15, and $\lim_{X\to\infty} g^*(X, U'(X)) = g_{\infty}$ follows from Lipschitz continuity and $\lim_{X\to\infty} U'(X) = 0$. Therefore, $\lim_{X\to\infty} \mu^*(X, U'(X))U'(X) = r(U_{\infty} - g_{\infty})$. By Lemma 27, this implies $U_{\infty} = g_{\infty}$. The case of $p = -\infty$ is analogous.

Step 2: Uniqueness of solution to optimality equation Suppose *U* and *V* are both linear growth (bounded) solutions to (7) and $U \neq V$, where without loss of generality, U(X) < V(X) for some interior *X*. Lemma 7 establishes that V - U does not have an interior maximum. But $\lim_{X\to p} V(X) - U(X) = 0$ for $p \in \{\underline{X}, \overline{X}\}$ in the case of compact \mathcal{X} and $p \in \{-\infty, \infty\}$ in the case of $\mathcal{X} = \mathbb{R}$ (by Lemmas 17, 25 and 29). Therefore, it cannot be that U(X) < V(X) for some interior *X*, as by continuity this would require the existence of an interior maximum in order to satisfy the boundary conditions. Therefore, *U* and *V* cannot differ, and there exists a unique linear growth (bounded) solution to (7).

Step 3: Uniqueness of PPE By Step 2, there is a unique linear growth (bounded) solution to (7). It remains to show that there are no other PPE. When there is a unique solution to (7), Theorem 2 implies that in any PPE with continuation values $(W_t)_{t\geq 0}$, $W_t = U(X_t)$ for all *t*. Therefore, the volatilities of the two continuation values are equal; otherwise, they both cannot be equal to $U(X_t)$. Given equal volatilities, actions are uniquely specified by $S^*(X, U'(X)/r)$. Therefore, there exists a unique PPE.

A.5 Proofs for Section 5

PROOF OF PROPOSITION 1. Let *U* be a bounded or linear growth solution to (7) and let $I \equiv [I_1, I_2] \subset \mathcal{X}$ denote a closed proper interval of states. At a state *X* corresponding to an interior extremum on *I*, U'(X) = 0. From (7), if *X* is an interior minimum on *I*, $g^*(X, 0) \leq U(X)$, and if *X* is an interior maximum on *I*, $U(X) \leq g^*(X, 0)$. Let *n* denote the number of (strict interior) interval extrema of *U* on *I* and let X_i denote the interval of states corresponding to the *i*th such extremum for i = 1, ..., n, where $I_1 < X_1 < X_2 < \cdots < X_n < I_2$ and $X_i < X_j$ corresponds to $\sup X_i < \inf X_j$. By the continuity of *U*, if n > 1 and X_i is a minimum for some i < n, then X_{i+1} must be a maximum and vice versa.

Item (i). Suppose $g^*(\cdot, 0)$ is constant on I and $n \ge 2$, so there are at least two interval extrema. If X_1 is a minimum, then X_2 must be a maximum with $U(x_1) < U(x_2)$ for $x_1 \in X_1$ and $x_2 \in X_2$. Therefore, $g^*(x_1, 0) \le U(x_1) < U(x_2) \le g^*(x_2, 0)$ for any $x_1 \in X_1$ and $x_2 \in X_2$. This is a contradiction because $g^*(\cdot, 0)$ is constant on I. The same logic holds for X_1 a maximum. Therefore, $n \le 1$.

Item (ii). Suppose $g^*(\cdot, 0)$ is strictly increasing on *I*. First show that a maximum cannot be followed by a minimum. Suppose $n \ge 2$, X_i is a maximum, and X_{i+1} is a minimum for some i < n. Then for $x_i \in X_i$ and $x_{i+1} \in X_{i+1}$, $U(x_{i+1}) < U(x_i)$, and, therefore, $g^*(x_{i+1}, 0) \le U(x_{i+1}) < U(x_i) \le g^*(x_i, 0)$. This is a contradiction because $g^*(\cdot, 0)$ is strictly increasing on *I*. Therefore, it is not possible to have a maximum followed by a minimum. Therefore, there can be at most two interval extrema, $n \le 2$. Further, if n = 2, X_1 is a minimum and X_2 is a maximum. The proof for the case of $g^*(\cdot, 0)$ strictly decreasing is analogous, where X_1 is a maximum and X_2 is a minimum when n = 2.

Suppose $g^*(\cdot, 0)$ is strictly increasing on *I* and that *U* is constant on *I*. Then U'(X) = 0 and U''(X) = 0 for all $X \in I$. From (7), $U(X) = g^*(X, 0)$ for all $X \in I$. Therefore, $g^*(X, 0)$ is constant on *I*, a contradiction. This establishes that *U* is not constant on *I*.

Item (iii). If X_i is a maximum (minimum) for some i < n, then X_{i+1} is a minimum (maximum) with $U(x_{i+1}) < U(x_i)$ ($U(x_i) < U(x_{i+1})$) for $x_i \in X_i$ and $x_{i+1} \in X_{i+1}$. This then follows directly from $U(X) \ge g^*(X, 0)$ at any interval minimum, $U(X) \le g^*(X, 0)$ at any interval maximum, and the Lipschitz continuity of g^* .

PROOF OF PROPOSITION 2. Let *U* be the unique bounded solution to (7). At a state *X* corresponding to an interior extremum on \mathcal{X} , U'(X) = 0. From (7), if *X* is an interior minimum, $g^*(X, 0) \leq U(X)$, and if *X* is an interior maximum, $U(X) \leq g^*(X, 0)$. Let *n* denote the number of (strict interior) interval extrema of *U* on \mathcal{X} and let n_g denote the number of (strict interior) interval extrema of $g^*(X, 0)$ on \mathcal{X} . First consider \mathcal{X} compact.

Item (i). Suppose $g^*(\cdot, 0)$ is constant on \mathcal{X} . Then $n_g = 0$ and there exists a $c \in \mathbb{R}$ such that $g^*(X, 0) = c$ for all $X \in \mathcal{X}$. By part (iv) (see below for proof), $n_g = 0$ implies n = 0. By the boundary conditions from Theorem 4, $U(\underline{X}) = c$ and $U(\overline{X}) = c$, which implies $U(\underline{X}) = U(\overline{X})$. Combined with n = 0, this implies that U is constant on \mathcal{X} . To establish the other direction, suppose $g^*(\cdot, 0)$ is not constant on \mathcal{X} . Then there exists a proper interval $I_1 \subset \mathcal{X}$ such that $g^*(\cdot, 0)$ is strictly monotone on I_1 . Take a closed proper subset $I_2 \subset I_1$. By Proposition 1(ii), U is not constant on I_2 . Therefore, U is not constant on \mathcal{X} .

Item (ii). Suppose $g^*(\cdot, 0)$ is monotonically increasing on \mathcal{X} and U is not monotonically increasing. Then $n_g = 0$ and U'(X) < 0 for some $X \in \mathcal{X}$. By Proposition 2(iv) (see below for proof), $n_g = 0$ implies n = 0. Therefore, it must be that U is monotonically decreasing on \mathcal{X} , i.e., $U'(X) \leq 0$ for all $X \in \mathcal{X}$. Given U'(X) < 0 for some $X \in \mathcal{X}$, this implies that $U(\overline{X}) < U(\underline{X})$. By the boundary conditions from Theorem 4, $U(\underline{X}) = g^*(\underline{X}, 0)$ and $U(\overline{X}) = g^*(\overline{X}, 0)$, and by the monotonicity of $g^*(\cdot, 0), g^*(\underline{X}, 0) \leq g^*(\overline{X}, 0)$. This implies $U(\underline{X}) \leq U(\overline{X})$, a contradiction. Therefore, U is monotonically increasing. The proof for U monotonically decreasing is analogous.

Item (iii). Suppose $g^*(\underline{X}, 0) = g^*(\overline{X}, 0)$ and suppose $g^*(\cdot, 0)$ is single-peaked with a unique interval extremum, a maximum. Let X_1^g denote the interval of states corresponding to this extremum. By Proposition 2(iv), $n_g = 1$ implies $n \le 1$. Given that there is a unique interval extremum and it is a maximum, $g^*(\cdot, 0)$ is monotonically increasing on $I_1 = [\underline{X}, \inf X_1^g]$ and strictly so on some proper interval $I'_1 \subset I_1$. Therefore, by Proposition 1(ii), U is not constant on I_1 . Similarly, $g^*(\cdot, 0)$ is monotonically decreasing on $I_2 = [\sup X_1^g, \overline{X}]$ and strictly so on some proper interval $I'_2 \subset I_2$, so U is not constant on I_2 . From the boundary conditions, $U(\underline{X}) = g^*(\underline{X}, 0)$ and $U(\overline{X}) = g^*(\overline{X}, 0)$. Therefore, $U(\underline{X}) = U(\overline{X})$. Since U is not constant and $U(\underline{X}) = U(\overline{X})$, by continuity U must have at least one interval extremum, $n \ge 1$. Given that it was already established that $n \le 1$, it must be that n = 1.

Suppose the unique interval extremum for U is a minimum. Let X_1 denote the interval of states corresponding to this extremum. Then $g^*(x_1, 0) \le U(x_1)$ for all $x_1 \in X_1$. Given that X_1 is a minimum and it is the unique interval extremum, U is monotonically decreasing on $[\underline{X}, \inf X_1]$ and strictly so on some some proper interval $I \subset [\underline{X}, \inf X_1]$. This implies that for all $x_1 \in X_1$, $U(x_1) < U(\underline{X}) = g^*(\underline{X}, 0)$. Therefore, for all $x_1 \in X_1$, $g^*(x_1, 0) \le U(x_1) < g^*(\underline{X}, 0)$. Further, since X_1^g is the unique interval extremum of $g^*(\cdot, 0)$ and a maximum, for all $x_1^g \in X_1^g$, $g^*(\underline{X}, 0) = g^*(\overline{X}, 0) < g^*(x_1^g, 0)$. But then $g^*(\cdot, 0)$ must have two interval extrema, since $g^*(x_1, 0) < g^*(\underline{X}, 0) = g^*(\overline{X}, 0)$ for $x_1 \in X_1$ and g^* is continuous. This is a contradiction. Therefore, U is single-peaked with a unique interval maximum. The proof for U single-peaked with a minimum is analogous.

Item (iv). This follows directly from $U(X) \ge g^*(X, 0)$ at an interior minimum of U, $U(X) \le g^*(X, 0)$ at an interior maximum of U, $U(\underline{X}) = g^*(\underline{X}, 0)$, $U(\overline{X}) = g^*(\overline{X}, 0)$, and the Lipschitz continuity of g^* .

For the case of $\mathcal{X} = \mathbb{R}$, replace $g^*(\overline{X}, 0)$ and $U(\overline{X})$ with $\lim_{X \to \infty} g^*(X, 0)$ and $\lim_{X \to \infty} U(X)$, and analogously for *X*. These limits exist by the proof of Theorem 4. \Box

PROOF OF PROPOSITION 3. Let U be the unique bounded solution to (7). Then U is continuous and bounded on a closed set. Therefore, either U attains a global maximum on \mathcal{X} , in which case $\overline{W} = U(X_H)$ for some $X_H \in \mathcal{X}$, or in the case where \mathcal{X} is unbounded, it is also possible that $\overline{W} = \limsup_{X \to X_H} U(X)$ for either $X_H = -\infty$ or $X_H = \infty$. Suppose U attains a global maximum at an interior state $X_H \in \operatorname{int} \mathcal{X}$. Then $U'(X_H) = 0$ and $U''(X_H) \leq 0$. From (7), this implies

$$U''(X_H) = \frac{2r(\overline{W} - g^*(X_H, 0))}{\sigma^*(X_H, 0)^2} \le 0,$$

and, therefore, $\overline{W} \leq g^*(X_H, 0)$. If \mathcal{X} is bounded and U attains a global maximum at boundary state $X_H = \underline{X}$ or $X_H = \overline{X}$, then by Theorem 4, $\overline{W} = g^*(X_H, 0)$. Similarly, if \mathcal{X} is unbounded and $\lim_{X \to X_H} U(X) = \overline{W}$ for either $X_H = -\infty$ or $X_H = \infty$, then by Theorem 4, $\overline{W} = g^*(X_H, 0)$. Therefore, $\overline{W} \leq \inf_{X_H \in \mathcal{X}_H} g^*(X_H, 0)$. The proof for \underline{W} is analogous.

References

Abreu, Dilip, Paul Milgrom, and David Pearce (1991), "Information and timing in repeated partnerships." *Econometrica*, 59, 1713–1733. [450, 452, 455]

Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. [458]

Board, Simon and Moritz Meyer-ter-Vehn (2013), "Reputation for quality." *Econometrica*, 81, 2381–2462. [453]

Bohren, J. Aislinn (2016), "Using persistence to generate incentives in a dynamic moral hazard problem." PIER Working Paper 16-024. [462]

Cisternas, Gonzalo (2018), "Two-sided learning and the ratchet principle." *The Review* of *Economic Studies*, 85, 307–351. [453]

Cripps, Martin, George Mailath, and Larry Samuelson (2004), "Imperfect monitoring and impermanent reputations." *Econometrica*, 72, 407–432. [452, 455, 469]

De Coster, Colette and Patrick Habets (2006), *Two-Point Boundary Value Problems: Lower and Upper Solutions*. Elsivier. [483]

Dilmé, Francesc (2019), "Reputation building through costly adjustment." *Journal of Economic Theory*, 181, 586–626. [452]

Doraszelski, Ulrich and Mark Satterthwaite (2010), "Computable Markov-perfect industry dynamics." *RAND Journal of Economics*, 41, 215–243. [453]

Duffie, Darrell, Andreu Mas-Colell John Geanakoplos, and Andrew McLennan (1994), "Stationary Markov equilibria." *Econometrica*, 62, 745–781. [454]

Dutta, Prajit K. (1995), "A folk theorem for stochastic games." *Journal of Economic Theory*, 66, 1–32. [453]

Dutta, Prajit K. and Rangarajan Sundaram (1992), "Markovian equilibrium in a class of stochastic games: Existence theorems for discounted and undiscounted models." *Economic Theory*, 2, 197–214. [454]

Ekmekci, Mehmet (2011), "Sustainable reputations with rating systems." *Journal of Economic Theory*, 146, 479–503. [452]

Ericson, Richard and Ariel Pakes (1995), "Markov-perfect industry dynamics: A framework for empirical work." *The Review of Economic Studies*, 62, 53–82. [453]

Faingold, Eduardo and Yuliy Sannikov (2011), "Reputation in continuous time games." *Econometrica*, 79, 773–876. [450, 452, 453, 454, 455, 458, 460, 461, 462, 464, 465, 466, 468, 469, 481, 482, 483, 487]

Faingold, Eduardo (2020), "Reputation and the Flow of Information in Repeated Games." *Econometrica*, 88, 1697–1723. [452]

Fudenberg, Drew and David Levine (1989), "Reputation and equilibrium selection in games with a patient player." *Econometrica*, 57, 759–778. [452]

Fudenberg, Drew and David Levine (1992), "Maintaining a reputation when strategies are imperfectly observed." *Review of Economic Studies*, 59, 561–579. [452]

Fudenberg, Drew and David Levine (2007), "Continuous time limits of repeated games with imperfect public monitoring." *Review of Economic Dynamics*, 10, 173–192. [452, 453, 454, 462, 466]

Fudenberg, Drew and David Levine (2009), "Repeated games with frequent signals." Quarterly Journal of Economics, 233–265. [452, 454, 462, 466]

Fudenberg, Drew, David Levine, and Eric Maskin (1994), "The folk theorem with imperfect public information." *Econometrica*, 62, 997–1039. [453]

Fudenberg, Drew and Yuichi Yamamoto (2011), "The folk theorem for irreducible stochastic games with imperfect public monitoring." *Journal of Economic Theory*, 146, 1664–1683. [453]

Hörner, Johannes, Takuo Sugaya, Satoru Takahashi, and Nicolas Vieille (2011), "Recursive methods in discounted stochastic games: An algorithm and a folk theorem." *Econometrica*, 79, 1277–1318. [453]

Karatzas, Ioannis and Steven Shreve (1991), *Brownian Motion and Stochastic Calculus*. Springer-Verlag, New York. [481]

Kreps, David, Paul Milgrom, John Roberts, and Robert Wilson (1982), "Rational cooperation in the finitely dilemma repeated Prisoners' dilemma." *Journal of Economic Theory*, 27, 245–252. [452]

Kreps, David and Robert Wilson (1982), "Reputation and imperfect information." *Journal of Economic Theory*, 27, 253–279. [452]

Mailath, George J. and Larry Samuelson (2001), "Who wants a good reputation?" *Review of Economic Studies*, 68, 415–441. [452, 473]

Milgrom, Paul and John Roberts (1982), "Predation, reputation and entry deterrence." *Journal of Economic Theory*, 27, 280–312. [452]

Nowak, Andrzej S. and T. E. S. Raghavan (1992), "Existence of stationary correlated equilibria with symmetric information for discounted stochastic games." *Math. Oper. Res.*, 17, 519–526. [454]

Sannikov, Yuliy (2007), "Games with imperfectly observable actions in continuous time." *Econometrica*, 75, 1285–1329. [458]

Sannikov, Yuliy and Andrzej Skrzypacz (2007), "Impossibility of collusion under imperfect monitoring with flexible production." *American Economic Review*, 97, 1794–1823. [450, 452, 455, 462]

Sannikov, Yuliy and Andrzej Skrzypacz (2010), "The role of information in repeated games with frequent actions." *Econometrica*, 78, 847–882. [450, 462]

Shapley, Lloyd S. (1953), "Stochastic games." *Proceedings of the National Academy of Sciences*, 39, 1095. [454]

Strulovici, Bruno and Martin Szydlowski (2015), "On the smoothness of value functions and the existence of optimal strategies." *Journal of Economic Theory*, 159, 1016–1055. [458]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 2 November, 2016; final version accepted 29 October, 2019; available online 3 March, 2023.