Continuous-Time Stochastic Games with Imperfect Public Monitoring

Benjamin Bernard*

This paper characterizes the set of perfect public equilibrium payoffs, the set of stationary Markov perfect equilibrium payoffs, and a simple class of semi-stationary equilibrium payoffs in continuous-time stochastic games with finitely many states and a publicly observable Brownian signal about past actions. Contrary to many discrete-time methods, the characterization does not rely on a convergence to the stationary distribution of the underlying state process. As a consequence, the correspondence of initial state to equilibrium payoffs is preserved, the characterization is possible for any level of discounting, and the characterization is applicable to games, in which the state process is not irreducible.

Keywords: *Stochastic games, continuous time, imperfect monitoring, perfect public equilibrium, stationary Markov equilibrium, computation of equilibria.*

1 Introduction

A wide range of economic applications can be cast as strategic interactions between different parties, where the environment changes in response to the parties' behavior. In the global financial crisis, financial institutions were unable to raise sufficient capital to meet their short-term liabilities because investors had lost confidence in the financial sector. The past performance of the financial sector had led credit lines to dry up and created an environment where financial institutions were not able to bridge short-term liquidity gaps in the usual way. The development of new technologies by competing research institutions exhibits a similar history-dependent environment. The successful discovery of a new technology changes the research environment forever: competing researchers will not be able to patent similar work anymore and any effort put into such a discovery was exerted in vain. It is impossible to forecast the exact time of a financial crisis or the discovery of a new technology. The occurrence of such a state change is random and the likelihood depends on the involved parties' actions. A game-theoretic model that accounts for these

*Department of Economics, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd, Da'an District, Taipei City, 10617 Taiwan; `benbernard@ntu.edu.tw`. Financial support from the Ministry of Science and Technology in Taiwan through grant 108-2410-H-002-249 and the Center for Research in Econometric Theory and Applications by the Ministry of Education in Taiwan through grant 109L900203 is gratefully acknowledged.

sudden state changes is a stochastic game. No deterministic-time dynamic game can capture these sudden and potentially drastic changes in the environment.[1]

Our understanding of perfect public equilibria (PPE) in repeated games owes an invaluable debt to the recursive techniques of Abreu, Pearce, and Stacchetti [1]. Their notion of self-generating payoff sets has lead to numerous applications, allowing us to characterize equilibria in a variety of settings. Chief among them, Fudenberg and Levine [10] developed a linear program that characterizes the limit PPE payoff set as the players' discount factor δ tends to 1. A key ingredient in this linear program is the fact that continuation payoffs of extremal PPE payoffs have to lie below the tangent hyperplane. In a stochastic game, this is no longer the case because the stage game played in the next state need not be the same. Nevertheless, Hörner, Sugaya, Takahashi, and Vieille [14] show that a linear program can be recovered as $\delta \to 1$ if the limiting PPE payoff set is independent of the initial state. This is the case, for example, if the underlying state process is irreducible.[2]

In continuous-time stochastic games with finitely many states, the effects of state changes on continuation payoffs can be cleanly separated from the effects of the public signal: continuous movements in the continuation value due to information from the public signal have to lie below the tangent hyperplane, whereas discontinuous movements at a time of a state transition are entirely caused by the state transition. This allows us to use the techniques developed in Sannikov [29] to relate the continuous movement of the continuation value to the boundary of the set $\mathcal{E}_y(r)$ of PPE payoff for any initial state y and any discount rate r. Sannikov [29] shows that in a repeated game with Brownian information, the boundary of the the PPE payoff set can be characterized by an ordinary differential equation (ODE) describing the curvature of the set. In this paper, we show that in a stochastic game, the family $\left(\mathcal{E}_y(r)\right)_y$ of PPE payoff sets can be described via a set of coupled ODEs, in which the differential equation describing the payoff set $\mathcal{E}_y(r)$ depends on the solutions to the equations describing $\mathcal{E}_{y'}(r)$ for all states y' that can be reached from state y. In a stochastic game without possible cycles among states, this readily yields a procedure to solve for the equilibrium payoff set. In general, however, solving this system of coupled ODEs is a non-trivial fixed-point problem.

We show that this set of coupled ODEs can be approximated by a sequence of uncoupled ODEs that arise from an iterative procedure over state transitions. Consider an auxiliary game in each state y that ends at the time of the first state transition such

[1] Stochastic games have been used also to model dynamic competition with inventories (e.g., Kirman and Sobel [21], the extraction of a common resource (e.g., Levhari and Mirman [22]), economic growth (e.g., Bernheim and Ray [5]), entry and exit dynamics in oligopolistic competition (e.g., Ericson and Pakes [8]), strategic pricing (e.g., Bergemann and Välimäki [3]), social uprisings (e.g., Acemoglu and Robinson [2]), and strategic experimentation (e.g., Keller, Rady, and Cripps [20]).

[2] In his proof of the folk theorem, Dutta [7] provides a weaker condition for the limiting PPE payoff set to be independent of the state.

that terminal payoffs after a transition to state y' come from a fixed payoff set $\mathcal{W}_{y'}$. Because the state is fixed for the duration of the auxiliary game, the auxiliary game is a repeated game with discontinuous and abrupt information arrival at the game's end. The equilibrium payoff set $\mathcal{B}_{r,y}(\mathcal{W})$ of the auxiliary game is thus characterized by the techniques in Bernard [4]. It is an extension of Sannikov's result to monitoring structures that contain both continuous and discontinuous signals. We can recover equilibrium payoffs of the stochastic game from equilibrium payoffs of the auxiliary games if the family $\mathcal{B}_r(\mathcal{W}) = \big(\mathcal{B}_{r,y}(\mathcal{W})\big)_y$ is self-generating, that is, if $\mathcal{W}_y \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ for each state y. Then each equilibrium payoff of the auxiliary game in state y can be attained with an incentive-compatible strategy profile until the transition to the next state y', at which point the continuation value comes from $\mathcal{W}_{y'} \subseteq \mathcal{B}_{r,y'}(\mathcal{W})$. Thus, the strategy profile can be concatenated with an equilibrium in the auxiliary game of state y' until the next state transition, and so on. An infinite concatenation of these equilibria of the stage games then yields a PPE of the stochastic game.

Similarly to the algorithm in Abreu, Pearce, and Stacchetti [1], an iterated application of the operator $\mathcal{B}_r$ to the set of feasible payoffs converges to the PPE payoff set. At first glance it may seem that the construction of PPE has been reduced to a discrete-time stochastic game, whose period lengths are determined by the state transitions. While incentives at state transitions are dealt with similarly, the continuous-time operator $\mathcal{B}_r$ also ensures that incentives are satisfied at all points between state transitions. This gives rise to two ODEs that describe the boundary of each set $\mathcal{B}_{r,y}(\mathcal{W})$: the *abrupt-information optimality equation* is a first-order ODE that captures optimal incentives from the absence of state transitions and the *optimality equation* is a second-order ODE that, similarly to Sannikov [29], captures optimal value transfers based on the public signal.

Because differential equations can be solved numerically rather easily, an ODE description of equilibrium payoffs also contributes to the literature on computation of equilibrium payoffs in stochastic games. While current algorithms restrict attention to patient players (e.g., Hörner et al. [14]) or perfect monitoring (e.g., Yeltekin, Cai, and Judd [31]), our techniques provide an algorithm for imperfect public monitoring. Different from the two aforementioned algorithms, our approach also yields the equilibrium strategies that attain extremal payoffs, as well as the incentives required to enforce those equilibria.

In many applications, particular attention is given to stationary Markov or Markov-perfect equilibria because of their simplicity.[3] Because stationary Markov equilibria are time independent, their characterization is very similar to the discrete-time counterpart;

[3]Existence of stationary Markov equilibria in behavior strategies in our setting with finitely many actions and states has been established in Fink [9]. Haller and Lagunoff [12] show genericity and Doraszelski and Escobar [6] show stability and purifiability. In more general settings, stationary Markov equilibria may not exist; see Levy [23] and Levy and McLennan [24] for a counterexample. Finally, see He and Sun [13] for the most general existence result to date.

see Haller and Lagunoff [12].[4] In addition to characterizing PPE and stationary Markov equilibrium payoffs, we characterize payoffs of what we call *semi-stationary* PPE. A semi-stationary PPE disregards the public signal completely and it depends on the public history only through the sequence of states that have been attained so far. Note that this requirement is weaker than it would be in a standard discrete-time setting because players can neither infer the time nor the time of state transitions from the sequence of states. Semi-stationary PPE thus strike a balance between simplicity of implementation and attaining a multitude of equilibrium payoffs. We illustrate the different equilibrium notions in an example of a regime-change game.

In the literature on discrete-time stochastic games, techniques often differ drastically between different classes of stochastic games. Analysis of games with an irreducible state process often rely on convergence to a steady state distribution (e.g., Hörner, Sugaya, Takahashi, and Vieille [14] and Peski and Toikka [27]), whereas the analysis of absorbing games often hinges on the fact that there is only one non-absorbing state. The techniques in this paper do not rely on any such properties, hence they are applicable to stochastic games with arbitrary state transitions. The continuous-time techniques do, however, require that the public signal satisfies a pairwise full rank condition to ensure that the aforementioned ODEs admit strong solutions. This condition is stronger than the pairwise identifiability condition in Sannikov [29] because admissible incentives related to state transitions may be discontinuous in the players' continuation value.

It is worth noting that the construction in this paper works for any discount rate and not just in the limit as the discount rate tends to 0. Thus, even for irreducible games, for which we have developed a decent understanding from discrete time, the techniques provide additional insights: since convergence to the steady state distribution is not required, the correspondence $y \mapsto \mathcal{E}_y(r)$ of initial states to equilibrium payoffs is preserved. The presence of absorbing states simplifies the computation as we illustrate in Section 7 as it decouples some of the ODEs.

The remainder of this paper is organized as follows. The model is presented in Section 2. Section 3 contains a motivating example of a contestable democracy. In Section 4, we introduce the notions of enforceability and self-generation in our setting. The set of stationary Markov equilibrium payoffs and semi-stationary equilibrium payoffs are characterized in Section 5. Section 6 describes the characterization of the set of all PPE payoffs. Section 7 discusses notions related to the computation of equilibrium payoffs and Section 8 concludes. Appendix A contains the mathematical foundation for the model and Appendices B and C contains the proofs.

[4]This paper restricts attention to pure-strategy equilibria. This is not a severe restriction for PPE because players can mix artificially by playing different pure actions for a fraction of the time. However, this restriction is more severe for stationary Markov equilibria because those are time independent.

2 Model

Consider a game between 2 players $i = 1, 2$ with a finite set of states $\mathcal{Y}$. State $y \in \mathcal{Y}$ determines the set $\mathcal{A}^i(y)$ of each player i's feasible actions. We assume that $\mathcal{A}^i(y)$ is finite for any state $y \in \mathcal{Y}$ and we denote by $\mathcal{A}(y) := \mathcal{A}^1(y) \times \mathcal{A}^2(y)$ the set of feasible action profiles in state y. We denote by $\mathcal{A} = \bigcup_{y \in \mathcal{Y}} \mathcal{A}(y)$ the set of all action profiles. The chosen action profile $a \in \mathcal{A}(y)$ affects the rate $\lambda_{y,y'}(a)$ at which the state process $S = (S_t)_{t \geq 0}$ transitions from state y to y'. We denote by $\lambda(a)$ the matrix containing the transition intensities from row state to column state. We denote by $\mathcal{Z}(a)$ the set of ordered pairs of states (y, y') with $\lambda_{y,y'}(a) > 0$ and by $\mathcal{Z}$ the set of those (y, y'), for which there exists some $a \in \mathcal{A}(y)$ with $\lambda_{y,y'}(a) > 0$. We say that y' is a *successor state* of y if $(y, y') \in \mathcal{Z}$.

In a game of imperfect monitoring, players cannot observe the actions chosen by their opponents, but only the impact of the chosen actions on the distribution of a publicly observable signal X as well as the state process S. Specifically, we assume that in state y, the chosen action profile a affects the drift rate $\mu(y, a)$ of a d-dimensional Brownian motion Z, that is, $dX_t = \mu(y, a)\, \mathrm{d}t + \mathrm{d}Z_t$.[5,6] The public information $\mathcal{F}_t$ at time t is a σ-algebra that contains the history of the processes S and Z up to time t. It may contain additional information that players can use for public randomization. We denote by $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$ the filtration of public information.

Definition 2.1. A *(public) strategy* A^i of player i is an $\mathbb{F}$-predictable process that takes values in $\mathcal{A}^i(S_-)$. We denote by $A = (A^1, A^2)$ a *(public) strategy profile.*

Remark 2.1. Stochastic games are tractable despite their generality because they are time homogeneous: the set of available actions, the payoffs, and the conditional distribution over future states and public signal realizations do not depend on the time. In a continuous-time setting, this imposes that, conditional on the current state and the chosen strategy profile, the increments of the signal and the state process are independent and identically distributed, that is, they are Lévy increments. It follows from Corollary 10.23 and Theorem 11.4 in Kallenberg [17] that the distribution of the state process and a continuously arriving public signal *have to be* of the chosen form.

Because at any time t, the chosen strategy profile affects the future distribution of the state process and the public signal, play of a strategy profile $A = (A^1, \ldots, A^n)$ induces

[5] In a continuous-time setting, the volatility of the public signal is perfectly observable through the quadratic variation process. Thus, the volatility does not depend on the chosen action profiles in a continuous-time game of imperfect monitoring. We normalize the volatility matrix to be the identity matrix for simplicity of notation.

[6] Changes in the state are also an imperfect public signal of past actions. Discrete arrival of information about past play without impact on the state can be accommodated by having duplicate states y_1, y_2 with identical sets of available actions, payoffs, drift rate of the public signal, and conditional transition probabilities to states in $\mathcal{Y} \setminus \{y_1, y_2\}$. The "transition" between duplicate states is then just a signal about past play. This framework thus allows a mixed signal structure as in Bernard [4].

a family of probability measures $Q^A = \big(Q_t^A\big)_{t\geq 0}$, under which players observe the game. On any interval $[0,T]$ for $T>0$, the public signal takes the form

$$X_t = \int_0^t \mu(S_s, A_s)\,\mathrm{d}s + Z_t^A, \tag{1}$$

under Q_T^A, where $Z^A = Z - \int \mu(S_s,A_s)\,\mathrm{d}s$ is a Q_T^A-Brownian motion describing noise in the continuous component. Moreover, under Q_T^A, state transitions to state y at time t occur with instantaneous intensity $\lambda_{S_{t-},y}(A_t)$. For a mathematical foundation of the model and details on the change of probability measures, see Appendix A.

Simon and Stinchcombe [30] and, more recently, Neyman [26] demonstrate that in continuous-time games of perfect monitoring, strategies may not lead to unique outcomes if they depend on the immediate past of the opponent's chosen actions. This is not a problem in our model: since the public signal has unbounded variation and state changes occur relatively rarely, actions taken by opponents do not immediately generate unambiguous information. Formally, in public monitoring games, one can identify the probability space Ω with the path space of all publicly observable processes, hence a realized path $\omega \in \Omega$ leads to the unique outcome $(A(\omega), S(\omega))$. In particular, each strategy in a public monitoring game is admissible in the sense of Neyman [26].

Definition 2.2.

(i) Each player i receives an unobservable expected flow payoff $g^i : \mathcal{Y}\times\mathcal{A}\to\mathbb{R}$.[7]

(ii) Player i's *discounted expected future payoff* (or *continuation value*) under strategy profile A at any time $t\geq 0$ is given by

$$W_t^i(S_t, A) = \int_t^\infty r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q_s^A}\big[g^i(S_s,A_s)\,\big|\,\mathcal{F}_t\big]\,\mathrm{d}s, \tag{2}$$

where $r>0$ is the discount rate of both players and the distribution of $(S_s)_{s\geq t}$ is determined uniquely by S_t and $(A_s)_{s\geq t}$.

(iii) A strategy profile A is a *perfect public equilibrium (PPE)* for discount rate r if for every state $y\in\mathcal{Y}$, for every player i and all possible deviations $\tilde{A}^i$,

$$W^i(y,A) \geq W^i\big(y, (\tilde{A}^i, A^{-i})\big)\ \text{ a.e.,}^8 \tag{3}$$

where A^{-i} denotes the strategy of player i's opponent in profile A.

[7]The expected flow payoff corresponds to the ex-ante stage game payoff in discrete-time games.

[8]Deviations of a PPE are not profitable *almost everywhere* (a.e.), that is, the inequality $W_t^i(y,A;\omega)\geq W_t^i\big(y,\tilde{A}^i, A^{-i};\omega\big)$ holds for every pair (ω,t) except on a set of $P\otimes$ *Lebesgue*-measure 0. Observe that (3) has to hold for every state $y\in\mathcal{Y}$ because every state occurs with positive probability at any time $t>0$ by the full-support assumption.

(iv) We denote the set of all payoff vectors that are achievable by a perfect public equilibrium when the initial state is y and the discount rate is r by

$$\mathcal{E}_y(r) := \left\{ w \in \mathbb{R}^2 \;\middle|\; \text{there exists a PPE } A \text{ with } W_0(y, A) = w \text{ a.s.} \right\}.$$

We denote by $\mathcal{E}(r)$ the family $\big(\mathcal{E}_y(r)\big)_{y \in \mathcal{Y}}$ of equilibrium payoffs.

The form of the players' continuation value in (2) shows that the players' strategies affect their expected payoffs directly through their expected flow payoff and indirectly, through the impact on the distribution of the public signal and the state process, which is reflected in the change of measure in the expectation operator. Because the weights $r\mathrm{e}^{-r(s-t)}$ in (2) integrate up to one, the continuation value of a strategy profile is a convex combination of "stage game" payoffs. The set of feasible payoff pairs $\mathcal{V} := \operatorname{conv}\{g(y, a) \mid a \in \mathcal{A}(y), y \in \mathcal{Y}\}$ is thus given by the convex hull of pure action payoff pairs. By deviating to their strategy of myopic best replies to their opponent's strategy profile, each player i can ensure an equilibrium payoff of at least

$$\underline{v}^i = \min_{y \in \mathcal{Y}} \min_{a^{-i} \in \mathcal{A}^{-i}(y)} \max_{a^i \in \mathcal{A}^i(y)} g^i\big(y, (a^i, a^{-i})\big).$$

It follows that $\mathcal{V}^* := \{w \in \mathcal{V} \mid w^i \geq \underline{v}^i \text{ for every player } i\}$ is an upper bouund for the set of equilibrium payoffs $\mathcal{E}_y(r)$ in each state $y \in \mathcal{Y}$. Moreover, each $\mathcal{E}_y(r)$ is convex because players are allowed to use public randomization. Indeed, for any two PPE A and A' with expected payoffs $W_0(y, A)$ and $W_0(y, A')$, respectively, any payoff vector $\nu W_0(y, A) + (1 - \nu) W_0(y, A')$ for $\nu \in (0, 1)$ can be attained by selecting either A or A' according to the outcome of a public randomization device at time 0.

Because the set of public perfect equilibria is typically very large, it is worth studying simple classes of PPE that might be easier to implement and, thus, perhaps more reasonable to expect in reality. The most frequently studied of such equilibria are stationary Markov or Markov-perfect equilibria, which depend on the entire public history only through the current state. Since there are only finitely many stationary Markov equilibria, those propose a sharp contrast to the multitude of PPE. The techniques in this model also allow us to easily characterize the payoffs of a family of equilibria between stationary Markov equilibria and general PPE. We call semi-stationary PPE.

Definition 2.3.

(i) A PPE A is a *stationary Markov (perfect) equilibrium* for initial state S_0 if there exists a map $a^* : \mathcal{Y} \to \mathcal{A}$ with $a_*(y) \in \mathcal{A}(y)$ for every state y such that $A = a^*(S_-)$.

(ii) A PPE is *semi-stationary* for initial state S_0 if there exists a map $a_* : \bigcup_{k=1}^{\infty} \mathcal{Y}^k \to \mathcal{A}$ with $a_*(y_1, \dots, y_k) \in \mathcal{A}(y_k)$ for any sequence of states of any length k such that $A = a_*(\hat{S}_-)$, where $\hat{S}_t$ is the sequence of states visited up to and including time t.

(iii) We denote by $\mathcal{E}^M(r)$ and $\mathcal{E}^S(r)$ the families of payoff pairs that are achievable in stationary Markov and semi-stationary equilibria, respectively. Note that

$$\mathcal{E}^M(r) \subseteq \mathcal{E}^S(r) \subseteq \mathcal{E}(r) \subseteq \mathcal{V}^*.$$

A semi-stationary PPE depends on the public history only through the sequence of states that have been attained so far. Note that this requirement is significantly weaker than it would be in a standard discrete-time setting because players can neither infer the time nor the time of state transitions from the sequence of states. This is a significant loss of information since players' actions affect the intensity of state transitions. Consider, for example, a strategy profile A and a deviation $\tilde{A}^i$ so that all state transitions occur ten times more frequently under the deviation than under A. Such a deviation leads to the same distribution over sequences of states as A. If players observe only the sequences of states, they cannot statistically distinguish $(\tilde{A}^i, A^{-i})$ from A, whereas the distinction would be easy given the full information of the state process.

Note that if the state process is cyclic, the sequence of state transitions is determined uniquely from the initial state. Thus, in a cyclic game, the sequence of action profiles taken—but not their duration—in semi-stationary PPE is known at the beginning of the game. This is the case in the following example of a regime-change game: if one ideological group is removed from office, it is because the other group has taken over.

3 Example of a Regime-Change Game

Suppose that proponents of two ideologies compete for power. State y_i for $i = 1, 2$ indicates that group i is in charge. At every instant, the incumbent can either work on policy reforms or manipulate the media, that is, $\mathcal{A}^i(y_i) = \{P, M\}$. The non-incumbent can either accommodate the incumbent or attempt to instigate a revolution, i.e., $\mathcal{A}^{-i}(y_i) = \{A, I\}$. The intensity of state changes is indicated in the left panel of Figure 1. State changes are much more frequent when the non-incumbent is actively trying to instigate a revolution. The incumbent can somewhat counteract this effect by manipulating the media. We suppose that the cooperative actions A and P have a flow cost of $2\,\mathrm{d}t$ and we normalize the cost of the non-cooperative actions I and M to zero.

Suppose that welfare X has two dimensions: X^1 measures components of welfare that are enjoyed equally by the two groups, whereas X^2 measures welfare of issues, on which the two groups have diametrically opposite views. In a sense, (X^1, X^2) is an orthogonal decomposition of welfare. This leads to a realized flow payoff of

$$\mathrm{d}V_t^1 := \mathrm{d}X_t^1 - \mathrm{d}X_t^2 - 2 \cdot 1_{\{A^1 \in \{A,P\}\}}\,\mathrm{d}t, \qquad \mathrm{d}V_t^2 := \mathrm{d}X_t^1 + \mathrm{d}X_t^2 - 2 \cdot 1_{\{A^2 \in \{A,P\}\}}\,\mathrm{d}t.$$

The expected flow payoff is thus $g^i(y, a) = \mu^1(y, a) + (-1)^i \mu^2(y, a) - 2 \cdot 1_{\{A^i \in \{A,P\}\}}$ for

$\lambda_{y,y'}(a)$	A	I
P	0.5	4.5
M	0.25	2.25

$\mu^1(y_i, a)$	A	I
P	7	1
M	2	0

$\mu^2(y_i, a)$	A	I
P	$(-1)^i$	$(-1)^i 3$
M	0	0

$g(y_1, a)$	A	I
P	$(5, 3)$	$(1, -2)$
M	$(2, -1)$	$(0, 0)$

$g(y_2, a)$	P	M
A	$(3, 5)$	$(-1, 2)$
I	$(-2, 1)$	$(0, 0)$

Figure 1: Impact of actions on state transitions and the drift rate of the public signal, as well as expected flow payoffs of the game.

group $i = 1, 2$. The pure action stage-game payoff pairs are summarized in Figure 1.

In addition to being payoff relevant, welfare changes carry information about past play. The chosen action profiles affect the drift rate of welfare as indicated in Figure 1. The total increase in welfare is highest if the incumbent works on policy reforms and the non-incumbent accommodates, which leads to compromises and to more moderate policies being implemented. If the non-incumbent actively instigates a revolution, the total increase in welfare is lower and more one-sided.

Note that each stage game has a unique Nash equilibrium in which both parties choose the cooperative action, giving each group its highest-possible payoffs of the stage game. Thus, if one group was to rule forever, the efficient PPE would demand that the non-incumbent accommodates at all times to ensure moderate policies are implemented. However, the possibility for state changes prevents permanent collaboration as the opposition is tempted to accelerate a state change by instigating a revolution. Figure 2 illustrates the different sets of equilibrium payoffs for this game. There are two stationary Markov equilibria, which involve one group exerting effort at all times and the other group working on policies when they are in charge but instigating a revolution when in opposition. The highest sum of payoffs is attained in two semi-stationary PPE, which exceed the sum of payoffs in the two stationary Markov equilibria. Since permanent collaboration is not an equilibrium, the restriction of stationarity forces that one group always instigates a revolution, hindering growth of welfare. By allowing players to make their actions dependent on the sequence of states, four consecutive terms of collaborative effort can be supported in equilibrium as illustrated in the zoom-in of Figure 2. In the fifth term, the opposition randomizes among its actions to select a stationary continuation equilibrium.

Outside the set of semi-stationary PPE payoffs, the boundary is given by the solution to two coupled ODEs. The solutions of the ODEs reveal the generically unique action profiles that are played on the boundary. Outside the set of semi-stationary PPE payoffs, the opposition mostly attempts to instigate a revolution. Doing so is temporarily costly, but it improves the likelihood of a state change that will put the opposition in power. On the upper part of the boundary $\partial\mathcal{E}_{y_1}(r)$, the incumbent is mostly willing to exert effort since an equilibrium reward for the opposition necessarily also rewards the incumbent.

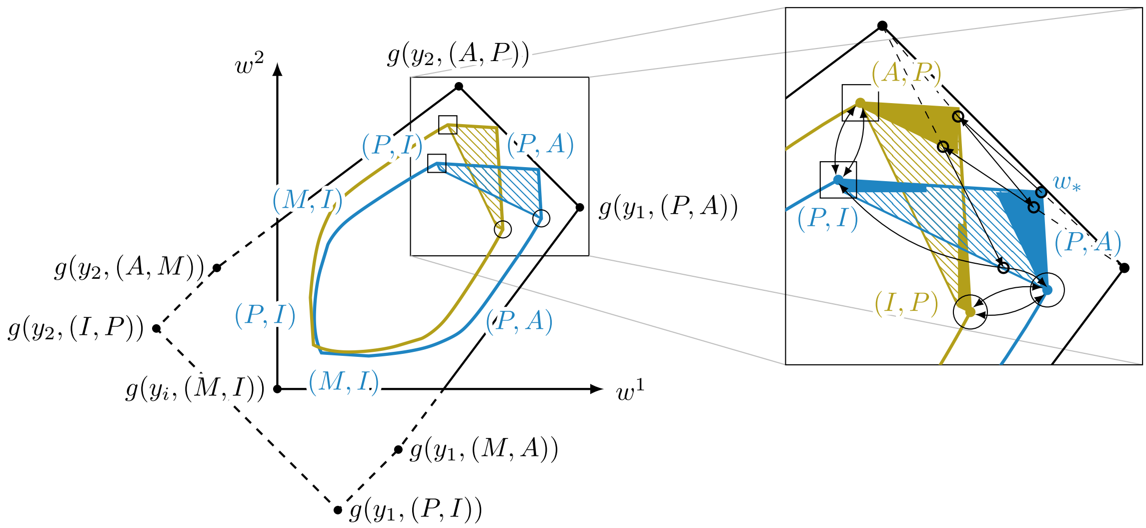

Figure 2: Depicted in color is the family $\mathcal{E}(r) = \{\mathcal{E}_{y_1}(r), \mathcal{E}_{y_2}(r)\}$ of PPE payoff sets. The hatched area are payoff pairs that can be attained semi-stationary PPE. The solid area of the zoom-in are payoff pairs that can be attained by play of a pure action until the next state transition. There are two stationary Markov equilibria, yielding payoff pairs in $\{\hat{w}_{\square}(y_1), \hat{w}_{\square}(y_2)\}$ and $\{\hat{w}_{\circ}(y_1), \hat{w}_{\circ}(y_2)\}$ indicated with squares and circles, respectively. The sample path of a semi-stationary PPE attaining w_* is shown.

This is not the case on the lower part of the boundary $\partial\mathcal{E}_{y_1}(r)$, where the incumbent exerts effort only when the non-incumbent accommodates.

4 Enforceability and Self-Generation

As in any game with imperfect public monitoring, players' incentives are tied to outcomes of the public signal. Thus, we first state the dependence of the continuation value on the public signal and the state process. It will be convenient to denote by

$$J_t^y = \sum_{0<s\leq t} 1_{\{S_s=y, S_{s-}\neq y\}} \tag{4}$$

the process that counts the number of state transitions to state y. Note that the increment $\mathrm{d}J_t^y$ is the increment of a counting process with instantaneous intensity $\lambda_{S_{t-},y}(A_t)$, which means $\mathrm{d}J^y \equiv 0$ if y cannot be reached directly from state S_{t-}. We obtain the following stochastic differential representation of the players' continuation value.

Lemma 4.1. *A semimartingale W is the continuation value of a strategy profile A if and only if W is bounded and for every player $i = 1, \ldots, n$, it holds that*

$$\begin{aligned} \mathrm{d}W_t^i &= r\big(W_t^i - g^i(S_t, A_t)\big)\,\mathrm{d}t + r\beta_t^i\left(\mathrm{d}Z_t - \mu(S_t, A_t)\,\mathrm{d}t\right) \\ &\quad + r\sum_{y\in\mathcal{Y}} \delta_t^i(y)\left(\mathrm{d}J_t^y - \lambda_{S_{t-},y}(A_t)\,\mathrm{d}t\right) + \mathrm{d}M_t^i \end{aligned} \tag{5}$$

for a martingale M^i orthogonal to Z and $(J^y)_{y\in\mathcal{Y}}$ with $M_0^i = 0$, predictable and locally square-integrable processes β^i and $\delta^i(y)$ for $y \in \mathcal{Y}$.

The first term in (5) is a drift term that describes the expected movement of the continuation value. It points away from the expected flow payoff: if $W_t^i < g^i(S_t, A_t)$, then player i extracts an instantaneous payoff rate that exceeds his continuation value and, therefore, doing so has to decrease his future payoff in expectation. The second term is a diffusion term that describes the exposure of the continuation value to the public signal, where the term $r\beta_t^i$ is the sensitivity of the exposure. The third term is a jump term that captures impacts of state changes on the continuation value of player i. The term $r\delta_t^i(y)$ is the instantaneous increase of player i's expected value when the state changes from S_{t-} to y. The final term is a martingale which captures public randomization.

In discrete-time games, incentives are provided by a *continuation promise* that maps the public signal to a promised continuation payoff for every player; see, for example, Abreu et al. [1]. The representation in (5) shows that in continuous-time games, the continuation value is linear in the public signal and hence, so is the continuation promise. To state the incentive-compatibility conditions, it will be convenient to denote by $\lambda(y,a) := \big(\lambda_{y,y_1}(a), \ldots, \lambda_{y,y_{|\mathcal{Y}|}}(a)\big)$ the vector of all transition intensities to successor states when the state is equal to y and action profile a is played.

Definition 4.2. An action profile $a \in \mathcal{A}(y)$ is *enforceable* in state y if there exists a *continuation promise* (β, δ) with $\beta = (\beta^1, \beta^2)^\top$ and $\delta = (\delta^1, \delta^2)^\top$ such that for every player i and every deviation $\tilde{a}^i \in \mathcal{A}^i(y)$,

$$g^i(y,a) + \beta^i \mu(y,a) + \delta^i \lambda(y,a) \geq g^i\big(y, (\tilde{a}^i, a^{-i})\big) + \beta^i \mu\big(y, (\tilde{a}^i, a^{-i})\big) + \delta^i \lambda\big(y, (\tilde{a}^i, a^{-i})\big). \quad (6)$$

A strategy profile A is *enforceable* for initial state y_0 if there exist processes $(\beta_t)_{t\geq 0}$, $(\delta_t)_{t\geq 0}$ such that (6) is satisfied a.e. for the induced state process S.

Action profile a is enforceable in state y if the continuation promise guarantees that the sum of expected instantaneous payoff rate $g^i(y,a)$ and promised continuation rate $\beta^i \mu(y,a) + \delta^i \lambda(y,a)$ for every player i is maximized in a^i over all unilateral deviations. Suppose that players keep their promises and the continuation promise used to enforce A are, in fact, the sensitivities of the continuation value to the public signal and the state process. Then, no player has an incentive to deviate at any point in time and the strategy profile is an equilibrium. This is formalized in the following lemma, which is the continuous-time analogue to the one-shot deviation principle.

Lemma 4.3. *A strategy profile A is a PPE for initial state y_0 if and only if (β, δ) related to A by* (5) *enforces A for the induced state process S with $S_0 = y_0$.*

For the characterization of PPE payoffs—but not stationary Markov and semi-stationary PPE payoffs—we make the following assumptions on stage-game payoffs and the distribution of the public signal.

Definition 4.4. For any state y, any action profile $a \in \mathcal{A}(y)$, and any player $i = 1, 2$, let $M^i_y(a)$ denote the $d \times (|\mathcal{A}^i(y)| - 1)$-dimensional matrix, whose column vectors are given by $\mu(y, \tilde{a}^i, a^{-i}) - \mu(y, a)$ for $\tilde{a}^i \in \mathcal{A}^i(y) \setminus \{a^i\}$. An action profile a is said to have *pairwise full rank* if the matrix $[M^1_y(a), M^2_y(a)]$ has rank $|\mathcal{A}^1(y)| + |\mathcal{A}^2(y)| - 2$.

Assumption 1. For each $y \in \mathcal{Y}$, every action profile $a \in \mathcal{A}(y)$ has pairwise full rank.

Remark 4.1. As in Fudenberg, Levine, and Maskin [11], an action profile a has pairwise full rank if and only if it satisfies two weaker conditions: a has *individual full rank* if $M^i_y(a)$ has rank $|\mathcal{A}^i(y)| - 1$ and a is *pairwise identifiable* if $\operatorname{span} M^1_y(a) \cap \operatorname{span} M^2_y(a) = \{0\}$. Assumption 1 ensures regularity of the boundaries of the family of PPE payoff sets. The individual full rank condition ensures that the optimality equation is locally Lipschitz continuous in incentives from state transitions. As in Sannikov [29], pairwise identifiability guarantees that the characterizing ODE is locally Lipschitz continuous in incentives from the public signal in non-coordinate directions. For local Lipschitz continuity in coordinate directions, we additionally need the following assumption.

Assumption 2. $\operatorname{span} M^1_y(a) \perp \operatorname{span} M^2_y(a)$ for each $y \in \mathcal{Y}$ and each $a \in \mathcal{A}(y)$.

Lemmas 4.1 and 4.3 motivate how we construct equilibrium profiles in continuous time—as solutions to the SDE (5), subject to the enforceability constraint (6). To do so, we use the fact that stochastic games are time homogeneous: since the continuation profile of a PPE after any time t is also an equilibrium of the entire game, the continuation value has to remain within the family of equilibrium payoff sets at all times. This property is known as self generation. In our setting, it is formalized as follows.

Definition 4.5. A family of sets $(\mathcal{W}_y)_{y \in \mathcal{Y}} \subseteq \mathbb{R}^n$ is *self-generating* if for every $y \in \mathcal{Y}$ and every $w \in \mathcal{W}_y$, there exists a solution $\big(W, A, \beta, \delta, M\big)$ to (5) such that (β, δ) enforces A, $S_0 = y$ a.s., $W_0 = w$ a.s., and $W_\tau \in \mathcal{W}_{S_\tau}$ a.s. for every stopping time τ.[9]

Remark 4.2. The stochastic differential equation (5) does not in general admit strong solutions, that is, it may not be possible to solve (5) for a fixed Brownian motion and a fixed state process. In the proofs that are contained in the appendices, we use weak solutions to (5), in which the Brownian motion, the state process, and the probability space are part of the solution. To keep the notation simple, we do not make this distinction in the main text. See Appendix B for details.

Similarly as in discrete-time games, the family of equilibrium payoff sets is the largest family of bounded self-generating sets.[10]

[9] The processes $(J^y)_{y \in \mathcal{Y}}$ appearing in (5) are defined from S as in (4).

[10] One can show that the union $(\mathcal{W}_y \cup \mathcal{V}_y)_{y \in \mathcal{Y}}$ of two families $(\mathcal{W}_y)_{y \in \mathcal{Y}}$, $(\mathcal{V}_y)_{y \in \mathcal{Y}}$ of self-generating sets is again self-generating. Thus, there is, in fact, a largest bounded self-generating family of payoff sets equal to the union of all such families of payoff sets.

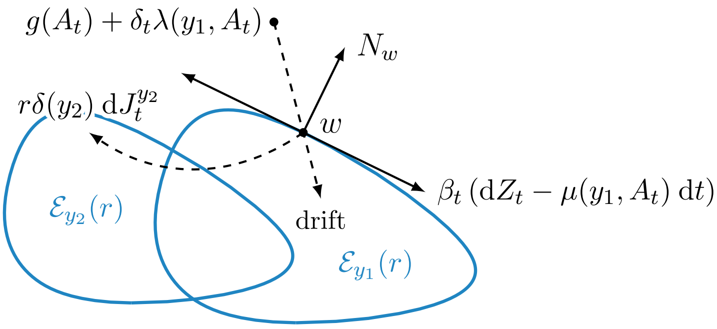

Figure 3: While the state is equal to y_1, the continuation value of a PPE moves continuously in $\mathcal{E}_{y_1}(r)$. At the boundary, self-generation implies that the continuation value has to move tangentially to the boundary, given through the incentives arising from the public signal. When the state changes to y_2, the continuation value changes discontinuously to a value in $\mathcal{E}_{y_2}(r)$.

Lemma 4.6. *The family $\big(\mathcal{E}_y(r)\big)_{y\in\mathcal{Y}}$ is the largest family of bounded self-generating sets.*

The characterization of $\mathcal{E}(r)$ as the largest family of bounded self-generating sets allows us to construct equilibria using a stochastic control approach. Self-generation implies that the continuation value of each PPE corresponds to an enforceable solution to (5) that locally remains in $\mathcal{E}_y(r)$ while the state is equal to y, which jumps to $\mathcal{E}_{y'}(r)$ when the state transitions to state y'. From (5), it follows that for any $w \in \partial\mathcal{E}_y(r)$ and any normal vector N_w to $w \in \partial\mathcal{E}_y(r)$, it is necessary that:

(I1) The drift points towards the interior of the set: ${N_w}^\top\big(g(y,a) + \delta\lambda(y,a) - w\big) \geq 0$,

(I2) The volatility is tangential to the set: ${N_w}^\top\beta = 0$,

(I3) State transitions are within $\mathcal{E}(r)$: $w + r\delta(y') \in \mathcal{E}_{y'}(r)$ for every y' with $(y, y') \in \mathcal{Z}(a)$.

Observe that restriction (I2) is necessary because of unbounded variation of Brownian motion: if incentives were provided in any non-tangential direction, the continuation value would escape $\mathcal{E}(r)$ immediately; see Figure 3.

Sannikov [29] has shown that in the special case of a repeated game, i.e., when $\mathcal{Y}$ is a singleton, one can relate the law of motion of the continuation value to the curvature of the equilibrium payoff set. Crucial in this regard is the fact that there are no state changes: the continuation value is continuous and has to move tangentially to the set. Incentives that can be provided at any given point on the boundary thus depend on the geometry of the payoff set only through local information, i.e., the direction of the tangent. This local dependence gives rise to an ordinary differential equation. Bernard [4] shows that when information also arrives discontinuously, such a local characterization is no longer possible because restriction (I3) depends on global information of the equilibrium payoff set. Nevertheless, one can approximate the boundary of the equilibrium payoff set through a sequence of ordinary differential equations. Key is the relaxation of restriction (I3) to a condition that requires continuation values after a state change to come from a fixed payoff set. We shall pursue a similar approach here.

Definition 4.7. Fix a family of payoff sets $\mathcal{W} = (\mathcal{W}_y)_{y\in\mathcal{Y}}$. Let σ denote the occurrence of the first state change.

(i) We say that a family $\mathcal{X} = (\mathcal{X}_y)_{y\in\mathcal{Y}}$ of payoff sets is $\mathcal{W}$*-relaxed self-generating* if for each $y \in \mathcal{Y}$ and each $w \in \mathcal{X}_y$, there exists a solution $(W, A, \beta, \delta, Z, S, M)$ to (5) such that $W_0 = w$ a.s., $S_0 = y$ a.s., $W_\sigma \in \mathcal{W}_{S_\sigma}$ a.s., and for almost every (ω, t) with $0 \le t < \sigma(\omega)$, we have $W_t(\omega) \in \mathcal{X}_{y_0}$ and $(\beta_t(\omega), \delta_t(\omega))$ enforces $A_t(\omega)$ in $S_t(\omega)$.

(ii) $\mathcal{B}_r(\mathcal{W}) = \big(\mathcal{B}_{r,y}(\mathcal{W})\big)_{y\in\mathcal{Y}}$ is the largest family of $\mathcal{W}$-relaxed self-generating payoff sets.

The operator $\mathcal{B}_r$ is a continuous-time extension of the standard set operator in Abreu, Pearce and Stacchetti [1] to stochastic games. Payoff pairs in $\mathcal{B}_{r,y}(\mathcal{W})$ for any state $y \in \mathcal{Y}$ can be attained by an enforceable strategy profile with continuation promise from $\mathcal{W}_{S_\sigma}$ at time σ. The additional requirement due to the continuous-time setting is that before the state transition, the continuation value remain in $\mathcal{B}_{r,y}(\mathcal{W})$. In particular, incentives related to the absence of state transitions and the arrival of the public signal have to enforce the strategy profile at all times before the state transition. Similarly to [1], the operator $\mathcal{B}_r$ is closely related to the concept of self-generation.

Lemma 4.8. *Let $\mathcal{W} = (\mathcal{W}_y)_{y\in\mathcal{Y}}$ such that $\mathcal{W}_y \subseteq \mathcal{V}$ for every state $y \in \mathcal{Y}$. If $\mathcal{W}$ is self-generating, then $\mathcal{W}_y \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ for every state $y \in \mathcal{Y}$. Conversely, if $\mathcal{W}_y \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ for every state $y \in \mathcal{Y}$, then $\mathcal{B}_r(\mathcal{W})$ is self-generating.*

Since payoff pairs in $\mathcal{B}_{r,y}(\mathcal{W})$ can be attained with a continuation promise in $\mathcal{W}_{S_\sigma}$ at time σ, the condition that $\mathcal{W}_{y'} \subseteq \mathcal{B}_{r,y'}(\mathcal{W})$ for every y' allows us to attain W_σ with an enforceable strategy profile that remains in $\mathcal{B}_{r,S_\sigma}(\mathcal{W})$ until the next state transition, and so on. Because Poisson processes have only countably many jumps, repeating this concatenation argument countably many times yields solutions that remain in $\mathcal{B}_{r,S}(\mathcal{W})$ forever. The proof in Appendix B additionally deals with some measurability issues of the concatenation. We obtain the following algorithm to approximate $\mathcal{E}(r)$.

Proposition 4.9. *Let $\mathcal{W}_0 = (\mathcal{W}_{0,y})_{y\in\mathcal{Y}}$ be the family of payoff sets with $\mathcal{W}_{0,y} = \mathcal{V}^*$ for every $y \in \mathcal{Y}$. Define the sequence $(\mathcal{W}_n)_{n\ge0}$ iteratively via $\mathcal{W}_n = \mathcal{B}_r(\mathcal{W}_{n-1})$ for $n \ge 1$. Then $(\mathcal{W}_{n,y})_{n\ge0}$ is decreasing in the set-inclusion sense for every $y \in \mathcal{Y}$ with $\bigcap_{n\ge0} \mathcal{W}_{n,y} = \mathcal{E}_y(r)$.*

This is an extension of the algorithm in Abreu, Pearce, and Stacchetti [1] to continuous-time stochastic games. We will show in the next section that the boundary of the resulting set in each step of the iteration admits a characterization by a differential equation. This is possible because in each step, the restriction on the use of abrupt information is fixed. Contrary to the algorithm in Hölder, Sugaya, Takahashi, and Vieille [14], our algorithm does not require players to be arbitrarily patient. As a consequence, the algorithm is applicable to wider class of games that violate Assumption A in [14] that the limit equilibrium payoff set is independent of the initial state. Their assumption A is typically encountered in the form of irreducibility of the state process. However, there are other sufficient conditions; see Dutta [7].

Remark 4.3. Such an approximation of the equilibrium payoff set via a notion of relaxed self-generating payoff sets was used also in Bernard [4] to approximate the equilibrium payoff set in continuous-time repeated games with discontinuous information arrival. Since there is only one state in a repeated game, a "state transition" simply corresponds to a discontinuous signal about past play. The repeated-game operator $\tilde{\mathcal{B}}_r$ requires that continuation values after every discontinuous signal come from the same payoff set, but is otherwise identical to $\mathcal{B}_r$. This similarity will help us greatly in the construction of $\mathcal{B}_r(\mathcal{W})$ for an arbitrary family of payoff sets $\mathcal{W}$ in Section 6.

5 Stationary Markov and Semi-Stationary Equilibria

In a semi-stationary equilibrium, the public signal is not used to provide incentives because the public signal does not affect future play. Therefore, incentives have to be tied to state changes, i.e., inequalities (6) have to be satisfied for $\beta = 0$. Since the time between two state changes is exponentially distributed, it may take an arbitrarily long time for a state change to occur. A solution to (5) with $\beta = 0$ is thus either locally constant (if the drift rate is 0) or it diverges with positive probability (if the drift rate is not 0). Because $\mathcal{E}^S(r)$ is bounded, the continuation value in a semi-stationary equilibrium must be locally constant. One could thus define a notion of self-generation similarly to Definition 4.5, requiring that any payoff be attainable by a locally constant enforceable solution to (5). We state the following equivalent condition in terms of primitives.

Definition 5.1. A family $\mathcal{W} = (\mathcal{W}_y)_{y \in \mathcal{Y}}$ of payoff sets is *mutually stationary* if for every $y \in \mathcal{Y}$ and every $w \in \mathcal{W}_y$, there exists $a \in \mathcal{A}(y)$ and δ such that $(0, \delta)$ enforces a with $w = g(y, a) + \delta\lambda(y, a)$, and $w + r\delta(y') \in \mathcal{W}_{y'}$ for every y' with $(y, y') \in \mathcal{Z}(a)$.

Similarly to Lemma 4.6, we obtain the following characterization of $\mathcal{E}^S(r)$.

Lemma 5.2. *The family of semi-stationary equilibrium payoffs $\mathcal{E}^S(r)$ is the largest bounded family of mutually stationary payoff sets.*

Mutual stationarity is the analogue to self-generation for semi-stationary PPE. The set $\mathcal{E}^S(r)$ can be approximated with a similar localized construction, where we impose that the continuation value after a state change comes from a fixed family of payoff sets.

Definition 5.3. Fix a family $\mathcal{W} = (\mathcal{W}_y)_{y \in \mathcal{Y}}$ of payoff sets.

(i) We say that a family $\mathcal{X} = (\mathcal{X}_y)_{y \in \mathcal{Y}}$ of payoff sets is $\mathcal{W}$-*stationary* if for each $y \in \mathcal{Y}$ and each $w \in \mathcal{X}_y$, there exist $a \in \mathcal{A}(y)$ and $\delta \in \mathbb{R}^{|\mathcal{Y}|}$ such that $(0, \delta)$ enforces a in state y, $w = g(y, a) + \delta\lambda(y, a)$, and $w + r\delta(y') \in \mathcal{W}_{y'}$ for every y' with $(y, y') \in \mathcal{Z}(a)$.

(ii) Let $\mathcal{S}_r(\mathcal{W}) = \big(\mathcal{S}_{r,y}(\mathcal{W})\big)_{y \in \mathcal{Y}}$ denote the largest family of $\mathcal{W}$-stationary payoff sets.

Observe that $\mathcal{S}_{r,y}(\mathcal{W}) \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ for any state $y \in \mathcal{Y}$: any payoff pair $w \in \mathcal{S}_{r,y}(\mathcal{W})$ can be attained by an enforceable solution to (5) that remains in w until the first state transition occurs, at which point the continuation value jumps to $\mathcal{W}_{S_\sigma}$. Therefore, the singleton $\{w\}$ is $\mathcal{W}$-relaxed self-generating and hence contained in $\mathcal{B}_{r,y}(\mathcal{W})$. Similarly to Lemma 4.8 and Proposition 4.9, we obtain the analogue results for semi-stationary PPE.

Lemma 5.4. *If $\mathcal{W}_y \subseteq \mathcal{S}_{r,y}(\mathcal{W})$ for every state y, then $\mathcal{S}_r(\mathcal{W})$ is mutually stationary.*

Proposition 5.5. *Let $\mathcal{W}_0$ be as in Proposition 4.9 and let $\mathcal{W}_n := \mathcal{S}_r(\mathcal{W}_{n-1})$ for any $n \geq 1$. Then $(\mathcal{W}_{n,y})_{n\geq 0}$ is decreasing for every $y \in \mathcal{Y}$ with $\bigcap_{n\geq 0} \mathcal{W}_{n,y} = \mathcal{E}^S_y(r)$.*

As stated, the construction in Proposition 5.5 converges to the set of semi-stationary equilibrium payoffs when players are not allowed to use public randomization. If, in each step of the algorithm, we instead take the convex hull of $\mathcal{S}_{r,y}$, then the resulting set is the set of all semi-stationary PPE payoffs with the use of public randomization.

5.1 Characterization of $\mathcal{S}_r(\mathcal{W})$

We begin by illustrating the computation of $\mathcal{S}_{r,y}(\mathcal{W})$ if y has at most one successor state. If y has no successor state, then $\mathcal{S}_{r,y}(\mathcal{W})$ is simply the set of static Nash equilibrium payoffs of stage game y. If y has exactly one successor state y', then δ^i in (5) is one-dimensional. Let $\mathcal{S}_{r,y,a}(\mathcal{W})$ denote the set of all stationary payoffs, for which there exists δ_0 such that $(0, \delta_0)$ enforces a, $w = g(y,a)+\delta_0\lambda(y,a)$, and $w+r\delta_0 \in \mathcal{W}_{y'}$. The stationarity condition of a payoff $w \in \mathcal{S}_{r,y,a}(\mathcal{W})$ implies that $w + r\delta_0 = g(y,a) + \delta_0\lambda(y,a) + r\delta_0 \in \mathcal{W}_y$. The set of all payoffs that can be reached from $\mathcal{S}_{r,y,a}(\mathcal{W})$ after a state transition is thus

$$\big(g(y,a) + \Psi_y(a)(\lambda(y,a) + r)\big) \cap \mathcal{W}_{y'}, \tag{7}$$

where $\Psi_y(a) := \{\delta \mid (0,\delta) \text{ enforces } a \text{ in state } y\}$. Note that (7) consists only of scalings, translates, and intersections of closed, convex sets and is thus easily computed. In order to get $\mathcal{S}_{r,y,a}(\mathcal{W})$, we simply shrink the set in (7) towards $g(y,a)$ by a factor $\frac{\lambda(y,a)}{\lambda(y,a)+r}$; see Figure 4. Repeating this construction for all action profiles $a \in \mathcal{A}(y)$ yields $\mathcal{S}_{r,y}(\mathcal{W})$.

If a state y has more than one successor state, an a analogue construction is carried out in the incentive space rather than in the payoff space. The condition that the continuation value after a transition to state y' comes from the set $\mathcal{W}_{y'}$ can be expressed as $g(y,a) + f_{y,y'}(\delta) \in \mathcal{W}_{y'}$, where $f_{y,y'}(\delta) = \delta\big(\lambda(y,a) + re_{y'}\big)$ and $e_{y'}$ is the unit vector in the direction of state y' among all successor states of y. This eliminates the variable w and allows us to parametrize the set of stationary payoffs via incentives δ. Such incentives have to come from $\Psi_y(a)$ and they have to satisfy the jump condition for every $y \in Y$, i.e., the set of all such incentives is given by

$$\mathcal{X}_{y,a}(\mathcal{W}) := \Psi_y(a) \cap \bigcap_{y':(y,y')\in\mathcal{Z}(a)} f^{-1}_{y,y'}\big(\mathcal{W}_{y'} - g(y,a)\big),$$

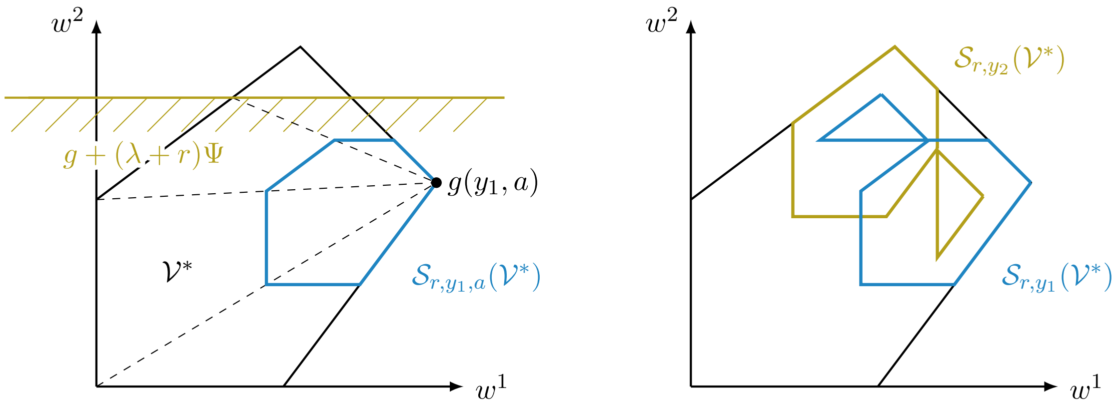

Figure 4: In the regime-change example, mutual effort $a = (P, A)$ can be enforced in state y_1 by any δ with $\delta^2 \leq 1.25$ and $\delta^1 \geq -12$. As illustrated in the left panel, the set $\mathcal{S}_{r,y_1,a}(\mathcal{V}^*)$ is given by shrinking the set of $(g + (\lambda + r)\Psi) \cap \mathcal{V}^*$ towards g. The right panel shows the family $\mathcal{S}_r(\mathcal{V}^*)$.

where $f_{y,y'}^{-1}$ denotes the inverse image under $f_{y,y'}$. It is now straightforward that

$$\mathcal{S}_{r,y}(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \big(g(y,a) + \mathcal{X}_{y,a}(\mathcal{W})\lambda(y,a)\big), \tag{8}$$

where $\mathcal{X}_{y,a}(\mathcal{W})\lambda(y,a) := \{v \in \mathbb{R}^2 \mid \exists\, \delta \in \mathcal{X}_{y,a}(\mathcal{W}) \text{ with } \delta\lambda(y,a) = v\}$ denotes the projection of $\mathcal{X}_{y,a}(\mathcal{W})$ onto $\mathbb{R}^2$ in the direction $\lambda(y,a)$.

Note that $\Psi_y(a)$ is a closed, convex polytope characterized by the affine inequalities in (6). Consider first the case, where $\mathcal{W}$ is a family of polygons $\mathcal{W}_y$ with extremal points $v_{y,1}, \ldots, v_{y,n_y}$ and corresponding normal vectors $N_{y,1}, \ldots, N_{y,n_y}$. Then the set $\mathcal{X}_{y,a}(\mathcal{W})$ is a convex polytope, characterized by the affine inequalities in (6) and

$$N_{y',j}{}^{\top}\big(\delta(\lambda(y,a) + re_{y'})\big) \leq N_{y',j}{}^{\top}x_{y',j}, \quad j = 1, \ldots, n_{y'},\ (y,y') \in \mathcal{Z}(a).$$

One can thus compute extremal points of $\mathcal{X}_{y,a}(\mathcal{W})\lambda(y,a)$ by projecting the extremal points of $\mathcal{X}_{y,a}(\mathcal{W})$ onto $\mathbb{R}^2$ in the direction $\lambda(y,a)$. Alternatively, one can maximize $N^{\top}\delta\lambda(y,a)$ over $\delta \in \mathcal{X}_{y,a}(\mathcal{W})$ for a sufficiently rich grid of normal vectors N. This is efficient numerically because it maximizes a linear function under a set of affine constrains. Since the projection of a polytope is a polygon, $\mathcal{M}_{r,y}(\mathcal{W})$ is a polygon again. In particular, $\mathcal{W}_n$ in the approximating sequence of Proposition 5.5 is a family of polygons for every n. However, it is not guaranteed that the limiting set $\mathcal{E}^S(r)$ is also a polygon.

For general $\mathcal{W}$, one can use inner and outer polygon approximations $\underline{\mathcal{W}}_y$ and $\overline{\mathcal{W}}_y$ of $\mathcal{W}_y$, respectively, as in Judd, Yeltekin, and Conklin [16] and then use the steps above.

5.2 Stationary Markov Equilibria

Stationary Markov equilibria are a special class of semi-stationary equilibria that depend on the sequence of observed states only through the current state. For initial state y_0, let $\mathcal{Y}(y_0)$ denote the set of all states that can be reached from y_0, including y_0 itself. The following is a verification result for whether or not a given map from states to action profiles is a stationary Markov perfect equilibrium.

Proposition 5.6. *A map $a_* : \mathcal{Y} \to \mathcal{A}$ with $a_*(y) \in \mathcal{A}(y)$ for every y is a stationary Markov perfect equilibrium for initial state y_0 if and only if for every state $y \in \mathcal{Y}(y_0)$, every player i, and every $a^i \in \mathcal{A}^i(y)$, we have*

$$w_*^i(y) \geq g^i(y, (a^i, a_*^{-i}(y))) + \sum_{y' \in \mathcal{Y}} \delta_*^i(y, y') \lambda_{y,y'}(a^i, a_*^{-i}(y)), \tag{9}$$

where we denote by $\Lambda_{y_0}(a_*)$ the matrix with entries $\lambda_{y,y'}(a_*(y))$ in row y' and column y for $(y, y') \in \mathcal{Y}(y_0)^2$, by $G_{y_0}(a_*)$ and w_* the matrices with entries $g^i(y, a_*(y))$ and $w_*^i(y)$, respectively, in row i and column corresponding to $y \in \mathcal{Y}(y_0)$, by $\mathbf{1}$ the $|\mathcal{Y}(y_0)|$-dimensional row vector containing all ones, and

$$w_* := rG_{y_0}(a_*)\big(\operatorname{diag}(r\mathbf{1} + \mathbf{1}\Lambda_{y_0}(a_*)) - \Lambda_{y_0}(a_*)\big)^{-1}, \quad \delta_*(y, y') := \frac{w_*(y') - w_*(y)}{r}. \tag{10}$$

Moreover, $W(a_*(S_-)) = w_*(S)$, where S is the state process starting in y_0.

It is worth noting that for a fixed map a_* from states to action profiles, w_* and δ_* are given explicitly by the respective expressions in (10). A naive algorithm is thus to verify the conditions of all such maps a_* since we restrict attention to pure-strategy equilibria in this paper. While this comes at almost no loss of generality for continuous-time PPE, it is a significant restriction for stationary Markov equilibria; see also Footnote 4.

6 Characterization of $\mathcal{B}_r(\mathcal{W})$

In Section 4, we motivated the construction of equilibrium profiles as enforceable solutions to (5). Due to the iterative procedure in Proposition 4.9, it is sufficient to construct such solutions up until the first state transition if we additionally impose that continuation values after a state transition come from a fixed family of payoff sets. This local construction of strategy profiles can be viewed as equilibria of an auxiliary game that ends at the time of a state transition with terminal payoffs from $\mathcal{W} = (\mathcal{W}_y)_{y \in \mathcal{Y}}$. A "state transition" in the auxiliary game does not really correspond a transition of states but simply to a discrete signal about past play that ends the game. The auxiliary game with initial state $y \in \mathcal{Y}$ is thus a *repeated game* with expected flow payoff rate $g_y(a) := g(y, a)$, drift rate $\mu_y(a) := \mu(y, a)$ of the public Brownian signal, and intensities $\lambda_y(y' \,|\, a) := \lambda_{y,y'}(a)$ of the discrete signals. The equilibrium payoff set of the auxiliary game is $\mathcal{B}_{r,y}(\mathcal{W})$.

The fact that this local construction is the same as in a continuous-time repeated game with public signal (X, S) means that we can characterize $\mathcal{B}_{r,y}(\mathcal{W})$ for a fixed state y

directly with the techniques from continuous-time repeated games; see also Remark 4.3. The construction of local incentives is thus identical to those in in Bernard [4], hence this section follows [4] rather closely.

We have already seen in the previous section that the family of $\mathcal{W}$-stationary payoffs is contained in $\mathcal{B}_r(\mathcal{W})$ and how one can characterize stationary payoffs. In a non-stationary payoff, the public signal as well as the history of the state process may be used to provide intertemporal incentives. Incentives at the boundary of $\mathcal{B}_{r,y}(\mathcal{W})$ must satisfy an analogue to informational restrictions (I1)–(I3), described in the following definition.

Definition 6.1. For a payoff pair $w \in \mathbb{R}^2$, a direction $N \in S^1$, and a family $\mathcal{W}$ of payoff sets, we say that a continuation promise (β, δ) from the set

$$\Xi_{y,a}(w, N, \mathcal{W}) := \left\{ (\beta, \delta) \;\middle|\; \begin{array}{l} (\beta, \delta) \text{ enforces } a \text{ in state } y, N^\top\big(g(y, a) + \delta\lambda(y, a) - w\big) \geq 0, \\ N^\top \beta = 0, \text{and } w + r\delta(y') \in \mathcal{W}_{y'} \text{ for all } y' \text{ with } (y, y') \in \mathcal{Z}(a). \end{array} \right\}$$

restricted-enforces a in state y. An action profile $a \in \mathcal{A}$ is *restricted-enforceable* for $(w, N, \mathcal{W})$ in state y if the set $\Xi_{y,a}(w, N, \mathcal{W})$ is non-empty.

We will distinguish between payoff pairs on the boundary, where the public signal is needed and where it is not needed to provide intertemporal incentives. Due to informational restriction (I1), necessity of the public signal at a boundary point depends not only on the location of the payoff pair but also on the direction of the outward normal vector.

Definition 6.2. Denote by $\Psi_{y,a}(w, \mathcal{W})$ the set of all δ such that $(0, \delta)$ enforces a in state y with $w + r\delta(y') \in \mathcal{W}_{y'}$ for all successor states y' of y. Define

$$\Gamma_y(\mathcal{W}) := \left\{ (w, N) \in \mathbb{R}^2 \times S^1 \;\middle|\; \begin{array}{l} \text{There exist } a \in \mathcal{A} \text{ and } \delta \in \Psi_{y,a}(w, \mathcal{W}) \\ \text{with } N^\top\big(g(y, a) + \delta\lambda(y, a) - w\big) \geq 0 \end{array} \right\}.$$

For any convex set $\mathcal{X}$ and any $w \in \partial\mathcal{X}$, let $\mathcal{N}_w(\mathcal{X})$ denote the set of outward normal vectors to $\mathcal{X}$ at w. Denote by $\mathcal{N}_{\mathcal{X}} := \{(w, N) \mid w \in \partial\mathcal{X} \text{ and } N \in \mathcal{N}_w(\mathcal{X})\}$ the *outward normal bundle* of $\mathcal{X}$. With slight abuse of terminology, we will refer to the boundary of $\mathcal{B}_{r,y}(\mathcal{W})$ within $\Gamma_y(\mathcal{W})$ when referring to boundary points $w \in \partial\mathcal{B}_{r,y}(\mathcal{W})$, for which there exists an outward normal vector N with $(w, N) \in \Gamma_y(\mathcal{W})$.

6.1 Characterization of $\partial\mathcal{B}_{r,y}(\mathcal{W})$ outside of $\Gamma_y(\mathcal{W})$

Consider, for a moment, that an enforceable solution W to (5) for initial state $y \in \mathcal{Y}$ remains on a continuously differentiable curve $\mathcal{C}$ before any state transitions occur.[11] Then

[11] A curve $\mathcal{C}$ is continuously differentiable if it has a unique normal vector N_w (up to orientation of the curve) at every point $w \in \mathcal{C}$ such that the Gauss map $w \mapsto N_w$ is continuous.

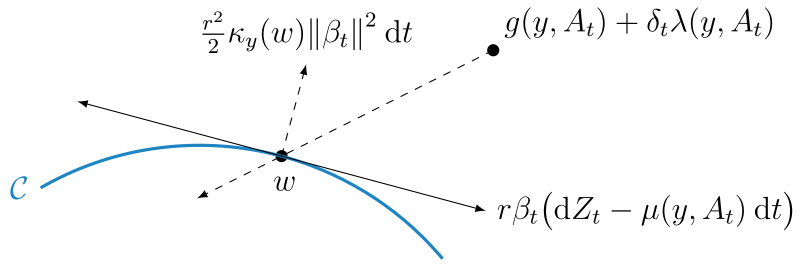

Figure 5: The infinitesimally quick oscillation in the continuation value due to unbounded variation of value transfers causes the continuation value to drift away from the curve in orthogonal direction.

the public signal must be used to provide incentives via tangential value transfers as illustrated in Figure 5. Due to unbounded variation of Brownian motion, players transfer value very rapidly. The infinitesimally quick tangential oscillation created by those transfers causes the continuation value to drift away from the curve in orthogonal direction, similar in spirit to the centrifugal force in physics. The larger the curvature is, the stronger is the outward drift created by the tangential transfers. The magnitude of the outward drift is given by Itō's formula and it is equal to $\frac{r^2}{2}\kappa_y(w)\|\beta_t\|^2$. For the continuation value to stay on the curve, the outward drift has to be offset precisely by the inward drift, hence

$$\frac{r^2}{2}\kappa_y(w)\|\beta_t\|^2 = rN_w{}^{\top}\big(g(y,A_t)+\delta_t\lambda(y,A_t)-w\big).$$

To complete the argument, it is necessary that whenever the continuation value revisits the same point on the curve, the tangential transfers induce the same curvature. This is the case when the chosen action profiles and the provided continuation promise are Markovian in the continuation value. This discussion is summarized in the following lemma.

Lemma 6.3. *Fix a state $y \in \mathcal{Y}$ and suppose that Assumptions 1–2 hold. Let $\mathcal{C}$ be a continuously differentiable curve oriented by the Gauss map $w \mapsto N_w$ such that:*

(i) There exist measurable selectors a^*, δ^*, and β^* on $\mathcal{C}$ such that the selections satisfy $\beta^*(w) \neq 0$ and $\big(\beta^*(w),\delta^*(w)\big) \in \Xi_{y,a^*(w)}\big(w,N_w,\mathcal{W}\big)$ for any $w \in \mathcal{C}$ and the curvature at any point $w \in \mathcal{C}$ is given by

$$\kappa_y(w) = \frac{2N_w{}^{\top}\big(g\big(y,a^*(w)\big)+\delta^*(w)\lambda\big(y,a^*(w)\big)-w\big)}{r\|\beta^*(w)\|^2}. \tag{11}$$

(ii) $\mathcal{C}$ is a closed curve or both of its endpoints are contained in $\mathcal{B}_{r,y}(\mathcal{W})$.

Then $\mathcal{C} \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ and the solution W to (5) with $M \equiv 0$, $A = a^*(W_-)$, $\delta = \delta^*(W_-)$, and $\beta = \beta^*(W_-)$ remains on $\mathcal{C}$ until an endpoint of $\mathcal{C}$ is reached or a state transition occurs.

The curvature of $\partial\mathcal{B}_{r,y}(\mathcal{W})$ outside of $\Gamma_y(\mathcal{W})$ arises from (11) by taking the maximum over all restricted-enforceable action profiles and their continuation promises.

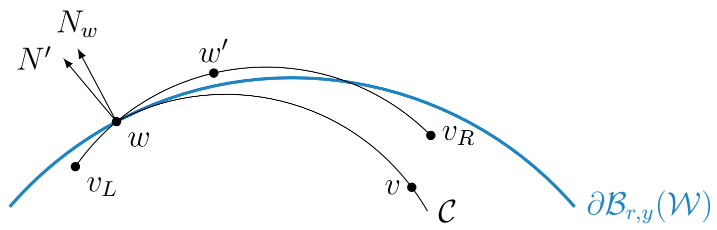

Figure 6: If a solution $\mathcal{C}$ to (12) starting at (w, N_w) falls into the interior of $\mathcal{B}_{r,y}(\mathcal{W})$, then a solution $\mathcal{C}'$ with initial conditions (w, N') for a slight rotation N' of N_w would leave and re-enter $\mathcal{B}_{r,y}(\mathcal{W})$. Lemma 6.3 then implies that $w' \in \mathcal{B}_{r,y}(\mathcal{W})$, which is a contradiction.

Lemma 6.4. *Suppose that Assumptions 1–2 hold. Outside of $\Gamma_y(\mathcal{W})$, the boundary $\partial\mathcal{B}_{r,y}(\mathcal{W})$ is continuously differentiable with curvature at almost every point w given by*

$$\kappa_y(w) = \max_{a\in\mathcal{A}} \max_{(\beta,\delta)\in\Xi_{y,a}(w,N_w,\mathcal{W})} \frac{2{N_w}^{\top}\big(g(y,a) + \delta\lambda(y,a) - w\big)}{r\|\beta\|^2}. \tag{12}$$

Suppose that a solution $\mathcal{C}$ to (12) starting at (w, N_w) falls into the interior of $\mathcal{B}_{r,y}(\mathcal{W})$, then a solution $\mathcal{C}'$ with initial conditions (w, N') for a slight rotation N' of N_w would leave and re-enter $\mathcal{B}_{r,y}(\mathcal{W})$; see Figure 6. That is, it would attain a payoff pair w' strictly outside of $\mathcal{B}_{r,y}(\mathcal{W})$. But Lemma 6.3 implies that $w' \in \mathcal{B}_{r,y}(\mathcal{W})$, which is a contradiction. Note that this argument requires continuity of solutions to (12) in initial conditions. On the other hand, a solution $\mathcal{C}$ to (12) cannot escape $\mathcal{B}_{r,y}(\mathcal{W})$ either because $\mathcal{C}$ maximizes the curvature over all restricted-enforceable action profiles and their incentives. Any other enforceable strategy profile that is played has to involve non-tangential value transfers, which causes the continuation value to grow arbitrarily large with positive probability.

Continuity of the solution $\mathcal{C}'$ in initial conditions (w, N') in this argument requires that κ_y is locally Lipschitz continuous in (w, N). One can show that for a fixed action profile, the right-hand side of (12) is locally Lipschitz continuous at (w, N) in the interior of the effective domain $\{(w, N) \mid \Xi_{y,a}(w, N, \mathcal{W}) \neq \emptyset\}$ of $\Xi_{a,y}$. The expression for the curvature may thus fail to be Lipschitz continuous only where the maximizing action profile in (12) fails to be restricted-enforceable within a small neighborhood. Assumptions 1 and 2 guarantee that the maximizing action profile is restricted-enforceable for small perturbations in w and N.

6.2 Characterization of $\partial\mathcal{B}_{r,y}(\mathcal{W})$ within $\Gamma_y(\mathcal{W})$

In some situations, the players receive sufficient information to provide incentives by observing the state process alone. Incentives are provided through lump rewards and punishments when a state transition does occur, as well as incremental punishments and rewards—typically in the opposite direction—for the lack of such state transitions. The provision of lump rewards/punishments upon a state transition is similar to the decomposition of payoffs in discrete time, hence we shall adopt the same terminology.

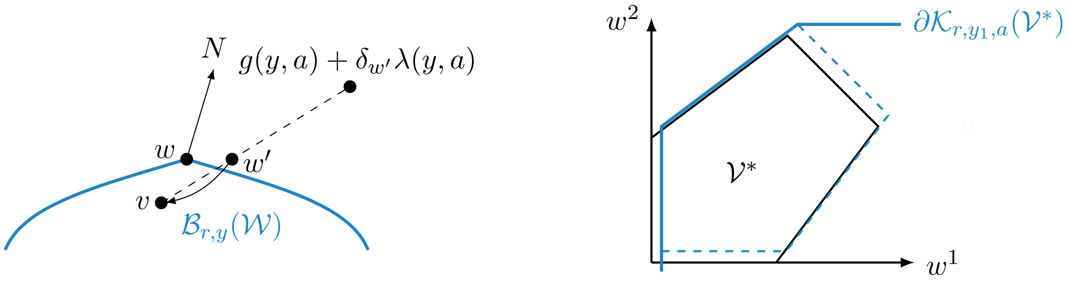

Figure 7: The left panel illustrates that it is not possible that both conditions of Lemma 6.6 hold simultaneously at $(w, N) \in \Gamma_y(\mathcal{W})$. The right panel illustrates the construction of $\mathcal{K}_{r,y_1,a}(\mathcal{V}^*)$ in the contestable democracy game for $a = (M, I)$. Extremal points in $\mathcal{K}_{r,y_1,a}(\mathcal{V}^*)$ or offset from extremal points in $\mathcal{V}^*$ precisely by the minimal jump size $r\delta_*$.

Definition 6.5. Consider $w \in \partial\mathcal{B}_{r,y}(\mathcal{W})$ and a set of action profiles $\mathcal{A}_w \subseteq \mathcal{A}(y)$.

(i) $\mathcal{A}_w$ *decomposes* w if for any $N \in \mathcal{N}_w\big(\mathcal{B}_{r,y}(\mathcal{W})\big)$, there exist $a \in \mathcal{A}_w$ and $\delta \in \Psi_{y,a}(w, \mathcal{W})$ with $N^\top\big(g(y,a) + \delta\lambda(y,a) - w\big) \geq 0$. Such w is said to be *decomposable.*

(ii) $\mathcal{A}_w$ *strictly decomposes* w if $\mathcal{A}_w$ decomposes w and for each $N \in \operatorname{ext}\mathcal{N}_w\big(\mathcal{B}_{r,y}(\mathcal{W})\big)$, there exist $a \in \mathcal{A}_w$ and $\delta \in \Psi_{y,a}(w, \mathcal{W})$ with $N^\top\big(g(y,a) + \delta\lambda(y,a) - w\big) > 0$.

(iii) $\mathcal{A}_w$ *minimally decomposes* w if $\mathcal{A}_w$ decomposes w and no proper subset of $\mathcal{A}_w$ decomposes w.

Referring back to the definition of $\Xi_{y,a}$ in Definition 6.1, we note that the defining characteristics of a decomposable payoff pair are the following: incentives are provided through state changes only, the drift rate points towards the interior of $\mathcal{B}_{r,y}(\mathcal{W})$, and the continuation payoff after a state transition to state y' is in $\mathcal{W}_{y'}$. The following lemma establishes that at most one of these defining conditions holds strictly.

Lemma 6.6. *For any* $(w, N) \in \mathcal{N}_{\mathcal{B}_{r,y}(\mathcal{W})} \cap \Gamma_y(\mathcal{W})$*, it is impossible that there exist* $a \in \mathcal{A}(y)$ *and* $\delta \in \Psi_{y,a}(w, \mathcal{W})$ *such that both of the following conditions hold:*

(i) $(0, \delta)$ *strictly enforces* a *in state* y*,*

(ii) $N^\top\big(g(y,a) + \delta\lambda(y,a) - w\big) > 0$.

The idea is that if such (w, N, a, δ) existed, then Conditions (i) and (ii) could be satisfied at any w' in a sufficiently small neighborhood of w. For a sufficiently small perturbation of incentives, it would then be possible to attain w' outside of $\mathcal{B}_{r,y}(\mathcal{W})$ with a locally enforceable strategy profile, that reaches $\mathcal{B}_{r,y}(\mathcal{W})$—or $\mathcal{W}_{y'}$ after a transition to state y'—with certainty. But then w' must lie in $\mathcal{B}_{r,y}(\mathcal{W})$, a contradiction.

Lemma 6.6 tells us that there are two kinds of boundary payoffs in $\Gamma_y(\mathcal{W})$. On the one hand, there are boundary payoffs that are not strictly decomposable, hence (w, N) has to satisfy the *abrupt-information optimality equation*

$$N^\top w = \max_{a \in \mathcal{A}} \max_{\delta \in \Psi_{y,a}(w,\mathcal{W})} N^\top\big(g(y,a) + \delta\lambda(y,a)\big). \tag{13}$$

Note that (13) is in some sense the limiting ODE of (12) as (w, N) approaches $\Gamma_y(r, \mathcal{W})$, corresponding to the case where both numerator and denominator of (12) are 0. Only for a solution to (13), does the expression for the curvature in (12) not explode for (w, N) in $\Gamma_y(r, \mathcal{W})$. On the other hand, there are strictly decomposable boundary payoffs, at which incentives for at least one player have to be binding with extremal rewards/punishments in $\mathcal{W}_{y'}$ for at least one successor state y'. Such a payoff pair thus lies on the boundary of

$$\mathcal{K}_{r,y,a}(\mathcal{W}) := \{w \mid \exists\, \delta \in \Psi_{y,a}(w, \mathcal{W})\}.$$

The boundary of the set $\mathcal{K}_{r,y,a}(\mathcal{W})$ consists of translates of $\partial\mathcal{W}_{y'}$ for successor states y' by the extremal punishments/rewards necessary to enforce a. This can be visualized particularly easily when there is a single type of events as we do in Figure 7 for the regime-change example of Section 3.

The implication of Lemma 6.6 is particularly pungent for the decomposition of corners of $\mathcal{B}_{r,y}(\mathcal{W})$, i.e., boundary points with more than one outward normal vector. Condition (ii) can be violated for all outward normal vectors only if $w = g(y, a) + \delta\lambda(y, a)$, i.e., if w is stationary. Thus, corners of $\mathcal{B}_{r,y}(\mathcal{W})$ are either stationary or in the set $\mathcal{K}_{r,y}(\mathcal{W}) := \bigcup_{a\in\mathcal{A}} \partial\mathcal{K}_{r,y,a}(\mathcal{W})$. This discussion is formalized in the following proposition.

Proposition 6.7. *Consider* $(w, N) \in \mathcal{N}_{\mathcal{B}_{r,y}(\mathcal{W})} \cap \Gamma_y(\mathcal{W})$ *with non-stationary* w. *Then exactly one of the following conditions holds:*

(i) There exists a set $\mathcal{A}_w \subseteq \mathcal{A}(y)$ *that strictly and minimally decomposes* w *such that* $w \in \partial\mathcal{K}_{r,y,a}(\mathcal{W})$ *for each* $a \in \mathcal{A}_w$ *and*

$$\mathcal{N}_w\big(\mathcal{B}_{r,y}(\mathcal{W})\big) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,y,\mathcal{A}_w}(\mathcal{W})\big), \tag{14}$$

where we denote $\mathcal{K}_{r,y,\mathcal{A}_w}(\mathcal{W}) := \bigcap_{a\in\mathcal{A}_w} \mathcal{K}_{r,y,a}(\mathcal{W})$.

(ii) (w, N) *satisfies* (13) *and either of the following conditions hold:*

(a) w *is in the interior of* $\mathcal{K}_{r,y,a_*}(\mathcal{W})$ *and* $\mathcal{N}_w\big(\mathcal{B}_{r,y}(\mathcal{W})\big) = \{N\}$,

(b) w *is decomposed by* a_* *and* $w \in \partial\mathcal{K}_{r,y,a_*}(\mathcal{W})$,

where a_* *is the action profile attaining the maximum in* (13).

Proposition 6.7 implies that within $\Gamma_y(\mathcal{W})$, the boundary of $\mathcal{B}_{r,y}(\mathcal{W})$ can have three kinds of continuously differentiable or smooth line segments and three kinds of corners. Consider a smooth segment $\mathcal{C} \subseteq \partial\mathcal{B}_{r,y}(\mathcal{W})$ with $\mathcal{N}_{\mathcal{C}} \subseteq \Gamma_y(\mathcal{W})$. Since the outward normal vector is unique, w must be minimally decomposable by a single action profile a. If a strictly decomposes w, then we are in case (i) and $\mathcal{C}$ must be contained in $\partial\mathcal{K}_{r,y,a}(\mathcal{W})$. If no action profile strictly decomposes w, then the drift rate of any enforceable strategy profile

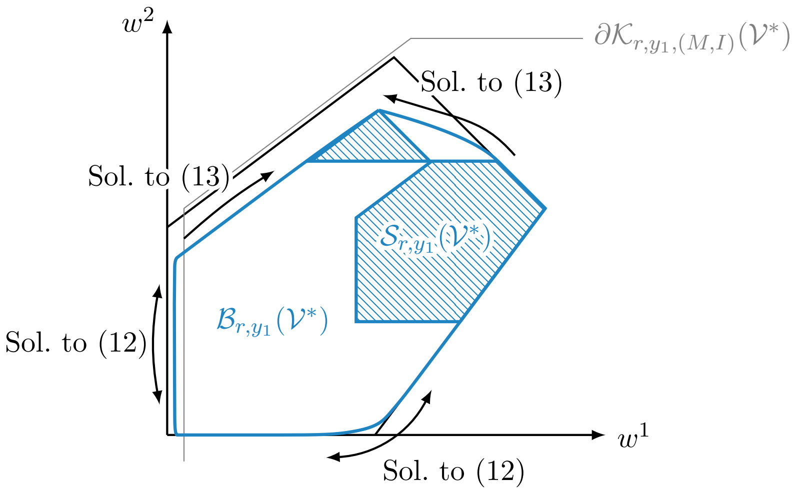

Figure 8: The set $\mathcal{B}_{r,y_1}(\mathcal{V}^*)$ of relaxed-generating payoff sets in the regime-change example after one step of the algorithm in Proposition 4.9. In this example, the boundary is a differentiable solution to (12) or (13) outside of $\mathcal{S}_{r,y_1}(\mathcal{V}^*)$. In general, it could have corners in $\mathcal{K}_{r,y_1}(\mathcal{V}^*)$.

attaining w must be parallel to the boundary. According to Proposition 6.7, $\mathcal{C}$ is thus either stationary or a solution to (13) with tangential drift as in case (ii.a). Similarly, if a corner w is strictly decomposable by $\mathcal{A}_w$, then w is a corner of $\mathcal{K}_{r,y,\mathcal{A}_w}(\mathcal{W})$ satisfying (14). Note that contrary to smooth line segments, corners are not necessarily decomposable by a single action profile. If w is not strictly decomposable, then it is either stationary or the starting point of a continuously differentiable solution to (13) as in case (ii.b).

6.3 PPE Payoff Set

Lemmas 6.4 and 6.7 motivate the following result, characterizing the family of payoff sets in each step of the algorithm in Proposition 4.9. Figure 8 illustrates the family of sets $\mathcal{B}_r(\mathcal{V}^*)$ after one step of the algorithm in the regime-change example in Section 3.

Theorem 6.8. *Let $\mathcal{W}$ be a family of convex and compact payoff sets. Then $\mathcal{B}_{r,y}(\mathcal{W})$ is the largest closed convex subset $\mathcal{X}$ of $\mathcal{V}^*$ that contains $\mathcal{S}_{r,y}(\mathcal{W})$ such that:*

(i) Outside of $\mathcal{S}_{r,y}(\mathcal{W}) \cup \mathcal{K}_{r,y}(\mathcal{W})$, the boundary $\partial\mathcal{X}$ is continuously differentiable and $(w, N_w) \in \mathcal{N}_{\mathcal{X}}$ solves (13) within $\Gamma_y(\mathcal{W})$ and (12) outside of $\Gamma_y(\mathcal{W})$,

(ii) Every corner $w \in \partial\mathcal{X}$ is either stationary, minimally decomposable by a maximizer a_* of (13) for (w, N) with $N \in \operatorname{ext}\mathcal{N}_w(\mathcal{X})$ and $w \in \partial\mathcal{K}_{r,y,a_*}(\mathcal{W})$, or minimally and strictly decomposable by some $\mathcal{A}_w \subseteq \mathcal{A}(y)$ with $\mathcal{N}_w(\mathcal{X}) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,y,\mathcal{A}_w}(\mathcal{W})\big)$.

Since $\mathcal{B}_r$ preserves compactness due to Theorem 6.8, it follows from Proposition 4.9 that $\mathcal{E}(r)$ is compact. An application of Theorem 6.8 for $\mathcal{W} = \mathcal{E}(r)$ thus provides a fixed-point characterization of $\mathcal{E}(r)$ since $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$. We refer to Section 8 in Bernard [4] for notes on the implementation of Theorem 6.8.

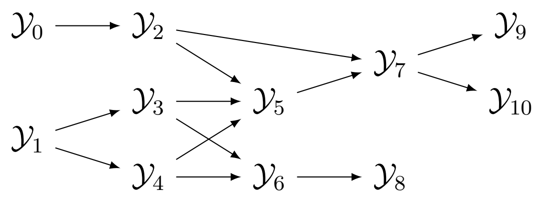

Figure 9: Communicating classes of a stochastic game form a directed acyclic graph.

Compared to repeated games, the proof of Theorem 6.8 has to deal with a more complex mathematical foundation of the model and additional arguments to show that the relevant differential equations are locally Lipschitz continuous. After establishing these, however, the construction of local incentives along the differential equations (12) and (13) are identical to Bernard [4] and, hence, omitted in this paper.

7 Computation

In the terminology of Markov processes, a set of states $\mathcal{Y}_0 \subseteq \mathcal{Y}$ is a *communicating class* if each state within $\mathcal{Y}_0$ can be reached from any other state in $\mathcal{Y}_0$ with positive probability (either directly or indirectly). Computation of the family of equilibrium payoff sets proceeds by communicating classes.

Communicating classes of a Markov process can be organized in a directed acyclic graph as illustrated in Figure 9.[12] A communicating class $\mathcal{Y}_0$ is a *direct predecessor class* of $\mathcal{Y}_1$, denoted $\mathcal{Y}_0 \prec \mathcal{Y}_1$, if some state in $\mathcal{Y}_1$ can be reached directly from some state in $\mathcal{Y}_0$. If $\mathcal{Y}_0 \prec \mathcal{Y}_1$, we also say that $\mathcal{Y}_1$ is a *direct successor class* of $\mathcal{Y}_0$. Consider a communicating class $\mathcal{Y}_e$ without direct successor class, that is, a class at the end of the directed graph. Since no states outside of $\mathcal{Y}_e$ can ever be reached, the subfamily $\big(\mathcal{E}_y(r)\big)_{y\in\mathcal{Y}_e}$ of equilibrium payoff sets can be computed from the algorithm in Proposition 4.9 without considering states in $\mathcal{Y}_e^c$. This is particularly simple if $\mathcal{Y}_e = \{y_e\}$ is a singleton, that is, y_e is an absorbing state. The continuation game is then just a repeated game and hence $\mathcal{E}_{y_e}(r)$ can be computed with Theorem 2 in Sannikov [29] if all enforceable action profiles are pairwise identifiable or with Theorem 6.8 in Bernard [4] otherwise. One can then proceed backwards in the directed graph: consider a communicating class $\mathcal{Y}'$, for which all subfamilies of equilibrium payoff sets of direct successor classes $\mathcal{Y}_{e_1}, \ldots, \mathcal{Y}_{e_n}$ have been computed already. In the computation of $\big(\mathcal{E}_y(r)\big)_{y\in\mathcal{Y}'}$, incentives from state transitions to states in $\mathcal{Y}_E := \bigcup_{k=1}^n \mathcal{Y}_{e_k}$ do not need to computed iteratively as in Proposition 4.9, but only incentives via state transitions within the communicating class $\mathcal{Y}'$ have to be accounted for in an iterative fashion. This becomes again particularly simple if $\mathcal{Y}' = \{y'\}$ is a singleton. Then $\mathcal{E}_{y'}(r) = \mathcal{B}_{r,y'}\big((\mathcal{E}_y(r))_{y\in\mathcal{Y}_E}\big)$ can be solved in a single application of Theorem 6.8 rather than an iterated application.

[12] Indeed, if there was a cycle of communicating classes, then every state within the cycle could be reached from any other state in the cycle, hence the union of all communicating classes in the cycle forms a single communicating class.

8 Conclusion

Based on recent developments in continuous-time repeated games, this paper provides a unifying framework for the analysis of stochastic games with imperfect public monitoring in a continuous-time setting. The methodology is not limited to irreducible games or absorbing games and it is applicable to any stochastic game, as long as the public signal satisfies Assumptions 1 and 2. The paper characterizes the set of all PPE payoffs, the set of all stationary Markov payoffs, as well as the payoffs of a notion of semi-stationary PPE that are simple to compute, yet not as limiting as stationary Markov equilibria.

The analysis relies on an iterative procedure over state transitions, allowing us to draw from the techniques of continuous-time repeated games in each step of the iteration. This is possible because we consider stochastic games with finitely many states and, hence, the state is constant almost everywhere. The techniques in the present paper cannot be extended in a straightforward manner to stochastic games with a continuum of states. Such an extension provides an interesting direction for future research.

References

[1] D. Abreu, D. Pearce, and E. Stacchetti: Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica*, **58** (1990), 1041–1063

[2] D. Acemoglu and J.A. Robinson: A theory of political transitions, *American Economic Review*, **91** (2001), 938–963

[3] D. Bergemann and J. Välimäki: Learning and strategic pricing, *Econometrica*, **64** (1996), 1125–1149

[4] B. Bernard: Continuous-time games with imperfect and abrupt information, 2018, revise and resubmit at *Review of Economic Studies*

[5] B.D. Bernheim and D. Ray: Markov perfect equilibria in altruistic growth economies with production uncertainty, *Journal of Economic Theory*, **47** (1989), 195–202

[6] U. Doraszelski and J.F. Escobar: A theory of regular Markov perfect equilibria in dynamic stochastic games: genericity, stability, and purification, *Theoretical Economics*, **5** (2010), 369–402

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

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

[9] A.M. Fink: Equilibrium in a stochastic n-person game, *Journal of Science of the Hiroshima University*, **28** (1964), 89–93

[10] D. Fudenberg, and D. K. Levine: Efficiency and observability with long-run and short-run players, *Journal of Economic Theory*, **62** (1994), 103–135

[11] D. Fudenberg, D. K. Levine, and E. Maskin: The folk theorem with imperfect public information, *Econometrica*, **62** (1994), 997–1039

[12] H. Haller and R. Lagunoff: Genericity and Markovian behavior in stochastic games, *Econometrica*, **68** (2000), 1231–1248

[13] W. He and Y. Sun: Stationary Markov perfect equilibria in discounted stochastic games, *Journal of Economic Theory*, **169** (2017), 35–61

[14] J. Hörner, T. Sugaya, S. Takahashi, and N. Vieille: Recursive methods in discounted stochastic games: an algorithm for $\delta \to 1$ and a folk theorem, *Econometrica*, **79** (2011), 1277–1318

[15] J. Jacod, A.N. Shiryaev: *Limit Theorems for Stochastic Processes*, 2nd edition, Springer, New York, 2002

[16] Kenneth L. Judd, Sevin Yeltekin, and James Conklin: Computing supergame equilibria, *Econometrica*, **71** (2003), 1239–1254

[17] O. Kallenberg: *Foundations of Modern Probability*, Springer, New York, 1997

[18] I. Karatzas and S. Shreve: *Methods of Mathematical Finance*, Applications of Mathematics 39, Springer, New York, 1998

[19] N. Kazamaki: *Continuous Exponential Martingales and BMO*, Lecture Notes in Mathematics 1579, Springer, Berlin, 1994

[20] G. Keller, S. Rady, and M. Cripps: Strategic experimentation with exponential bandits, *Econometrica*, **73** (2005), 39–68

[21] A.P. Kirman and M.J. Sobel: Dynamic oligopoly with inventories, *Econometrica*, **42** (1974), 279–287

[22] D. Levhari and L. Mirman: The great fish war: an example using a dynamic Cournot-Nash solution, *The Bell Journal of Economics*, **11** (1980), 322–334

[23] J. Levy: A discounted stochastic game with no stationary Nash equilibrium: two examples, *Econometrica*, **81** (2013), 1973–2003

[24] J. Levy, A. McLennan: Corrigendum to "Discounted stochastic games with no stationary Nash equilibrium: two examples", *Econometrica*, **83** (2015), 1237–1252

[25] J.R. Munkres: *Topology*, 2nd ed., Prentice Hall, 2000

[26] A. Neyman: Continuous-time stochastic games, *Games and Economic Behavior*, **104** (2017), 92–130

[27] M. Peski and J. Toikka: Value of persistent information, *Econometrica*, **85** (2017), 1921–1948

[28] P. Protter: *Stochastic Integration and Differential Equations*, 2nd ed., Stochastic Modelling and Applied Probability 21, Springer, New York, 2005

[29] Y. Sannikov: Games with imperfectly observable actions in continuous time, *Econometrica*, **75** (2007), 1285–1329

[30] L.K. Simon and M.B. Stinchcombe: Extensive form games in continuous time: pure strategies, *Econometrica*, **57** (1989), 1171–1214

[31] S. Yeltekin, Y. Cai, K.L. Judd: Computing equilibria of dynamic games, *Operations Research*, **65** (2017), 337–356

[32] N.D. Yen: Lipschitz continuity of solutions of variational inequalities with parametric polyhedral constraints, *Mathematics of Operations Research*, **20** (1995), 695–708

A Mathematical Foundation

A mathematical foundation for the model in Section 2 can be provided as follows. Let $\mathcal{Z}$ denote the set of ordered pairs (y, y') of states, where $\lambda_{y,y'}(a) > 0$ for some action profile $a \in \mathcal{A}(y)$. Let $(\Omega, \mathcal{F}, P)$ be a probability space containing a Brownian motion Z and a Poisson process $J^{y,y'}$ with intensity 1 for each $(y, y') \in \mathcal{Z}$ such that $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ and Z are mutually independent. The state process S is then defined by

$$S_0 = y_0, \qquad S_t := \begin{cases} S_{t-} & \text{if } \Delta J_t^{S_{t-},y} = 0 \text{ for all } y \text{ with } (S_{t-}, y) \in \mathcal{Z}, \\ y & \text{if } \Delta J_t^{S_{t-},y} = 1, \end{cases} \tag{15}$$

i.e., S is a piecewise constant stochastic process that jumps to y at time t if and only if $\Delta J_t^{S_{t-},y} = 1$. When we require existence of a state process S with certain properties in Section 2, the mathematically precise statement is to require existence of independent Poisson processes $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ such that the induced process S has the desired properties. The family $Q^A = \left(Q_t^A\right)_{t\geq 0}$ of probability measures is defined via its density process

$$\frac{\mathrm{d}Q_t^A}{\mathrm{d}P} := \exp\left(\int_0^t \mu(S_s, A_s)\,\mathrm{d}Z_s - \int_0^t \left(\frac{1}{2}|\mu(S_s, A_s)|^2 + \sum_{(S_{s-},y)\in\mathcal{Z}(A_{s-})} \lambda_{S_{s-},y}(A_{s-}) - 1\right)\mathrm{d}s\right)$$
$$\prod_{\substack{0<s\leq t\\(S_{s-},y)\in\mathcal{Z}(A_{s-})}} \left(1 + (\lambda_{S_{s-},y}(A_{s-}) - 1)\Delta J_s^{S_{s-},y}\right). \tag{16}$$

Formally, the filtration $\mathbb{F}$ contains the filtration generated by Z and $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$, and not just the filtration generated by Z and S so that the density process is adapted to $\mathbb{F}$. Note that the instantaneous intensities of the processes $J^{y,y'}$ are equal to 1 under Q^A, where $S \neq y$ so that players learn nothing from these processes. A mathematical foundation based on processes $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ ensures that each Q_t^A is absolutely continuous with respect to some reference measure P. It follows from Girsanov's theorem (e.g., Theorem III.3.11 in Jacod and Shiryaev [15]) that the public signal under Q^A indeed takes the form (1) and that state transitions occur with instantaneous intensities $\lambda_{S_{t-},y}(A_t)$.

A.1 Continuation Value of PPE

The proofs in this appendix are adaptations of the proofs in Bernard [4] to the setting of stochastic games. Because the public signal and the state process generate the same information as the public signal in Bernard [4], those adaptations are minor.

Proof of Lemma 4.1. Fix a strategy profile A first and observe that $W := W(S, A)$ is bounded as it remains in $\mathcal{V}$ at all times. Fix a player i and a time $T > 0$, and define the bounded $\mathcal{F}_T$-measurable random variable $w_T^i := W_T^i - r\int_0^T \big(W_t^i - g^i(S_t, A_t)\big)\,\mathrm{d}t$. Because $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ are pairwise orthogonal and orthogonal to Z, the stable subspace generated by Z and $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ is the space of all stochastic integrals with respect to these processes by Theorem IV.36 in Protter [28]. Therefore, we obtain the unique martingale representation property for a square-integrable martingale by Corollary 1 to Theorem IV.37 in [28]. That is, for a bounded $\mathcal{F}_T$-measurable random variable w_T^i, there exist an $\mathcal{F}_0$-measurable c_T^i, predictable and square-integrable processes $\big(\beta_{t,T}^i\big)_{0\le t\le T}$ and $\big(\delta_{t,T}^i(y,y')\big)_{0\le t\le T}$ for $(y,y')\in\mathcal{Z}$ with $\delta_{t,T}^i(y,y') = 0$ for $(y,y')\notin\mathcal{Z}(A)$, $\mathbb{E}_{Q_T^A}\Big[\int_0^T \big|\beta_{t,T}^i\big|^2\,\mathrm{d}t\Big] < \infty$ and $\mathbb{E}_{Q_T^A}\Big[\int_0^T \big|\delta_{t,T}^i(y,y')\big|^2 \lambda_t^{y,y'}\,\mathrm{d}t\Big] < \infty$ for all $(y,y')\in\mathcal{Z}$ and a Q_T^A-martingale M^i orthogonal to Z and $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$ with $M_0^i = 0$ such that

$$w_T^i = c_T^i + r\int_0^T \beta_{t,T}^i(\mathrm{d}Z_t - \mu(S_t,A_t)\,\mathrm{d}t) + \sum_{(y,y')\in\mathcal{Z}} r\int_0^T \delta_{t,T}^i(y,y')(\mathrm{d}J_t^{y,y'} - \lambda_t^{y,y'}\,\mathrm{d}t) + M_{T,T}^i.$$

To prove that (5) holds, we need to show that c_T^i, $\beta_{t,T}^i$, $\delta_{t,T}^i(y,y')$ and $M_{t,T}^i$ do not depend on T. It follows from (2) and Fubini's theorem that

$$w_T^i = W_T^i + r\int_0^T g^i(S_t,A_t)\,\mathrm{d}t - r\int_0^\infty \int_0^{s\wedge T} r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q_s^A}\big[g^i(S_s,A_s)\;\big|\;\mathcal{F}_t\big]\,\mathrm{d}t\,\mathrm{d}s. \tag{17}$$

Let $\tilde{T} \le T$ and take conditional expectations on $\mathcal{F}_{\tilde{T}}$ under Q_T^A of (17) to deduce that

$$\begin{aligned}
\mathbb{E}_{Q_T^A}\big[w_T^i\,\big|\,\mathcal{F}_{\tilde{T}}\big] - w_{\tilde{T}}^i &= \mathbb{E}_{Q_T^A}\big[W_T^i\,\big|\,\mathcal{F}_{\tilde{T}}\big] - W_{\tilde{T}}^i + r\int_{\tilde{T}}^T \mathbb{E}_{Q_t^A}\big[g^i(S_t,A_t)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}t \\
&\quad - r\int_{\tilde{T}}^\infty \int_{\tilde{T}}^{s\wedge T} r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q_s^A}\big[g^i(S_s,A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}t\,\mathrm{d}s \\
&= \mathbb{E}_{Q_T^A}\big[W_T^i\,\big|\,\mathcal{F}_{\tilde{T}}\big] - W_{\tilde{T}}^i - \int_T^\infty r\mathrm{e}^{-r(s-T)}\mathbb{E}_{Q_s^A}\big[g^i(S_s,A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}s \\
&\quad + \int_{\tilde{T}}^\infty r\mathrm{e}^{-r(s-\tilde{T})}\mathbb{E}_{Q_s^A}\big[g^i(S_s,A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}s \\
&= 0.
\end{aligned}$$

Taking $\tilde{T} = 0$, this shows that $c^i_T = W^i_0$ does not depend on T. It also implies

$$w^i_{\tilde{T}} = W^i_0 + r\int_0^{\tilde{T}} \beta^i_{t,T}\big(\mathrm{d}Z_t - \mu(S_t, A_t)\,\mathrm{d}t\big) + \sum_{(y,y')\in\mathcal{Z}} r\int_0^{\tilde{T}} \delta^i_{t,T}(y,y')\big(\mathrm{d}J^{y,y'}_t - \lambda^{y,y'}_t\,\mathrm{d}t\big) + M^i_{\tilde{T},T},$$

which yields $\beta^i_{\cdot,T} = \beta^i_{\cdot,\tilde{T}}$ and $\delta^i_{\cdot,T}(y,y') = \delta^i_{\cdot,\tilde{T}}(y,y')$ for every $(y,y') \in \mathcal{Z}$ a.e. on $[0,\tilde{T}]$ and $M^i_{\tilde{T},T} = M^i_{\tilde{T},\tilde{T}}$ a.s. by the uniqueness of the orthogonal decomposition. Taking $\mathcal{F}_t$-conditional expectations, we deduce $M^i_{t,\tilde{T}} = M^i_{t,T}$ a.s. for $t \in [0,\tilde{T}]$, proving that the integral representation is independent of T. We thus omit the subscript T and $\tilde{T}$ of the constructed processes β^i, $\big(\delta^i(y,y')\big)_{(y,y')\in\mathcal{Z}}$, and M^i. To arrive at (5), we set

$$\tilde{M}^i = M^i + \sum_{(y,y')\in\mathcal{Z}} r\int_0^{\cdot} \delta^i_t(y,y')1_{\{S_{t-}\neq y\}}\big(\mathrm{d}J^{y,y'}_t - 1\,\mathrm{d}t\big)$$

and $\delta^i(y) := \delta^i(S_-, y)1_{\{(S_-,y)\in\mathcal{Z}\}}$ for any $y \in \mathcal{Y}$. Because M^i is orthogonal of Z and $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$, and the processes $J^{y,y'}$ are independent of each other and of Z, it follows that $\tilde{M}^i$ is a martingale orthogonal to Z and to $J^y = \sum_{0<s\leq t}\Delta J^{S_{s-},y}_t$ for every $y \in \mathcal{Y}$. Since the processes $\delta^i(y,y')$ are square-integrable by construction, so is $\delta^i(y)$. Moreover, $\delta^i(y)$ is predictable because $\delta^i(y,y')$ and S_- are. By construction,

$$\tilde{M}^i + \sum_{y\in\mathcal{Y}} r\int_0^{\cdot} \delta^i(y)\big(\mathrm{d}J^y_t - \lambda_{S_{t-},y}(A_t)\,\mathrm{d}t\big) = M^i + \sum_{(y,y')\in\mathcal{Z}} r\int_0^{\tilde{T}} \delta^i_{t,T}(y,y')\big(\mathrm{d}J^{y,y'}_t - \lambda^{y,y'}_t\,\mathrm{d}t\big),$$

hence W satisfies (5) for processes β^i, $\big(\delta^i(y)\big)_{y\in\mathcal{Y}}$, and $\tilde{M}^i$.

To show the converse, we derive from Itō's formula that

$$\begin{aligned}\mathrm{d}(\mathrm{e}^{-rt}W^i_t) &= -r\mathrm{e}^{-rt}g^i(S_t, A_t)\,\mathrm{d}t + r\mathrm{e}^{-rt}\beta^i_t\left(\mathrm{d}Z_t - \mu(S_t,A_t)\,\mathrm{d}t\right) \\ &\quad + r\mathrm{e}^{-rt}\sum_{y\in\mathcal{Y}}\delta^i_t(y)\big(\mathrm{d}J^y_t - \lambda_{S_{t-},y}(A_t)\,\mathrm{d}t\big) + \mathrm{e}^{-rt}\,\mathrm{d}M^i_t. \end{aligned} \tag{18}$$

Since M^i is strongly orthogonal to Z and $(J^y)_{y\in\mathcal{Y}}$, it is also strongly orthogonal to the density process given in (16). Therefore, M^i is a martingale also under Q^A. Integrating (18) from t to T and taking Q^A_T-conditional expectations on $\mathcal{F}_t$ thus yields

$$W^i_t = \int_t^T r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q^A_s}\left[g^i(S_s, A_s) \mid \mathcal{F}_t\right]\mathrm{d}s + \mathrm{e}^{-r(T-t)}\mathbb{E}_{Q^A_T}\left[W^i_T \mid \mathcal{F}_t\right].$$

Since W is bounded, the second summand converges to zero a.s. as T tends to ∞, hence W^i_t is indeed i's continuation value under strategy profile A in S_t. □

Proof of Lemma 4.3. Fix a strategy profile A and let $\tilde{A}$ be a strategy profile involving a unilateral deviation of some player i, that is, $\tilde{A}^{-i} = A^{-i}$ a.e. For (β, δ) related to $W(S, A)$ by (5), integrating (18) from t to u yields

$$W_t^i(S_t, A) = -\int_t^u e^{-r(s-t)}\Big(\beta_s^i \left(dZ_s - \mu(S_s, A_s)\, ds\right) - g^i(S_s, A_s)\, ds - dM_s^i\Big)$$
$$- \sum_{y\in\mathcal{Y}} \int_t^u e^{-r(s-t)}\delta_s^i(y)\left(dJ_s^y - \lambda_{S_{s-},y}(A_s)\, ds\right) + e^{-r(u-t)}W_u^i(S_u, A).$$

Note that the term $e^{-r(u-t)}W_u^i(S_u, A)$ vanishes as we let $u \to \infty$ because $W(S, A)$ remains in the bounded set $\mathcal{V}$. Since M is a martingale up to time u also under $Q_u^{\tilde{A}}$, taking conditional expectations yields

$$W_t^i\big(S_t, \tilde{A}\big) = \lim_{u\to\infty} \mathbb{E}_{Q_u^{\tilde{A}}}\left[\int_t^u re^{-r(s-t)}g^i\big(S_s, \tilde{A}_s\big)\, ds \,\middle|\, \mathcal{F}_t\right]$$
$$= W_t^i(S_t, A) + \lim_{u\to\infty} \mathbb{E}_{Q_u^{\tilde{A}}}\left[\int_t^u re^{-r(s-t)}\Big(\big(g^i\big(S_s, \tilde{A}_s\big) - g^i(S_s, A_s)\big)\, ds\right.$$
$$\left. + \beta_s^i\left(dZ_s - \mu(S_s, A_s)\, ds\right) + \sum_{y\in\mathcal{Y}}\delta_s^i(y)\left(dJ_s^y - \lambda_{S_{s-},y}(A_s)\, ds\right)\Big) \,\middle|\, \mathcal{F}_t\right] \quad \text{a.s.}$$

Note here that the state process S is the same process under strategy profile A and $\tilde{A}$, but it has a different distribution under Q_u^A and $Q_u^{\tilde{A}}$. Because β is constructed using a martingale representation result for the bounded random variable w_T^i in (17), the process $\int_t^\cdot re^{-r(s-t)}\beta_s^i\left(dZ_s - \mu(S_s, A_s)\, ds\right)$ is a bounded mean oscillation (BMO) martingale under the probability measure Q_u^A up to any time $u > t$. It follows from Theorem 3.6 in Kazamaki [19] that $\int_t^\cdot re^{-r(s-t)}\beta_s^i\left(dZ_s - \mu(S_s, \tilde{A}_s)\, ds\right)$ is a martingale under $Q_u^{\tilde{A}}$. Since $W(A)$ lies in $\mathcal{V}$ $P \otimes$ *Lebsgue*-a.e., it follows from Lemma 4.1 that $\|\delta(y)\| \leq |V|/r$ $P \otimes$ *Lebsgue*-a.e. In particular, each $\delta(y)$ is uniformly bounded P-a.s. and hence also $Q_u^{\tilde{A}}$-a.s. for any $u > t$. The lemma after Theorem IV.29 in Protter [28] thus implies that $\int_t^\cdot re^{-r(s-t)}\delta_s^i(y)\left(dJ_s^y - \lambda_{S_{s-},y}(\tilde{A}_s)\, ds\right)$ is a $Q_u^{\tilde{A}}$ up to any time $u > t$. Together with Fubini's theorem, this implies

$$W_t^i\big(S_t, \tilde{A}\big) - W_t^i(S_t, A) = \int_t^\infty e^{-r(s-t)}\mathbb{E}_{Q_s^{\tilde{A}}}\Big[g^i\big(S_s, \tilde{A}_s\big) - g^i(S_s, A_s) + \beta_s^i\left(\mu\big(S_s, \tilde{A}_s\big)\right. \tag{19}$$
$$\left. - \mu(S_s, A_s)\right) + \delta_s^i\left(\lambda\big(S_s, \tilde{A}_s\big) - \lambda(S_s, A_s)\right) \,\Big|\, \mathcal{F}_t\Big]\, ds \quad \text{a.s.}$$

If (β, δ) enforces A in state S, the above conditional expectation is non-positive, hence A is a PPE. To show the converse, assume towards a contradiction that there exist a player i

and a set $\Xi \subseteq \Omega \times [0,\infty)$ with $P \otimes Lebesgue(\Xi) > 0$, such that for some strategy $\hat{A}^i$

$$g^i\big(S, \hat{A}\big) - g^i(S, A) + \beta^i\left(\mu\big(S, \hat{A}\big) - \mu(S, A)\right) + \delta^i\left(\lambda\big(S, \hat{A}\big) - \lambda(S, A)\right) > 0$$

on the set Ξ, where we denoted $\hat{A} = \big(\hat{A}^i, A^{-i}\big)$ for the sake of brevity. Because β and δ are predictable, we can and do choose Ξ predictable as well. Thus, $\tilde{A}^i := \hat{A}^i 1_{\Xi} + A^i 1_{\Xi^c}$ is predictable and, in particular, a strategy for player i. For $\tilde{A} = \big(\tilde{A}^i, A^{-i}\big)$, the expectation in (19) is strictly positive for $t = 0$, which is a contradiction. □

A.2 Self-Generation

The main idea behind the proof of stochastic self generation in Proposition 4.6 is the following: Any payoff in $\mathcal{E}(r)$ can be attained by an enforceable solution W to (5) by Lemma 4.3. Thus, for any stopping time τ, W_τ is trivially attained by an enforceable solution to (5). We would then like to conclude that, by virtue of Lemma 4.3, W_τ must lie in $\mathcal{E}_{S_\tau}(r)$. This last conclusion, however, is subject to some subtle measurability issues. Without restrictions on β and δ, solutions to (5) are weak solutions, that is, the public signal Z, the Poisson processes $(J^{y,y'})_{(y,y')\in\mathcal{Z}}$, and the probability space are part of the solution. Thus, the equivalence in Lemma 4.3 can be stated as

$$\mathcal{E}_{y_0}(r) = \left\{ w \in \mathcal{V} \;\middle|\; \begin{array}{l} \text{There exists a } \textit{stochastic framework } (\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}}) \\ \text{containing a solution } (W, A, \beta, \delta, M) \text{ to (5) for } S \text{ defined in (15)} \\ \text{such that } W_0 = w \ P\text{-a.s. and } (\beta, \delta) \text{ enforces } A \text{ in state } S. \end{array} \right\},$$

In general, the probability space may thus depend on the payoff that is being attained. To conclude that the random variable W_τ is in $\mathcal{E}_{S_\tau}(r)$, we construct a regular conditional probability that allows us to aggregate the different probability spaces. This is possible because the probability spaces for different realizations of W_τ share the same path space.

Lemma A.1. *For an $\mathcal{F}_0$-measurable random variable W^* in a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}}\big)$ the following are equivalent:*

(a) $W^* \in \mathcal{E}_{y_0}(r)$ P-a.s.,

(b) There exist a solution (W, A, β, δ, M) to (5) in the given stochastic framework for S defined in (15) such that $W_0 = W^*$ P-a.s. and (β, δ) enforces A in state S.

Proof. Let first $W^* \in \mathcal{E}_{y_0}(r)$ a.s. for initial state $y_0 \in \mathcal{Y}$. Although a PPE may exist on different probability spaces in (a) for each realization $W^* = w$, the path space in each w for the strategy profile A and its stochastic framework is given by $\mathcal{A}^{[0,\infty)} \times \mathcal{D}^{d+|\mathcal{Z}|}$, where $\mathcal{D}^{d+|\mathcal{Z}|}$ is the space of càdlàg functions from $[0,\infty)$ into $\mathbb{R}^{d+|\mathcal{Z}|}$. By Theorem A.2.2 in Kallenberg [17], there exists a metric on $\mathcal{D}^{d+|\mathcal{Z}|}$ that induces the Skorohod topology, under which $\mathcal{D}^{d+|\mathcal{Z}|}$ is complete and separable. Since $\mathcal{A}^{[0,\infty)}$ is compact by Tychonoff's theorem (see Theorem 37.3 in Munkres [25]), it follows that $\Omega = \mathcal{V} \times \mathcal{A}^{[0,\infty)} \times \mathcal{D}^{d+|\mathcal{Z}|}$ is complete and separable. Thus, by Theorem V.3.19 in Karatzas and Shreve [18], there exists a regular conditional probability $(w, F) \mapsto P_w(F)$ for $(w, F) \in \mathcal{V} \times \mathcal{F}$, that is,

(i) for each $w \in \mathcal{V}$, P_w is a probability measure on $(\Omega, \mathcal{F})$,

(ii) for each $F \in \mathcal{F}$, the mapping $w \mapsto P_w(F)$ is Borel($\mathcal{V}$)-measurable,

(iii) $P_w(F) = P(F \,|\, W^* = w)$ for each $F \in \mathcal{F}$ and ν-a.e. $w \in \mathcal{V}$, where ν is the distribution of W^*.

By Lemma 4.3, for each $w \in \mathcal{E}_{y_0}(r)$ there exists a solution $(W^w, A^w, \beta^w, \delta^w, M^w)$ to (5) for S^w defined from the stochastic framework and A^w via (15) with $W_0^w = w$ and (β^w, δ^w) enforcing A^w in S^w. Define the processes (W, A, β, δ, M) pointwise as $(W^w, A^w, \beta^w, \delta^w, M^w)$ on $\{W^* = w\}$ for each $w \in \mathcal{V}$. Let Ξ denote the set on which (W, A, β, δ, M) satisfies (5) such that $W_0 = W^*$ and (β, δ) enforces A in state S. It follows from the properties of a regular conditional probability that

$$P(\Xi) = \int_{\mathcal{E}_{y_0}(r)} P(\Xi \,|\, W^* = w) \,\mathrm{d}\nu(w) = 1.$$

To show the converse, let (W, A, β, δ, M) be a solution to (5) for S defined in (5) such that $W_0 = W^*$ a.s. and (β, δ) enforces A in S. Suppose that $W^* \notin \mathcal{E}_{y_0}(r)$ on an $\mathcal{F}_0$-measurable set Ξ with $\nu(\Xi) > 0$. For each $w \notin \mathcal{E}_{y_0}(r)$, there exists a player with a profitable deviation and hence, by Lemma 4.3, there exists a set $\tilde{\Xi}(w)$ with positive P_w-measure, on which (β, δ) does not enforce A in state S. We thus obtain a contradiction via

$$P((\beta, \delta) \text{ enforces } A) \le \nu\big(\mathcal{E}_{y_0}(r)\big) + \int_{\mathcal{E}_{y_0}(r)^c} P_w((\beta, \delta) \text{ enforces } A) \,\mathrm{d}\nu(w) < 1. \qquad \square$$

Proof of Lemma 4.6. Consider any family $(\mathcal{W}_y)_{y \in \mathcal{Y}}$ of bounded stochastic self-generating sets. For any state $y \in \mathcal{Y}$, any payoff pair $w \in \mathcal{W}_y$ can be attained by an enforceable solution to (5) with initial state y. It follows that $w \in \mathcal{E}_y(r)$ by Lemma 4.3, i.e., any family of self-generating sets is contained in $\mathcal{E}(r)$. Since $\mathcal{E}_y(r) \subseteq \mathcal{V}$ is bounded for each $y \in \mathcal{Y}$, it remains to show that $\mathcal{E}(r)$ is self-generating. Take $w \in \mathcal{E}_y(r)$ so that Lemma 4.1 yields the existence of a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y') \in \mathcal{Z}}\big)$ and a solution (W, A, β, δ, M) to (5) for S defined in (5) such that $W_0 = w$ a.s. and (β, δ) enforces A in state S. We now fix a stopping time τ and show that $W_\tau \in \mathcal{E}_{S_\tau}(r)$ a.s. To do so, we set $W^* = W_\tau$, $\tilde{\mathcal{F}}_t = \mathcal{F}_{\tau+t}$, $\tilde{Z}_t = Z_{\tau+t} - Z_\tau$, $\tilde{J}_t^{y,y'} := J_{\tau+t}^{y,y'} - J_\tau^{y,y'}$ for every $(y, y') \in \mathcal{Z}$, $\tilde{M}_t = M_{\tau+t} - M_\tau$, $\tilde{W}_t = W_{\tau+t}$, $\tilde{\beta}_t = \beta_{\tau+t}$, $\tilde{\delta}_t = \delta_{\tau+t}$, $\tilde{A}_t = A_{\tau+t}$, and $\tilde{S}_t = S_{\tau+t}$. Because the tilde-processes and -filtrations satisfy condition (b) in Lemma A.1 for $W^* = W_\tau$, we obtain by the equivalence in Lemma A.1 that $W_\tau = W^* \in \mathcal{E}_{\tilde{S}_0}(r) = \mathcal{E}_{S_\tau}(r)$ a.s. $\square$

B Concatenations of Solutions and Convergence of Algorithm

In this appendix we prove the convergence of the algorithm in Proposition 4.9 to $\mathcal{E}(r)$. We first show that $\mathcal{B}_r(\mathcal{W})$ is monotone in $\mathcal{W}$.

Lemma B.1. *Let $\mathcal{W}_y \subseteq \mathcal{W}'_y$ for every $y \in \mathcal{Y}$. Then $\mathcal{B}_{r,y}(\mathcal{W}) \subseteq \mathcal{B}_{r,y}(\mathcal{W}')$ for every $y \in \mathcal{Y}$.*

Proof. Fix any state $y \in \mathcal{Y}$ and any payoff pair $w \in \mathcal{B}_{r,y}(\mathcal{W})$. By definition, there exists a solution (W, A, β, δ, M) to (5) for S defined in (15) with initial state y such that $W_0 = w$ a.s., $W_\sigma \in \mathcal{W}_{S_\sigma}$ a.s., and on $[\![0, \sigma)\!)$, (β, δ) enforces A in state y and $W \in \mathcal{B}_{r,y}(\mathcal{W})$.[13] Since this implies that also $W_\sigma \in \mathcal{W}'_{S_\sigma}$ a.s., it follows that $\mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}'$-relaxed self-generating. In particular, $\mathcal{B}_{r,y}(\mathcal{W}) \subseteq \mathcal{B}_{r,y}(\mathcal{W}')$ by maximality of $\mathcal{B}_r(\mathcal{W}')$. □

For the proof of Lemma 4.8 it will be necessary to concatenate enforceable solutions to (5) at transition times of the state process. Concatenating enforceable solutions to (5) will be an essential tool also in the proofs of our other results. For ease of later use, we state the concatenation procedure in a separate lemma.

Lemma B.2. *Fix a family $\mathcal{W} = (\mathcal{W}_y)_{y \in \mathcal{Y}}$ of payoff sets. Consider a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y') \in \mathcal{Z}}\big)$ containing a solution (W, A, β, δ, M) to (5) for S defined in (15). Fix an $\mathbb{F}$-stopping time $\tau \le \sigma$ P-a.s., where σ is the time of the first state transition. Let $\hat{\sigma}$ denote the first state transition strictly after τ. Suppose that (β, δ) enforces A in state S_0 on $[\![0, \tau)\!)$ and $W_\tau \in \mathcal{B}_{r,S_\tau}(\mathcal{W})$. Then there exists a solution $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ to (5) on the same stochastic framework such that $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ coincides with (W, A, β, δ, M) on $[\![0, \tau)\!)$, (β, δ) enforces A in state S on $[\![0, \hat{\sigma})\!)$, $W \in \mathcal{B}_{r,\mathcal{Y}_\tau}(\mathcal{W})$ on $[\![\tau, \hat{\sigma})\!)$ and $W_{\hat{\sigma}} \in \mathcal{W}_{S_{\hat{\sigma}}}$ P-a.s.*

Concatenating solutions requires an analogous result to Lemma A.1 for a family of relaxed self-generating payoff sets. The proof works analogously and is omitted.

Lemma B.3. *For an $\mathcal{F}_0$-measurable random variable W^* in a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y') \in \mathcal{Z}}\big)$ the following are equivalent:*

(a) $W^* \in \mathcal{B}_{r,y_0}(\mathcal{W})$ P-a.s.,

(b) There exist a solution (W, A, β, δ, M) to (5) in the given stochastic framework for S defined in (15) such that $W_0 = W^*$ P-a.s., $W_\sigma \in \mathcal{W}_{S_\sigma}$ P-a.s., and on $[\![0, \sigma)\!)$, continuation promise (β, δ) enforces A in state S_0 and $W \in \mathcal{B}_{r,y_0}(\mathcal{W})$.

[13] For two stopping times σ and τ, the set $[\![\sigma, \tau)\!) := \{(\omega, t) \in \Omega \times [0, \infty) \mid \sigma(\omega) \le t < \tau(\omega)\}$ is called the (left-closed, right-open) stochastic interval from σ to τ. Closed, open, and left-open, right-closed stochastic intervals are defined analogously.

Proof of Lemma B.2. Fix a stochastic framework, $\left(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}}\right)$, a solution (W, A, β, δ, M) to (5), and stopping times, τ, σ, and $\hat{\sigma}$. Define the processes

$$\tilde{Z} := Z_{\cdot+\tau} - Z_\tau, \qquad \tilde{J}^{y,y'} := J^{y,y'}_{\cdot+\tau} - J^{y,y'}_\tau, \text{ for every } (y,y') \in \mathcal{Z}$$

and the filtration $\tilde{\mathbb{F}} := \left(\tilde{\mathcal{F}}_t\right)_{t\geq 0}$ defined by $\tilde{\mathcal{F}}_t := \mathcal{F}_{t+\tau}$. Because Brownian motion and Poisson processes have independent and identically distributed increments, $\tilde{Z}$ is an $\tilde{\mathbb{F}}$-Brownian motion and $\tilde{J}^{y,y'}$ for any $(y,y') \in \mathcal{Z}$ is an $\tilde{\mathbb{F}}$-Poisson process. Therefore, $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde{Z}, (\tilde{J}^{y,y'})_{(y,y')\in\mathcal{Z}})$ is a stochastic framework with $W_\tau \in \tilde{\mathcal{F}}_0$. Since $\mathcal{W}_\tau \in \mathcal{B}_{r,S_\tau}(\mathcal{W})$ by assumption, the equivalence in Lemma B.3 implies the existence of a solution $\left(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}, \tilde{M}\right)$ to (5) on the stochastic framework $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde{Z}, (\tilde{J}^{y,y'})_{y,y'\in\mathcal{Z}})$ such that $\tilde{W}_0 = W_\tau$ P-a.s., $\tilde{W}_{\tilde{\sigma}} \in \mathcal{W}$ P-a.s., and on $[\![0, \tilde{\sigma})\!)$, $(\tilde{\beta}, \tilde{\delta})$ enforces $\tilde{A}$ in state $\tilde{S}_0$ and $\tilde{W} \in \mathcal{B}_r(\mathcal{W})$, where $\tilde{\sigma}$ is the time of the first state transition of $\tilde{S}$ defined by (15) from processes $(\tilde{J}^{y,y'})_{y,y'\in\mathcal{Z}}$ with initial state S_τ. We define the concatenated processes $\hat{W}$, $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, $\hat{M}$, and $\hat{S}$ by setting

$$W' = W 1_{[\![0,\tau)\!)} + \tilde{W}_{\cdot-\tau} 1_{[\![\tau,\infty)\!)}$$

and similarly for $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, $\hat{M}$, and $\hat{S}$. The concatenation $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ is a solution to (5) in the concatenated stochastic framework defined by

$$\hat{Z} = Z 1_{[\![0,\tau)\!)} + (\tilde{Z} + Z_\tau) 1_{[\![\tau,\infty)\!)}, \qquad \hat{J}^{y,y'} = J^{y,y'} 1_{[\![0,\tau)\!)} + (\tilde{J}^{y,y'} + J^{y,y'}_\tau) 1_{[\![\tau,\infty)\!)}.$$

Observe that $\hat{Z} = Z$ and $\hat{J}^{y,y'} = J^{y,y'}$ for all events $(y,y') \in \mathcal{Z}$, hence the concatenated stochastic framework is identical to the original framework. Thus, $\left(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}\right)$ is a solution to (5) in $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}})$. Since $\tau + \tilde{\sigma} = \hat{\sigma}$, it follows that $(\hat{\beta}, \hat{\delta})$ enforces $\hat{A}$ in state $\hat{S}$ on $[\![0, \hat{\sigma})\!)$. Moreover, $\hat{W}$ is contained in $\mathcal{B}_{r,S_\tau}(\mathcal{W})$ on the stochastic interval $[\![\tau, \hat{\sigma})\!)$ on which $\hat{W}$ coincides with $\tilde{W}$. Finally, $\hat{W}_{\hat{\sigma}} = \tilde{W}_{\tilde{\sigma}} \in \mathcal{W}_{\tilde{S}_{\tilde{\sigma}}} = \mathcal{W}_{S_{\hat{\sigma}}}$ P-a.s. □

Proof of Lemma 4.8. We first show that $\mathcal{W}_y \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ implies that $\mathcal{B}_r(\mathcal{W})$ is self-generating. To that end, fix any state $y \in \mathcal{Y}$ and any payoff pair $w \in \mathcal{B}_{r,y}(\mathcal{W})$. By definition of $\mathcal{B}_{r,y}(\mathcal{W})$, there exists a solution (W, A, β, δ, M) to (5) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}})$ such that $W_0 = w$ P-a.s., $W_\sigma \in \mathcal{W}_{S_\sigma}$ P-a.s., and on $[\![0, \sigma)\!)$, we have $W \in \mathcal{B}_{r,y}(\mathcal{W})$ and (β, δ) enforcing A for state process S defined by (15). Let σ_n denote the time of the n^{th} state transition. Since $W_\sigma \in \mathcal{W}_{S_\sigma} \subseteq \mathcal{B}_{r,S_\sigma}(\mathcal{W})$, we can apply Lemma B.2 to $\tau = \sigma$ to obtain a solution $\left(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}\right)$ to (5) on the same stochastic framework such that $\hat{W}_0 = w$ P-a.s., $\hat{W}_{\sigma_2} \in \mathcal{W}_{S_{\sigma_2}}$ P-a.s., and on $[\![0, \sigma_2)\!)$, $(\hat{\beta}, \hat{\delta})$ enforces $\hat{A}$ in state S and $\hat{W} \in \mathcal{B}_{r,S}(\mathcal{W})$. With an iteration of this concatenation procedure, we can construct solutions to (5) on $[\![0, \sigma_n)\!)$ that remains in $\mathcal{B}_{r,S}(\mathcal{W})$ up until σ_n for any n. Because Poisson processes have only countably many jumps, an iteration of this procedure constructs an enforceable solution to (5) that remains in $\mathcal{B}_r(\mathcal{W})$ forever, showing that $\mathcal{B}_r(\mathcal{W})$ is self-generating.

For the converse, suppose that $\mathcal{W} = (\mathcal{W}_y)_{y\in\mathcal{Y}}$ is self-generating. By definition, for any state $y \in \mathcal{Y}$ and any $w \in \mathcal{W}_y$, there exists an enforceable solution W to (5) that is in $\mathcal{W}_S$ a.e. In particular, $W_\sigma \in \mathcal{W}_{S_\sigma}$ a.s. It follows that $\mathcal{W}_y \subseteq \mathcal{B}_{r,y}(\mathcal{W})$ by maximality of $\mathcal{B}_{r,y}(\mathcal{W})$. Since the state was arbitrary the result follows. □

Proof of Proposition 4.9. Lemma 4.8 implies that $\mathcal{E}_y(r) \subseteq \mathcal{B}_{r,y}\big(\mathcal{E}(r)\big)$ for every state $y \in \mathcal{Y}$ and that $\mathcal{B}_r\big(\mathcal{E}(r)\big)$ is self-generating. Since $\mathcal{E}(r)$ is the largest family of bounded self-generating sets, it follows that $\mathcal{B}_{r,y}\big(\mathcal{E}(r)\big) \subseteq \mathcal{E}_y(r)$ for every $y \in \mathcal{Y}$ and hence $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$. Next, we show that $\mathcal{B}_{r,y}(\mathcal{V}^*) \subseteq \mathcal{V}^*$ for every $y \in \mathcal{Y}$, i.e., each player i attains at least $\underline{v}^i$ in $\mathcal{B}_{r,y}(\mathcal{V}^*)$. Suppose towards a contradiction that this is not the case, i.e., there exists $w \in \mathcal{B}_{r,y}(\mathcal{V}^*)$ with $w^i < \underline{v}^i$ for some i. Since $w \in \mathcal{B}_{r,y}(\mathcal{V}^*)$, there exists a solution (W, A, β, δ, M) to (5) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^{y,y'})_{(y,y')\in\mathcal{Z}})$ with $W_0 = w$ a.s., $W_\sigma \in \mathcal{V}^*$ a.s., such that on $[\![0,\sigma)\!)$, we have $W \in \mathcal{B}_{r,y}(\mathcal{V}^*)$ and (β, δ) enforces A in state y. By Lemma 4.1, we can write

$$w^i = \int_0^\sigma r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^A}\big[g^i(S_s, A_s)\big]\,\mathrm{d}s + \mathrm{e}^{-r\sigma}\mathbb{E}_{Q_\sigma^A}\big[W_\sigma^i\big]. \tag{20}$$

Since $w^i < \underline{v}^i$ but $W_\sigma^i \geq \underline{v}^i$ a.s., it follows that the first term in (20) is smaller than $(1 - \mathrm{e}^{-r\sigma})\underline{v}^i$. Define now the strategy $\tilde{A}^i = \arg\max_{a^i\in\mathcal{A}^i(y)} g^i(y, \cdot\,, A^{-i})1_{[\![0,\sigma)\!)} + A^i 1_{[\![\sigma,\infty)\!)}$. Because $\tilde{A}^i$ is i's myopic best reply to A^{-i} on $[\![0,\sigma)\!)$, it follows that $g^i(y, \tilde{A}^i, A^{-i}) \geq \underline{v}^i$ on that same stochastic interval. We deduce that

$$\int_0^\sigma r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^{\tilde{A}^i,A^{-i}}}\Big[g^i(S_s, \tilde{A}_s^i, A_s^{-i})\Big]\,\mathrm{d}s \geq (1 - \mathrm{e}^{-r\sigma})\underline{v}^i > \int_0^\sigma r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^A}\big[g^i(S_s, A_s)\big]\,\mathrm{d}s.$$

This contradicts the fact that (β, δ) enforces A in state y on $[\![0,\sigma)\!)$.[14] We have thus shown that $\mathcal{W}_{1,y} = \mathcal{B}_{r,y}(\mathcal{V}^*) \subseteq \mathcal{V}^* = W_{0,y}$ for every $y \in \mathcal{Y}$. Since $\mathcal{E}_y(r) \subseteq \mathcal{W}_{0,y}$ for every $y \in \mathcal{Y}$, monotonicity of $\mathcal{B}_r$ implies that $\mathcal{E}_y(r) \subseteq \mathcal{W}_{1,y} \subseteq \mathcal{W}_{0,y}$. An iterated application of Lemma B.1 thus shows that $(W_{n,y})_{n\geq 0}$ is decreasing in the set-inclusion sense and that it is bounded from below by $\mathcal{E}_y(r)$ for any $y \in \mathcal{Y}$. Those sequences must converge to a limit $\mathcal{W}_\infty$ with $W_{\infty,y} \supseteq \mathcal{E}_y(r)$ for every $y \in \mathcal{Y}$ such that $\mathcal{W}_{\infty,y} = \mathcal{B}_{r,y}(\mathcal{W}_\infty)$. The limit $\mathcal{W}_\infty$ is thus self-generating and hence $\mathcal{W}_\infty = \mathcal{E}(r)$ by Lemma 4.8. □

C Stationary Markov and Semi-Stationary Equilibria

We begin with the proofs of the results characterizing semi-stationary PPE.

Proof of Lemma 5.2. We first show that any payoff pair in a family $\mathcal{W}$ of bounded mutu-

[14] This follows along the same lines as the proof of Lemma 4.3 by restricting attention to $[\![0,\sigma)\!)$.

ally stationary payoff sets can be attained by a semi-stationary PPE. Fix initial state y_0 and $w_0 \in \mathcal{W}_{y_0}$. By definition of mutual stationarity, there exist $a_0 \in \mathcal{A}(y_0)$ and δ_0 such that $(0, \delta_0)$ enforces a_0, $w_0 = g(y_0, a_0) + \delta_0 \lambda(y_0, a_0)$, and $w_0 + r\delta_0(y) \in \mathcal{W}_y$ for every successor state y of y_0. A solution to (5) for $A \equiv a_0$, $\beta \equiv 0$, $\delta \equiv \delta_0$, and $M \equiv 0$ starting at w_0 thus remains in w_0 until the first time σ_1, when a state transition occurs. It follows from the choice of δ_0 that $W_{\sigma_1} \in \mathcal{W}_{S_{\sigma_1}}$. Note that w_0, δ_0, and a_0 can be viewed as a function from the initial state. Therefore, W_{σ_1} is fully determined by w_0, δ_0, and the state S_σ, that is, W_σ is a function of (y_0, S_{σ_1}). Since $W_{\sigma_1} \in \mathcal{W}_{S_{\sigma_1}}$, by definition of mutual stationarity, we can decompose W_{σ_1} by a_1 and δ_1, which are functions of W_{σ_1}, hence functions of (y_0, S_{σ_1}). We can thus get an enforceable solution to (5) on $[\![\sigma_1, \sigma_2)\!)$ that remains in W_{σ_1} until σ_2 and jumps to $\mathcal{W}_{\sigma_2}$ at time σ_2. Similarly to the proof of Lemma B.2, we can concatenate the solutions at state transitions. Since there are countably many state transitions, a concatenation will yield an enforceable solution to (5) on $[0, \infty)$, which is a PPE by Lemma 4.3. Moreover, it is semi-stationary because a_k is a function of $(y_0, S_{\sigma_1}, \ldots, S_{\sigma_k})$.

Since the union of two families of mutually stationary payoff sets is again mutually stationary, this shows that $\mathcal{E}^S(r)$ contains the largest bounded family of mutually stationary payoff sets. Thus, it remains to show that $\mathcal{E}^S(r)$ is mutually stationary.

Fix an initial state y_0 and a payoff pair $w_0 \in \mathcal{E}^S_{y_0}(r)$. Let A be a semi-stationary PPE attaining w_0 for initial state y_0, corresponding to selector a_*. Let τ be a stopping with $\tau \leq \sigma_1$ a.s. It follows along the same lines as in the proof of Lemma 5.6 that $W_\tau(A) = w_0$. Since τ was arbitrary, W is locally constant on $[\![0, \sigma)\!)$. Lemma 4.1 thus implies that $\beta \equiv 0$, $M \equiv 0$, and $\delta \equiv \delta_0$ for some δ_0 such that

$$w_0 = g(y, a_*(y_0)) + \sum_{y' \in \mathcal{Y}(y_0)} \delta_0(y') \lambda_{y_0, y'}(a_*(y_0)). \tag{21}$$

Semi-stationarity implies that also the continuation profile after the first state transition is semi-stationary and hence $W_{\sigma_1} \in \mathcal{E}^S_y(r)$ on the event $\{S_{\sigma_1} = y\}$ for all states y, for which $\{S_{\sigma_1} = y\}$ has positive measure, i.e., for all y such that $(y_0, y) \in \mathcal{Z}$. The SDE representation of the continuation value implies that $W_{\sigma_1} - w_0 = r\delta_0(y)$ on $\{S_{\sigma_1} = y\}$ for all successor states y of y_0 and hence $w_0 + r\delta_0(y) \in \mathcal{E}^S_y(r)$. Finally, since A is a PPE, Lemma 4.3 implies that $(0, \delta_0)$ enforces $a_*(y_0)$. □

Proof of Lemma 5.4. Since $\mathcal{S}_r(\mathcal{W})$ is $\mathcal{W}$-stationary, for each $y \in \mathcal{Y}$ and each $w \in \mathcal{S}_{r,y}(\mathcal{W})$, there exist $a \in \mathcal{A}(y)$ and δ such that $(0, \delta)$ enforces a in state y, $w = g(y, a) + \delta\lambda(y, a)$, and $w + r\delta(y') \in \mathcal{W}_{y'}$ for every $(y, y') \in \mathcal{Z}$. If $\mathcal{W}_{y'} \subseteq \mathcal{S}_{r,y'}(\mathcal{W})$ for every $(y, y') \in \mathcal{Z}$, then $\mathcal{S}_r(\mathcal{W})$ is also $\mathcal{S}_r(\mathcal{W})$-stationary and hence mutually stationary. □

Proof of Proposition 5.5. This proof mimics the proof of Proposition 4.9. We first note that $\mathcal{S}_r$ is monotone by definition, that is, $\mathcal{S}_{r,y}(\mathcal{W}) \subseteq \mathcal{S}_{r,y'}(\mathcal{W}')$ for any two families $\mathcal{W}, \mathcal{W}'$ of payoff sets with $\mathcal{W}_y \subseteq \mathcal{W}'_y$ for every state y. Note that $\mathcal{S}_{r,y}(\mathcal{W}) \subseteq \mathcal{B}_{r,y}(\mathcal{W}) \subseteq \mathcal{V}^*$ for any family $\mathcal{W}$ of feasible payoff sets, hence $\mathcal{W}_{1,y} \subseteq \mathcal{W}_{0,y}$. Since $\mathcal{E}^S_y(r) \subseteq \mathcal{V}^*$, monotonicity of $\mathcal{S}_{r,y}$ implies that $\mathcal{E}^S_y(r) \subseteq \mathcal{W}_{1,y} \subseteq \mathcal{W}_{0,y} = \mathcal{V}^*$ for every state y. An iterated application of $\mathcal{S}_r$ thus yields a family of payoff sets that is decreasing in the set-inclusion sense. Each member of the family must thus converge to a limit set $\mathcal{W}_{\infty,y} = \bigcap_{n\geq 0} \mathcal{W}_{n,y}$ such that $\mathcal{W}_{\infty,y} \supseteq \mathcal{E}^S_y(r)$. Since $\mathcal{S}_r(\mathcal{W}_\infty) = \mathcal{W}_\infty$, Lemma 5.4 implies that $\mathcal{W}_\infty$ is mutually stationary and hence $\mathcal{W}_\infty = \mathcal{E}^S(r)$. □

Finally, we prove the characterization of stationary Markov equilibria in Proposition 5.6.

Proof of Proposition 5.6. Suppose first that PPE $A = a_*(S_-)$ is a stationary Markov perfect equilibrium for initial state y_0. For any state $y \in \mathcal{Y}(y_0)$, let S^y denote the state process defined in (15) with initial state y and set $w_y := W_0(y, a_*(S^y_-))$. Fix an arbitrary stopping time τ and define the processes $J^{\tau,y,y'} := J^{y,y'}_{\cdot+\tau} - J^{y,y'}_\tau$ for every $(y,y') \in \mathcal{Z}$. Let S^τ be defined as in (15) from processes $(J^{\tau,y,y'})_{y,y'\in\mathcal{Z}}$ with initial state S_τ. Since $J^{y,y'}$ is a Lévy process for every $(y,y') \in \mathcal{Z}$, the process $J^{\tau,y,y'}$ is identically distributed as $J^{y,y'}$. In particular, on $\{S_\tau = y\}$, the process $\tilde{S}$ is identically distributed as S^y. This implies that

$$\tilde{A} := A_{\cdot+\tau-} = a_*(S_{\cdot+\tau-}) = a_*(\tilde{S}_-) = \sum_{y\in\mathcal{Y}(y_0)} a_*(\tilde{S}_-)1_{\{S_\tau=y\}} \overset{d}{=} \sum_{y\in\mathcal{Y}(y_0)} a_*(S^y_-)1_{\{S_\tau=y\}}$$

and hence

$$W_\tau(A) = W_0(a_*(\tilde{S})) \overset{d}{=} \sum_{y\in\mathcal{Y}(y_0)} W_0(a_*(S^y_-))1_{\{S_\tau=y\}} = \sum_{y\in\mathcal{Y}(y_0)} w_y 1_{\{S_\tau=y\}}.$$

Since τ was arbitrary, this shows that $W(A) = w_*(S)$, where $w_*(y) = w_y$ for every $y \in \mathcal{Y}(y_0)$. In particular, $W(A)$ is locally constant where S is constant. Together with Lemma 4.1, this implies that $\beta = 0$, $M = 0$, and $\delta(y') = \delta_*(S_-, y')$ for every y' such that

$$w_*(y) = g(y, a_*(y)) + \sum_{y'\in\mathcal{Y}(y_0)} \delta_*(y,y')\lambda_{y,y'}(a_*(y)) \tag{22}$$

for every state $y \in \mathcal{Y}(y_0)$. Moreover, since W has to jump from $w_*(y)$ to $w_*(y')$ when a state transition from y to y' occurs, it follows that $\delta_*(y,y') = \frac{1}{r}\big(w_*(y') - w_*(y)\big)$ for every $(y,y') \in \mathcal{Y}(y_0)^2$. Substituting the expression for $\delta_*(y,y')$ into (22) and solving for the vector w_* yields (10). Note here that

$$B = \operatorname{diag}\left(r\mathbf{1} + \mathbf{1}\Lambda_{y_0}(a_*)\right) - \Lambda_{y_0}(a_*)$$

is indeed invertible since dividing each row y' by $r + \Lambda_{y_0}(a_*)e_{y'}$ turns B into the identity matrix minus a strictly substochastic matrix. Finally, Lemma 4.3 implies (9).

Suppose that the converse holds, i.e., there exists $a_* : \mathcal{Y} \to \mathcal{A}$ with $a_*(y) \in \mathcal{A}(y)$ such that for w_* and δ_* defined in (10) for states $\mathcal{Y}(y_0)$, inequalities (9) hold. Let S be the state process starting in initial state y_0. Let W be a solution to (5) for $A = a_*(S_-)$, $\beta \equiv 0$, $\delta = \delta_*(S_-)$, and $M \equiv 0$, starting in initial state $w_*(y_0)$. Since (10) is equivalent to (22), it follows that W is locally constant unless a state change from y to y' occurs, at which point there is a jump of size $w(y') - w(y)$ in W and hence $W = w_*(S)$. Due to (9), (β, δ) enforces A and hence Lemma 4.3 implies that A is a PPE. □

D Characterization of $\mathcal{B}_r(\mathcal{W})$

Since the state is locally constant for enforceable solutions to (5) attaining payoffs in $\mathcal{B}_r(\mathcal{W})$, the characterization of $\mathcal{B}_{r,y}(\mathcal{W})$ for a fixed state y is very similar to the case of repeated games with signal (S, X). There are two differences. First, the set $\mathcal{B}_{r,y}(\mathcal{W})$ is not necessarily contained within $\operatorname{conv} g(y, \mathcal{A}(y))$ because continuation payoffs partially come from $\operatorname{conv} g(y', \mathcal{A}(y'))$ for successor states y' of y. Second, for repeated games, the approximation in algorithm in Proposition 4.9 leads to a decreasing sequence of payoff sets, each satisfying $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$. In stochastic games, incentives related to state transitions may not come from a superset of $\mathcal{B}_{r,y}(\mathcal{W})$: while $\mathcal{B}_{r,y}(\mathcal{W}) \subseteq \mathcal{W}_y$ is still true for the approximating sequence in Proposition 4.9, jump incentives come from the sets $\mathcal{W}_{y'}$ for successor states y' of y. These two differences affect local Lipschitz continuity of (12), which we shall prove here. Once we have local Lipschitz continuity, the remainder of the proof works analogously as in Bernard [4] and we omit it here.

For the proof, we need a slightly more general version of the optimality equation. For a convex set $\mathcal{X}$ and any $h \geq 0$, let $\mathcal{X}_h := \{v \in \mathbb{R}^2 \mid \min_{w \in \mathcal{X}} \|v - w\| \leq h\}$ denote the set of all payoff pairs within distance h from $\mathcal{X}$. Call a family $\mathcal{L} = (\mathcal{L}_y)_{y \in \mathcal{Y}}$ of set-valued maps $\mathcal{L}_y : \mathbb{R}^2 \rightrightarrows \mathbb{R}^{2 \times |Y|}$ a *Lipschitz expansion* of $\mathcal{W} = (\mathcal{W}_y)_{y \in \mathcal{Y}}$ if

$$\mathcal{L}_y(w) = \left\{ \delta \in \mathbb{R}^{2 \times |Y|} \;\middle|\; w + r\delta(y') \in \mathcal{W}_{y', h(w)} \text{ for all } y' \text{ with } (y, y') \in \mathcal{Z} \right\}$$

for a non-negative Lipschitz-continuous function $h : \mathbb{R}^2 \to \mathbb{R}$. Observe that each $\mathcal{L}_y$ is Lipschitz continuous in the sense of set-valued maps. We denote by $\mathcal{L}^0$ the trivial Lipschitz expansion with $h \equiv 0$. Let $\Psi_{y,a}(w, N, \mathcal{L})$ denote the set of all $\delta \in \mathcal{L}_y(w)$, for which there exists β with $N^\top \beta = 0$ such that (β, δ) enforces a. Denote by

$$E_{y,a}(\mathcal{L}) := \left\{ (w, N) \in \mathbb{R}^2 \times S^1 \;\middle|\; \Psi_{y,a}(w, N, \mathcal{L}) \neq \emptyset \right\}$$

the effective domain of of $(w, N) \mapsto \Psi_{y,a}(w, N, \mathcal{L})$. For any N and any $\delta \in \Psi_{y,a}(w, N, \delta)$, let $\Phi_{y,a}(N, \delta)$ denote the set of all $\phi \in \mathbb{R}^d$, for which $(T\phi, \delta)$ enforces a, where T is the clockwise orthogonal vector to N. Let $\phi_y(a, N, \delta)$ denote the shortest vector in $\Phi_{y,a}(N, \delta)$. We can now write (12) in the following form, for which we prove local Lipschitz continuity.

Lemma D.1. *Suppose that Assumptions 1 and 2 are satisfied. Then*

$$\kappa_{y,\mathcal{L}}(w, N) := \max_{a\in\mathcal{A}(y)} \max_{\delta\in\Psi_{y,a}(w,N,\mathcal{L})} \frac{2{N_w}^\top\big(g(y,a) + \delta\lambda(y,a) - w\big)}{r\|\phi_y(a, N, \delta)\|^2} \vee 0 \tag{23}$$

is locally Lipschitz continuous outside of $\Gamma_y(\mathcal{L}) = \bigcup_{a\in A(y)} \Gamma_{y,a}(\mathcal{L})$*, where*

$$\Gamma_{y,a}(\mathcal{L}) := \{(w, N) \mid \textit{there exists } \delta \in \mathcal{L}_y(w) \textit{ with } (0, \delta) \textit{ enforces } a\}.$$

The basic idea behind the proof of Lemma D.1 is to show that the right-hand side of (23) is locally Lipschitz continuous for fixed a on $E_{y,a}(\mathcal{L}) \setminus \Gamma_{y,a}(\mathcal{L})$. Then $\kappa_{y,\mathcal{L}}$ is locally Lipschitz continuous except at $\bigcup_{a\in\mathcal{A}(y)} \partial E_{y,a}(\mathcal{L})$. The first lemma shows that $\partial E_{y,a}(\mathcal{L})$ is contained in $\mathbb{R}^2 \times \{\pm e_1, \pm e_2\}$ for any $a \in \mathcal{A}(y)$.

Lemma D.2. *Suppose that Assumption 1 is satisfied. For any Lipschitz expansion* $\mathcal{L}$*, any state* y*, and any* $a \in \mathcal{A}(y)$*, we have* $\mathbb{R}^2 \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \subseteq E_{y,a}(\mathcal{L})$*. Moreover,* $(w, N) \mapsto \Psi_{y,a}(w, N, \mathcal{L})$ *is locally Lipschitz continuous on* $\mathbb{R}^2 \times (S^1 \setminus \{\pm e_1, \pm e_2\})$*.*

Proof. Fix any $\mathcal{W}$, y, a, $w \in \mathbb{R}^2$, and any non-coordinate N. Choose any δ in $\mathcal{L}_y(w)$. Since $M^i(a)$ has individual full rank by Assumption 1, there exists β^i that solves (6) for player i with equality. Since a is pairwise identifiable by Assumption 1, it follows from Lemma 2 in Sannikov [29] applied to payoff function $\tilde{g}(y, a) = g(y, a)+\delta\lambda(y, a)$ that there exists β' with $N^\top\beta' = 0$ such that (β', δ) enforces a. In particular, $\Psi_{y,a}(w, N, \mathcal{L}) \neq \emptyset$ and hence $(w, N) \in E_{y,a}(\mathcal{L})$. Moreover, this construction shows that $\Psi_{y,a}(w, N, \mathcal{L}) = \mathcal{L}_y(w)$ for non-coordinate N, which is locally Lipschitz continuous. □

Proof of Lemma D.1. Let $G^i_y(a)$ denote the row vector with entries $g^i(y, \tilde{a}^i, a^{-i})-g^i(y, a)$ and let $\Lambda^i_y(a)$ denote the matrix with column vectors $\lambda(y, \tilde{a}^i, a^{-i}) - \lambda(y, a)$. For non-coordinate N, $\Phi_{y,a}(N, \delta)$ is the set of all ϕ that satisfy $\phi M^i_y(a) \le -\frac{1}{T^i}\big(G^i_y(a) + \delta^i\Lambda^i_y(a)\big)$ for $i = 1, 2$. Since $M^i_y(a)$ does not depend on N or δ and the right-hand side is locally Lipschitz continuous in N, δ, the set $\Phi_{y,a}(N, \delta)$ is a constant-rank polyhedron with locally Lipschitz continuous right-hand side. It follows from the main result of Yen [32] that the projection of 0 onto $\Phi_a(N, \delta)$ is locally Lipschitz continuous in (N, δ). In particular, $\phi_y(a, N, \delta)$ is locally Lipschitz continuous in (N, δ) for non-coordinate N.

Consider now a coordinate direction $N = \pm e_i$. Then any β with $N^\top\beta$ cannot provide any incentives to player i, hence any $\delta \in \Psi_{y,a}(w, N, \mathcal{L})$ must satisfy $G^i_y(a) + \delta^i\Lambda^i_y(a) \le 0$. Denote by $\Phi^{-i}_{y,a}(\delta)$ the set of all ϕ that satisfy $\phi M^{-i}_y(a) \le -\big(G^i_y(a) + \delta^{-i}\Lambda^{-i}_y(a)\big)$. As above, the shortest vector $\phi^{-i}_y(a, \delta)$ in $\Phi^{-i}_{y,a}(\delta)$ is locally Lipschitz continuous in δ as the projection of 0 onto $\Phi^{-i}_{y,a}(\delta)$. The shortest vector in $\Phi_{y,a}(N, \delta) = \frac{1}{T^1}\Phi^1_{y,a}(\delta) \cap \frac{1}{T^2}\Phi^2_{y,a}(\delta)$ is thus at least as long as $\frac{1}{T^{-i}}\phi^{-i}_y(a, \delta)$. Since $\operatorname{span} M^i_y(a)$ is orthogonal to $\operatorname{span} M^{-i}_y(a)$ it follows that $\frac{1}{T^{-i}}\phi^{-i}_y(a, \delta)M^i_y(a) = 0$, hence $\phi_y(a, N, \delta) = \frac{1}{T^{-i}}\phi^{-i}_y(a, \delta)$, which is locally Lipschitz continuous in a neighborhood of (N, δ).

Since $\Gamma_y(\mathcal{W})$ is closed, for any $(w, N) \notin \Gamma_y(\mathcal{W})$, there exists a neighborhood on which $\|\phi_y(a, N, \delta)\|$ is bounded away from 0. For non-coordinate N, it follows from Lemma D.2 that $\kappa_{y,\mathcal{L}}(w, N)$ is locally Lipschitz continuous as the maximum of a locally Lipschitz continuous function over a locally Lipschitz continuous set; see Lemma B.2 in Bernard [4]. For (w, N) with coordinate N, fix any action profile a with $\Psi_{y,a}(w, N, \mathcal{L}) = \emptyset$. This means that there exists no $\delta \in \mathcal{L}(w)$ with $0 \in \Phi^i_{y,a}(\delta)$. Since $\Phi^i y, a(\delta)$ is closed, this implies that 0 is bounded away from $\Phi^i_{y,a}(\delta)$, hence $\left\|\phi_y(a, \tilde{N}, \delta)\right\|$ converges to ∞ as $\tilde{N}$ converges to N because it is the shortest vector in $\Phi_{y,a}(\tilde{N}, \delta) = \frac{1}{\tilde{T}^1}\Phi^1_{y,a}(\delta) \cap \frac{1}{\tilde{T}^2}\Phi^2_{y,a}(\delta)$. Let $\mathcal{A}_y(w, N)$ denote the action profiles $a \in \mathcal{A}(y)$, for which

$$\kappa_{y,a,\mathcal{L}}(w, N) := \max_{\delta \in \Psi_{y,a}(w,N,\mathcal{L})} \frac{2N_w{}^\top\big(g(y, a) + \delta\lambda(y, a) - w\big)}{r\|\phi_y(a, N, \delta)\|^2} \vee 0$$

is strictly positive. Then $\kappa_{y,a,\mathcal{L}}(w', N') > 0$ for (w', N') in a neighborhood of (w, N) and hence $\kappa_{y,\mathcal{L}}(w', N') = \max_{a \in \mathcal{A}_y(w',N')} \kappa_{y,a,\mathcal{L}}(w, N)$ in a neighborhood of (w, N). In particular, $\kappa_{y,\mathcal{L}}$ is locally Lipschitz continuous at (w, N). □