CONTINUOUS-TIME GAMES WITH IMPERFECT AND ABRUPT INFORMATION*

Benjamin Bernard†

This paper studies two-player games in continuous time with imperfect public monitoring, where information may arrive both gradually and continuously, governed by a Brownian motion, and abruptly and discontinuously, according to Poisson processes. For this general class of two-player games, we characterize the equilibrium payoff set via a convergent sequence of differential equations. The resulting differential equation takes a different form depending on whether or not incentives can be provided through abrupt information exclusively. Moreover, in the presence of abrupt information, the boundary of the equilibrium payoff set may not be smooth outside the set of static Nash payoffs. Equilibrium strategies that attain extremal payoff pairs as well as their intertemporal incentives are elicitable from the limiting solution.

Keywords: Repeated games, continuous time, information arrival, imperfect observability, equilibrium characterization.

1 Introduction

A major goal of the theory of repeated games is to explain how the structure of intertemporal incentives in extremal equilibria depends on features of the underlying information structure. For two-player games with public monitoring, there are two qualitative features of monitoring structures whose impact on the optimal provision of incentives is reasonably well understood: (i) whether deviations by a player from a given profile can be statistically distinguished from deviations by the other player

*I would like to thank Christoph Frei, Tilman Klumpp, participants of the 11$^{\text{th}}$ World Congress of the Econometric Society, the 5$^{\text{th}}$ World Congress of the Game Theory Society, the Asian Meeting of the Econometric Society 2016, the 2017 Midwest Economic Theory Conference, and seminar participants at the University of Amsterdam, the University of Bern, the University of Bonn, Carnegie Mellon University, National Taiwan University, and Academia Sinica for valuable comments and interesting discussions. Financial support by the Natural Sciences and Engineering Research Council of Canada through grant RGPIN/402585-2011, the Swiss National Science Foundation through grant P2SKP1_171737, 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.

†Department of Economics, National Taiwan University, No. 1, Sec. 4, Roosevelt Rd, Da'an District, Taipei City, 10617 Taiwan; benbernard@ntu.edu.tw.

and (ii) how much overall noise is involved in the latter inference. The literature on discrete-time repeated games has shed a lot of light on the impact of (i) and (ii).[1] However, monitoring structures can also differ in other important attributes, chief among them the nature of the *flow of information*: whether information arrives gradually and continuously over time or abruptly and discontinuously. In order to study the impact of the flow of information on the structure of intertemporal incentives, it is natural to consider a continuous-time model so that information may arrive at *any* time and not just at multiples of a fixed time interval.

Starting with Abreu, Milgrom, and Pearce [1], most existing models on the timing of public information augment the standard discrete-time framework by allowing players to continuously observe the public signal. In his seminal paper, Sannikov [22] has shown that additional insights on intertemporal incentives can be gained by studying models, in which not only the public signal is observed continuously, but actions are taken continuously as well.[2] Specifically, in two-player games it is possible to relate intertemporal incentives of extremal equilibria to the geometry of the equilibrium payoff set via an ordinary differential equation. In Sannikov's model, information about past actions arrives gradually and continuously through the observation of a drifted Brownian motion. The purpose of this paper is to extend the continuous-time framework to a monitoring structure, in which public signals may arrive continuously as in Sannikov [22], but also abruptly and discontinuously through the observation of Poisson processes. This allows us to contrast how the flow of information affects payoffs and intertemporal incentives in extremal equilibria.

In many applications, there is a clear distinction between the two types of information arrival. Consider, for example, a partnership between two firms in which the firms lack the expertise to directly evaluate the quality of the component contributed by the partnering firm. Each firm has an incentive to produce a lower-quality component at a lower cost as it may do so undetectedly. The partnering firm only receives noisy information from the performance of the product on the market. Demand moves continuously due to day-to-day fluctuations and it is subject to demand shocks when the product receives negative media coverage. As illustrated in Figure 1, a decomposition of total demand leads to two separate signals that are both indicative of the quality of the partner's component: the continuous increase in demand without the impact of demand shocks and the frequency at which demand shocks occur. How the available information should be used to write the optimal effort-inducing agreement between the two parties crucially depends on the flow of information.

[1]See Abreu, Milgrom, and Pearce [1], Abreu, Pearce, and Stacchetti [2, 3], Fudenberg, Levine, and Maskin [11], Fudenberg and Levine [9], and Sannikov and Skrzypacz [24].

[2]In a continuous-time framework, actions are interpreted as the rate, at which players execute a given task: the rate of production in an oligopoly or an effort rate in partnership models. While extremal equilibria typically require players to adjust actions infinitesimally quickly, doing so is not necessary to attain asymptotic efficiency; see Bernard and Frei [8].

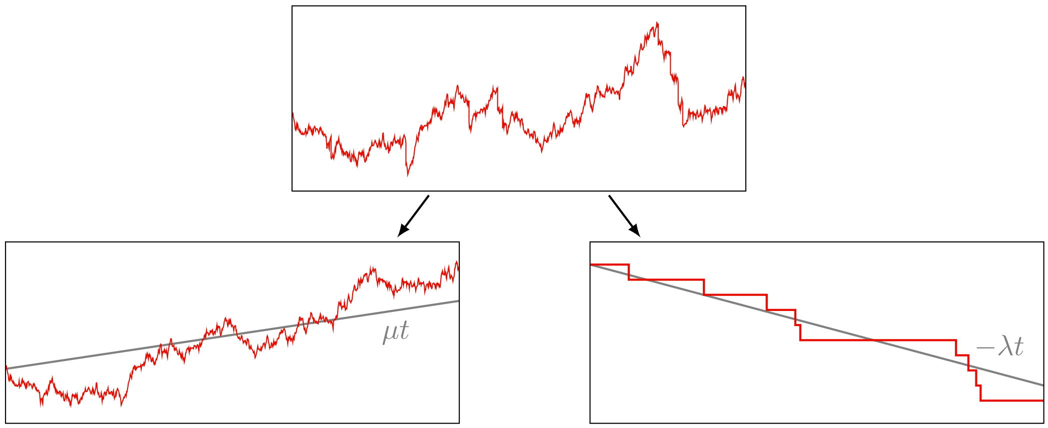

Figure 1: The total demand for the product (top panel) responds to normal market fluctuations (bottom left panel) and shocks due to negative media coverage (bottom right panel). Both components carry important information about past play: the expected rise in continuous revenue μ and the frequency λ of negative media coverage depend on the chosen effort levels.

A first distinction between the two types of information arrival is the distinct set of intertemporal incentives that can be sustained for equilibria on the efficient frontier. Abreu, Pearce, and Stacchetti [2] have shown that mutual punishments after undesirable signal realizations as in Green and Porter [12] can be an efficient way of providing equilibrium incentives. These incentives, to which we will refer as *value burning*, cannot be sustained on the efficient frontier using only continuous information: because Brownian information is so noisy, the chance of wrongfully entering a punishment phase is prohibitively large. Burning value after undesirable outcomes of Brownian information thus destroys too much value in expectation.[3] Brownian information is therefore used only to transfer value tangentially between players as in Fudenberg, Levine, and Maskin [11]. Abrupt information is not restricted in this way and can be used to transfer value tangentially and/or to burn value. On the efficient frontier, value burning corresponds to an inward movement and value transfers to a tangential movement of the continuation value. For ease of reference, we say that value is burnt/transferred whenever the continuation moves inwards/tangentially also for inefficient extremal equilibrium payoffs. See Figure 2 for an illustration.

Continuous-time methods provide additional tractability that may give rise to explicit results that are not available in discrete time. Sannikov [22] shows that for two-player games with Brownian information, it is possible to express the curvature of the equilibrium payoff set in terms of the unique equilibrium incentives at that point. As a result, the boundary of the equilibrium payoff set can be described by an ordinary differential equation (ODE). Such an explicit characterization of the equilibrium payoff set for any level of discounting is a result without analogue in discrete time. Discrete-time results on equilibrium payoffs are often limited to folk theo-

[3] Sannikov and Skrzypacz [24] provide an excellent discussion of this point.

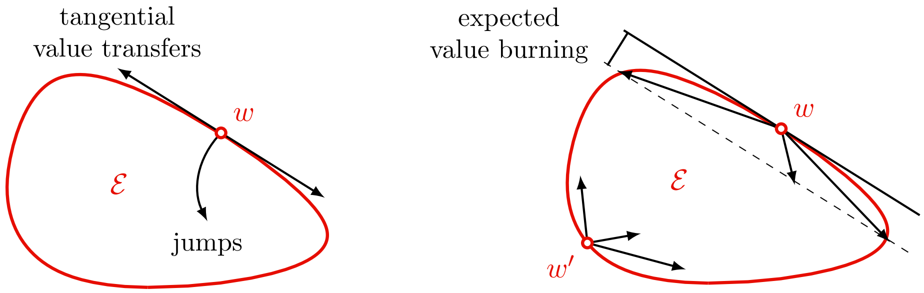

Figure 2: Because Brownian information is used to transfer value between players along tangents, only local information about the geometry of the equilibrium payoff set $\mathcal{E}$ is needed to structure incentives at a payoff pair w on the boundary. In contrast, abrupt information may be used to transfer value tangentially as well as to burn value—requiring global information about $\mathcal{E}$. For ease of reference, we refer to inward movements as value burning even though at the inefficient payoff pair w' an inward movement may correspond to a reward for one or both players.

rems and payoff bounds as players get arbitrarily patient; see Fudenberg, Levine and Maskin [11] and Fudenberg and Levine [9], respectively. Crucial for Sannikov's result is the assumption that information is Brownian. Equilibrium incentives at the boundary then consist of tangential value transfers only and, in particular, they depend on the geometry of the equilibrium payoff set only through the direction of the tangent. Therefore, only local information about the boundary of the equilibrium payoff set is needed to characterize the set of equilibrium payoffs, in the same way that any ODE uses only local information of a function to encode its global properties. In contrast, when information arrives according to Poisson processes, any punishments or rewards are consistent with equilibrium behavior as long as the continuation value remains within the equilibrium payoff set. This restriction on incentives, however, depends on the entire geometry of the equilibrium payoff set. To describe the boundary locally through equilibrium incentives, global information about the set of equilibrium payoffs is required. A differential equation describing the boundary of the equilibrium payoff set via equilibrium incentives is thus self-referential.

We show that this self-referential differential equation can be approximated by a convergent sequence of explicit ODEs that are obtained with an iterative procedure over the arrival times of infrequent events. This requires the introduction of the new concept of *relaxed self-generating* payoff sets, where in each step, the continuation value after a punishment/reward due to the arrival of an infrequent event has to come from the payoff set from the previous step. An iterated computation of the largest relaxed self-generating payoff set, starting with set of feasible and individually rational payoffs, converges to the set of equilibrium payoffs. This algorithm is similar in spirit to the discrete-time algorithm in Abreu, Pearce and Stacchetti [3]. However, contrary to its discrete-time counterpart, the sets in every step of the continuous-time algorithm can be computed efficiently as the numerical solution to an explicit ODE. This paper thus also contributes to the literature on computing equilibrium payoffs, providing an alternative to Judd, Yeltekin and Conklin [15] for two-player games and an extension of Abreu and Sannikov [4] to imperfect monitoring. The concept of relaxed self-generating payoff sets will be of interest in itself for future works in continuous-time games, in which information arrives according to Poisson processes as in, for example, continuous-time stochastic games with finitely many states.

The ODE that characterizes the boundary of the largest relaxed self-generating payoff set $\mathcal{B}$ takes two different forms depending on whether or not incentives via the Brownian information are required. If they are required, the noisiness of the Brownian signal makes the boundary smooth and the boundary is characterized by the *optimality equation*, a 2nd-order ODE that extends the ODE in Sannikov [22]. Boundary payoffs of $\mathcal{B}$, where incentives via Brownian information are not required, fall into one of three categories. First, the majority of such payoffs are a differentiable solution to a 1st-order ODE, called the *abrupt-information optimality equation*, which equates the distance of the expected flow payoff from the separating tangent with the minimal expected value that has to be burnt to provide incentives. It is in some sense a limiting ODE of the optimality equation as incentives provided via the Brownian signal disappear. Despite this similarity, the techniques used in the derivation are quite different from the optimality equation since the smoothness property from the Brownian signal is missing, on which the techniques in Sannikov [22] build. Second, boundary payoffs are *stationary* if they are the sum of expected flow payoff and expected rewards/punishments provided upon the realization of discrete events. Those payoffs are called stationary because the impact of instantaneous payoff extraction and expected rewards/punishments on the continuation value precisely offset each other. The continuation value of any enforceable strategy profile attaining a stationary boundary payoff is thus locally constant until the arrival of a discrete event. Until such an event occurs, players receive no new information that warrants an adjustment of their chosen actions, hence the strategy profile is locally constant as well. The set $\mathcal{B}$ may have corners at stationary payoffs, which includes the set of static Nash payoffs. Finally, $\mathcal{B}$ may have corners where incentives provided via abrupt information are both binding and extremal. Most often, such corners correspond to starting points of boundary segments that solve the abrupt-information optimality equation.

We then proceed to analyze how much the abrupt information contributes to intertemporal incentives at boundary payoffs, where both continuous and abrupt information are used. Because punishments/rewards after a discrete event lead to a discontinuous change in the continuation value, the continuation value lies below the tangent where the boundary of the equilibrium payoff set is curved; see the right panel of Figure 2. Thus, where the boundary is curved, incentives provided through the discontinuous information necessarily burn a positive amount of value. As a result, the continuation value after a discrete even is chosen close to tangentially to reduce the amount of value burnt. This is reminiscent of the proof of the folk theorem in Fudenberg, Levine, and Maskin [11]: to attain asymptotic efficiency, players cannot rely on incentives that burn *any* value. This paper does not prove a folk theorem but instead characterizes the equilibrium payoff set for any positive discount rate. Therefore, some value burning is possible even in extremal equilibria, but used sparingly unless incentives cannot be provided through value transfers.

The impact of information arrival on equilibrium payoffs has been analyzed by a variety of papers using a discrete-time framework with increasingly frequent actions. Abreu, Milgrom, and Pearce [1] demonstrate that intertemporal incentives may be destroyed in games with Poisson signals if the signals are not sufficiently informative. Sannikov and Skrzypacz [23] show that intertemporal incentives are destroyed in a Cournot duopoly with Brownian information. In a model where players observe only cumulative signals, Fudenberg and Levine [10] demonstrate that Brownian information arrival supports a larger (smaller) set of equilibrium behavior than Poisson information arrival if actions affect the volatility (drift) of the Brownian motion. The first paper to model continuous observation of both continuous and discontinuous information arrival is Sannikov and Skrzypacz [24]. They establish a bound for the limiting equilibrium payoff set as players get arbitrarily patient. Our paper deviates from this stream of literature by directly studying the continuous-time game. This allows us to address questions not only on asymptotic efficiency, but also on equilibrium behavior for any level of discounting: the solution to the approximating sequence of ODEs reveals the unique action profiles and intertemporal incentives that are used at any given extremal equilibrium payoff. This allows us to show with relative ease that any equilibrium payoff in the partnership game with sufficiently informative demand shocks can be attained by play of one-sided effort, followed by a forgiving grim trigger profile or permanent Nash reversion. We can elicit equilibria in this way not just in the partnership game, but in any finite-action two-player game.

The remainder of the paper is organized as follows. We introduce the continuous-time model with the general information structure in Section 2. We provide a detailed example of a partnership game with a preview of our results in Section 3. Section 4 contains the important concepts of enforceability and self-generation in our setting. We develop the concept of relaxed self-generating payoff sets and show how it is applied to approximate the equilibrium payoff set in Section 5. In Section 6, we present our main results: the characterization of a relaxed-self generating payoff set and the iterative construction of equilibria. In Section 7, we discuss the model assumptions as well as how abrupt information is used in more detail. A description of how to implement the numerical solution of our algorithm is presented in Section 8 and we conclude in Section 9. The vast majority of the proofs are contained in Appendices A–F.

2 The setting

Consider a game where two players $i = 1, 2$ continuously choose actions from the finite sets $\mathcal{A}^i$ at each point in time $t \in [0, \infty)$. The set of all pure action profiles

$a = (a^1, a^2)$ is denoted by $\mathcal{A} = \mathcal{A}^1 \times \mathcal{A}^2$. Players cannot directly observe each other's actions and instead see only the impact of the chosen actions on the distribution of a public signal. The signal contains gradual and continuous information modeled by a d-dimensional Brownian motion Z, as well as abrupt and discontinuous information through the observation of infrequent events. There are finitely many (possibly zero) different types of infrequent events, indexed by $y \in Y$. Events arrive according to Poisson processes $(J^y)_{y \in Y}$ that are independent of each other and independent of the Brownian motion Z. An event of type y leads to a jump in the public signal of size $h(y)$ so that the public signal is given by $X = Z + \sum_{y \in Y} h(y)\, J^y$.

The public information at time t is a σ-algebra $\mathcal{F}_t$ that contains the history of the processes $Z, (J^y)_{y \in Y}$ up to time t, as well as orthogonal information that players may use for public randomization. Because we study perfect public equilibria, a player's choice of action at time t must be based solely on information in $\mathcal{F}_t$.

Definition 2.1. A *(public) pure strategy* A^i for player i is an $(\mathcal{F}_t)_{t \geq 0}$-predictable stochastic process with values in $\mathcal{A}^i$.

The game primitives $\mu : \mathcal{A} \to \mathbb{R}^d$ and $\lambda(y \,|\, \cdot\,) : \mathcal{A} \to (0, \infty)$ determine the impact of a chosen action profile on the drift rate of the public signal and the intensity of events of type $y \in Y$, respectively. We denote by $\lambda(a) := \big(\lambda(y_1 \,|\, a), \ldots, \lambda(y_{|Y|} \,|\, a)\big)^\top$ the vector of intensities of all types of events. We assume that an event of any type y is possible after any history, that is, it is a game of full-support public monitoring.

Assumption 1 (Full support)**.** $\lambda(y \,|\, a) > 0$ for all $a \in \mathcal{A}$ and all $y \in Y$.

Because at any time t, the chosen strategy profile affects the future distribution of the public signal, play of a strategy profile $A = (A^1, A^2)$ induces a family of probability measures $Q^A = \big(Q_t^A\big)_{t \geq 0}$, under which players observe the public signal.[4] On $[0, T]$ for any $T > 0$, the public signal signal takes the form

$$X_t = \int_0^t \mu(A_s)\, \mathrm{d}s + Z_t^A + \sum_{y \in Y} h(y)\, J_t^y,$$

under Q_T^A, where $Z^A = Z - \int \mu(A_s)\, \mathrm{d}s$ is a Q_T^A-Brownian motion describing noise in the continuous component and J^y has instantaneous intensity $\lambda(y \,|\, A)$ under Q_T^A.

[4]Formally, Z and $(J^y)_{y \in Y}$ are defined on a probability space $(\Omega, \mathcal{F}, P)$ for a preliminary probability measure P. Under P, Z is a standard Brownian motion and J^y has intensity 1 for every event $y \in Y$. The family $Q^A = \big(Q_t^A\big)_{t \geq 0}$ is defined via its density process relative to P, given by

$$\exp\left(\int_0^t \mu(A_s)\, \mathrm{d}Z_s - \int_0^t \left(\frac{1}{2}|\mu(A_s)|^2 + \sum_{y \in Y} \lambda(y \,|\, A_{s-}) - 1\right) \mathrm{d}s\right) \prod_{y \in Y} \prod_{0 < s \leq t} \big(1 + (\lambda(y \,|\, A_{s-}) - 1)\Delta J_s^y\big).$$

Remark 2.1. With the techniques in this paper it is possible to consider signals of the slightly more general form $X = \sigma Z + \sum_{y\in Y} h(y)\, J^y$ for a k-dimensional Brownian motion Z and covariance matrix $\sigma \in \mathbb{R}^{d\times k}$ with rank d. Then σ has right-inverse $\sigma^\top(\sigma\sigma^\top)^{-1}$ and the game is equivalent to the game with public signal

$$\tilde{X}_t = \int_0^t \sigma^\top(\sigma\sigma^\top)^{-1}\mu(A_s)\,\mathrm{d}s + Z_t^A + \sum_{y\in Y}\sigma^\top(\sigma\sigma^\top)^{-1}h(y)\,J_t^y.$$

Indeed, the information carried by $\tilde{X}$ is identical to the information in $X = \sigma\tilde{X}$.[5]

Remark 2.2. Anderson [5] and Simon and Stinchcombe [25] demonstrate that seemingly simple strategies need not necessarily lead to unique outcomes in continuous-time games of perfect monitoring. The problem arises when strategies depend on actions chosen in the immediate past by either player. This is not a problem in our model because we restrict attention to public strategies. Indeed, one can identify the probability space with the path space of all publicly observable processes, hence a realized path $\omega \in \Omega$ leads to the unique outcome $A(\omega)$.

Definition 2.2. Let $r > 0$ be a discount rate common to both players. Each player i receives an unobservable expected flow payoff $g^i : \mathcal{A} \to \mathbb{R}$.[6]

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

$$W_t^i(A) = \int_t^\infty r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q_s^A}\left[g^i(A_s)\,\middle|\,\mathcal{F}_t\right]\mathrm{d}s. \tag{1}$$

(ii) A pure strategy profile A is a *perfect public equilibrium (PPE)* for discount rate r if for every player i and all possible deviations $\tilde{A}^i$,

$$W^i(A) \geq W^i\big(\tilde{A}^i, A^{-i}\big) \text{ a.e.,}^7$$

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

(iii) We denote the set of all payoff pairs achievable by pure-strategy PPE by

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

[5] Contrary to discrete-time models with frequent actions (c.f. Fudenberg and Levine [10]), the volatility of the continuous component of the public signal is perfectly observable in a continuous-time setting through its quadratic variation process. Therefore, players' actions cannot affect the volatility of the diffusion in a continuous-time game with imperfect monitoring.

[6] The expected flow payoff is the expectation of an observable ex-post flow payoff that carries no information about past actions beyond the public signal. Similarly to discretely repeated games, incentives and equilibrium payoffs depend only on the ex-ante flow payoffs.

[7] Deviations of a PPE are not profitable *almost everywhere* (a.e.), that is, the inequality $W_t^i(A;\omega) \geq W_t^i\big(\tilde{A}^i, A^{-i};\omega\big)$ holds for every pair (ω, t) except on a set of $P \otimes$ *Lebesgue*-measure 0.

The form of the players' continuation value in (1) 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, which is reflected in the change of measure in the expectation operator. Because the weights $re^{-r(s-t)}$ in (1) 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 is thus given by the convex hull of pure action payoff pairs $\mathcal{V} := \text{conv}\,\{g(a) \mid a \in \mathcal{A}\}$. Because we restrict attention to pure strategies, by deviating to his strategy of myopic best responses, each player i can ensure that his payoff in equilibrium dominates his minmax payoff

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

The set of equilibrium payoffs is thus contained in the set of all feasible and individually rational payoffs $\mathcal{V}^* := \{w \in \mathcal{V} \mid w^i \geq \underline{v}^i \text{ for all } i\}$. Let $\mathcal{A}^N \subseteq \mathcal{A}$ denote the set of sage-game Nash equilibria and denote by $\mathcal{V}^N := \text{conv}\left\{g(a) \,\middle|\, a \in \mathcal{A}^N\right\}$ the corresponding payoff pairs. Because indefinite play of a stage-game Nash profile is a PPE, we obtain the inclusions $\mathcal{V}^N \subseteq \mathcal{E}(r) \subseteq \mathcal{V}^* \subseteq \mathcal{V}$. Observe that $\mathcal{E}(r)$ is convex because players are allowed to use public randomization. Indeed, for any two PPE A and A' with expected payoffs $W_0(A)$ and $W_0(A')$, respectively, any payoff pair $\nu W_0(A) + (1-\nu)W_0(A')$ 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.

3 Example of a partnership game

Consider a partnership game between two players, where each player continuously chooses an effort level from the set $\mathcal{A}^i = \{0,1\}$ at every point in time t. We suppose that players receive an expected flow payoff of $g^i(a) = 4(a^1 + a^2) - a^1a^2 - 5a^i$, that is, players enjoy output but dislike effort. Players only imperfectly observe each other's effort levels through a continuous stream of revenue X^γ and the arrival of demand shocks $J^{A,\gamma}$, where the parameter $\gamma \in [0,1]$ captures the relative informativeness of the two signals as defined below. There are two separate continuous streams of revenue, given by $\mathrm{d}X_t^\gamma = \mu_\gamma(A_t)\,\mathrm{d}t + \mathrm{d}Z_t^{A,\gamma}$ for a 2-dimensional Brownian motion $Z_t^{A,\gamma}$, where $\mu_\gamma(a) = 2.4(1-\gamma)a$. The arrival of demand shocks is governed by a Poisson process $J^{A,\gamma}$ with instantaneous intensity $\lambda_\gamma(A_t)$, where

$$\lambda_\gamma(a) = \gamma\big(21 - 4(a^1 + a^2) - 12a^1a^2\big).$$

Figure 3 shows the equilibrium payoff sets for $r = 5$ and different values of γ. For low values of γ, the continuous information is at its most informative, whereas the abrupt information gets more informative as γ increases.

As the figure shows, neither type of information arrival dominates the other in terms of impact on equilibrium payoffs, but the two types serve different purposes.

	0	1
0	0, 0	4, −1
1	−1, 4	2, 2

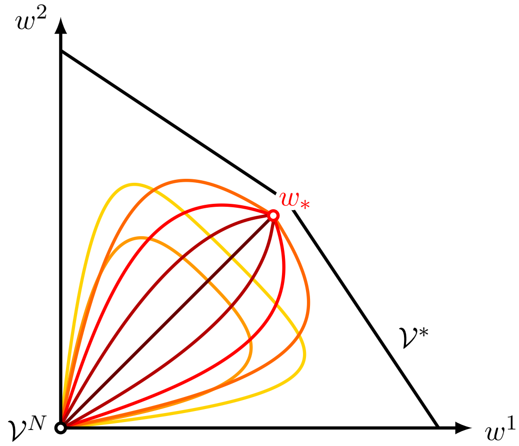

Figure 3: The left panel shows the matrix of stage game payoffs. The right panel shows the equilibrium payoff set $\mathcal{E}_\gamma(5)$ as a subset of $\mathcal{V}^*$ for γ ranging from $\gamma = 0$ (yellow) to $\gamma = 1$ (dark red) in increments of size 0.2. If the demand shocks are sufficiently informative, stationary payoff w_* can be attained in equilibrium. As the informativeness of the continuous signal increases, players can transfer larger amounts of value in equilibrium, which widens the equilibrium payoff set.

Abrupt information may significantly reduce the patience required to enforce mutual effort. If the demand shocks are sufficiently informative ($\gamma \geq 1/3$), mutual effort can be enforced using abrupt information exclusively by burning at least $5/(8\gamma)$ payoff units upon the arrival of a demand shock. Doing so by, say, entering a Nash punishment phase for a certain amount of time after a demand shock, allows the players to attain the highest symmetric equilibrium payoff pair $w_* = (1.875, 1.875)$; see Figure 4. If the demand shocks are relatively uninformative ($\gamma < 1/3$), the payoff pair w_* can no longer be attained in equilibrium only for discount rate $r = 5$. Brownian information in this example helps to attain non-stationary payoffs. Because information arrives continuously, players can adjust tangential transfers dynamically to enforce asymmetric effort levels. Figure 3 shows that the equilibrium payoff set gets wider as the Brownian information gets more informative because players are able to transfer larger amounts of value between each other in equilibrium. For this specific parametrization of the partnership game, the equilibrium payoff $\mathcal{E}_\gamma(5)$ for $\gamma = 0.3$ coincides with the set of static Nash payoffs: neither continuous nor discontinuous information are sufficiently informative to provide intertemporal incentives.

Equilibrium actions and incentives are unique on the boundary. This allows us to elicit equilibrium profiles that attain extremal equilibrium payoff pairs. The only behavior consistent with equilibrium at w_* is mutual exertion of effort, followed by a punishment phase upon the occurrence of a demand shock that yields an expected payoff of $w_*^i - 5/(8\gamma)$ for each player $i = 1, 2$. For $\gamma > 1/3$, this can be achieved by a forgiving grim-trigger profile as illustrated in Figure 4, but other forms of punishment are possible in equilibrium. For lower values of γ, the demand shocks arrive less

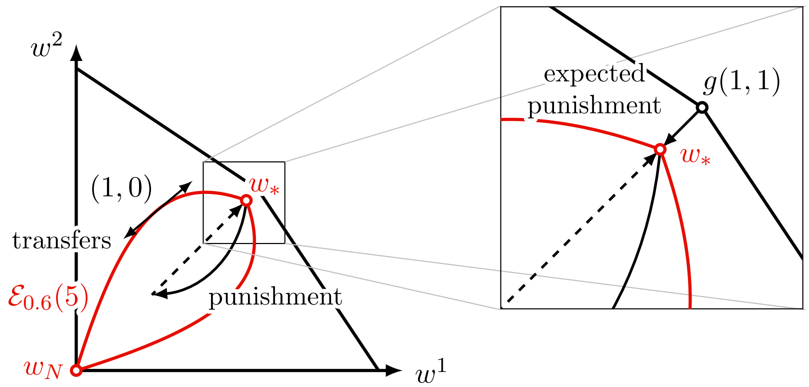

Figure 4: The equilibrium payoff set $\mathcal{E}_{0.6}(5)$ is shown together with the path of the continuation value of a forgiving grim-trigger profile attaining w_*. The solid arrow indicates the minimal amount of value that has to be burnt upon the arrival of a demand shock to enforce mutual effort. As the punishment phase progresses, the players' continuation value increases along the dashed arrow until it is equal to w_* again. The zoom-in in shows that w_* lies below the stage-game payoff pair $g(1,1)$ precisely by the minimum expected value burnt by the punishment. On curved parts of the boundary, one-sided effort is enforced mostly through tangential value transfers.

frequently, hence more severe punishments are necessary to deter deviations. The only sufficiently strong punishment for $\gamma = 1/3$ is a permanent reversion to the static Nash profile, i.e., w_* is attainable only by an unforgiving grim-trigger profile. For $\gamma < 1/3$, no feasible punishment deters shirking by either player, hence w_* is not attainable in equilibrium anymore.

On the curved parts of the boundary, equilibrium behavior prescribes one-sided exertion of effort; see Figure 4. The action profiles of one-sided effort are enforced by transferring value between players tangentially to the set. The continuation value locally remains on the boundary until w_* or w_N is reached or a demand shock occurs. By concatenating one-sided effort with a forgiving grim trigger profile when w_* is reached, it follows that any extremal equilibrium payoff can be attained by the equilibrium of one-sided effort, followed by a grim-trigger profile or permanent Nash reversion. Since any payoff in the interior of $\mathcal{E}(r)$ can be attained by local play of one-sided effort until the continuation value reaches the boundary and hence w_* or w_N, *any* equilibrium payoff can be attained by the equilibrium of one-sided effort, followed by a grim-trigger profile or permanent Nash reversion.

4 Enforceability and self-generation

As in any game with public monitoring, intertemporal incentives are tied to the public signal. We thus start this section by stating the dependence of the continuation value on the public signal, described by the following stochastic differential equation.

Lemma 4.1. *For a two-dimensional process W and a pure strategy profile A, the following are equivalent:*

(i) W is the continuation value of A.

(ii) W is a bounded semimartingale which satisfies for $i = 1, 2$ that

$$\begin{aligned} \mathrm{d}W_t^i &= r\big(W_t^i - g^i(A_t)\big)\,\mathrm{d}t + r\beta_t^i\,\big(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t\big) \\ &\quad + r\sum_{y\in Y}\delta_t^i(y)\,\big(\mathrm{d}J_t^y - \lambda(y|A_t)\,\mathrm{d}t\big) + \mathrm{d}M_t^i \end{aligned} \tag{2}$$

for a martingale M^i (strongly) orthogonal to Z and all J^y with $M_0^i = 0$, predictable processes β^i and $\delta^i(y)$ for $y \in Y$, satisfying $\mathbb{E}_{Q_T^A}\Big[\int_0^T |\beta_t^i|^2\,\mathrm{d}t\Big] < \infty$ and $\mathbb{E}_{Q_T^A}\Big[\int_0^T |\delta_t^i(y)|^2 \lambda(y\,|A_t)\,\mathrm{d}t\Big] < \infty$ for any $T \geq 0$.

The first term in (2) describes the expected movement of the continuation value. It points away from the expected flow payoff: if $g^i(A) > W^i$, 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 and third term in (2) characterize the exposure of the continuation value to the public signal. The exposure to the continuous component is linear with sensitivity $r\beta^i$. Arrival of abrupt information changes the continuation value discontinuously, where $r\delta^i(y)$ are the impacts on player i's continuation value when an event of type y occurs. The second and third term are martingale increments, which means that their impact on the continuation value averages out in expectation. Nevertheless, the exposure to the public signal is important for the provision of incentives as we shall see below. Finally, M captures moves in the continuation value due to public randomization: they average out in expectation and M is 0 if players do not use public randomization.

In discrete-time games, incentives are provided by a continuation promise that maps the outcome of the public signal to a promised continuation payoff for every player; see, for example, Abreu, Pearce, and Stacchetti [3]. The representation in (2) shows that in continuous-time games, such a continuation promise has to be provided via the sensitivities β and δ to the public signal.

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

$$g^i(a) + \beta^i\mu(a) + \delta^i\lambda(a) \geq g^i\big(\tilde{a}^i, a^{-i}\big) + \beta^i\mu\big(\tilde{a}^i, a^{-i}\big) + \delta^i\lambda\big(\tilde{a}^i, a^{-i}\big). \tag{3}$$

We say such a pair (β, δ) *enforces* a. A continuation promise (β, δ) *strictly enforces* a if (3) holds with strict inequality for both players. A strategy profile is *enforceable* if and only if it takes values in enforceable action profiles almost everywhere.

When a player deviates from a given strategy profile, his deviation affects the expected flow payoff as well as the distribution of the public signal. Condition (3) guarantees that for a given continuation promise, the sum of the immediate benefits from the change in flow payoff and the promised future benefits from the change in the signal's distribution are non-positive in expectation. If players keep their continuation promise (β, δ), then β and δ are the sensitivities of the players' continuation value to the public signal in (2). A strategy profile is then a PPE if and only if promises are kept and deviations are deterred. This is formalized in the following lemma.

Lemma 4.3. *A strategy profile A is a PPE if and only if (β^1, β^2) and $\left(\delta^1(y), \delta^2(y)\right)_{y \in Y}$ related to $W(A)$ by (2) enforce A.*

Lemmas 4.1 and 4.3 motivate how we construct equilibrium profiles in continuous time—as the solution to (2) subject to the enforceability constraint in (3). To do so, we use the fact that repeated 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 equilibrium payoff set at all times. This property is known as self generation. In our setting, it is formalized as follows.

Definition 4.4. A payoff set $\mathcal{W} \subset \mathbb{R}^2$ is called *self-generating* if for every payoff pair $w \in \mathcal{W}$, there exists a solution (W, A, β, δ, M) to (2) such that (β, δ) enforces A a.e., $W_0 = w$ a.s., and $W_\tau(A) \in \mathcal{W}$ a.s. for every stopping time τ.

Remark 4.1. The stochastic differential equation (2) does not in general admit strong solutions, that is, it may not be possible to solve (2) for a fixed Brownian motion Z and a fixed set $(J^y)_{y \in Y}$ of Poisson processes. In the proofs that are contained in the appendices, we use weak solutions to (2), in which the Brownian motion, the Poisson 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 A.2 for details.

Similarly as in discrete time, the equilibrium payoff set is the largest bounded self-generating set. For the sake of reference, we state this property as a lemma.

Lemma 4.5. *The set $\mathcal{E}(r)$ is the largest bounded self-generating payoff set.*

The characterization of the equilibrium payoff set as the largest bounded self-generating set allows us to construct equilibria using a stochastic control approach. Because the equilibrium payoff set is self-generating, the continuation value of a PPE has to remain within the set at all times. At the boundary, the law of motion given in (2) thus places certain restrictions on admissible continuation promises; see Figure 5. Sections 5 and 6 illustrate how those restrictions give rise to a local description of the boundary. The arguments in those sections are mostly geometric in nature and are to be taken as a motivation as well as an outline of the main result's proof. All the proofs are contained in the appendices.

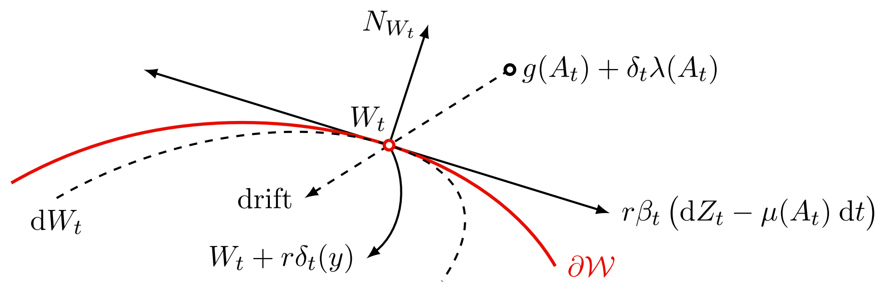

Figure 5: For continuation value W to remain in $\mathcal{W}$, it is necessary that intertemporal incentives satisfy informational restrictions (I1)–(I3) at the boundary.

5 Iteration over arrival of abrupt information

To formalize the informational restrictions at the boundary, we introduce the following notation. For a convex set $\mathcal{W}$ and any payoff pair $w \in \partial\mathcal{W}$, denote by $\mathcal{N}_w(\mathcal{W}) := \left\{N \in S^1 \;\middle|\; N^\top(w - v) \geq 0 \text{ for all } v \in \mathcal{W}\right\}$ the set of all outer-pointing normal vectors to $\partial\mathcal{W}$ at w, where the unit circle S^1 is the set of all directions. If the boundary is continuously differentiable at w, the normal vector is unique and we denote it by N_w. The restrictions on the continuation promise (β, δ) used to provide incentives at the boundary of a self-generating set $\mathcal{W}$ are the following:

(I1) Inward-pointing drift: $N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$ for any $N \in \mathcal{N}_w(\mathcal{W})$,

(I2) Tangential volatility: $N^\top\beta = 0$ for any $N \in \mathcal{N}_w(\mathcal{W})$,

(I3) Jumps within the set: $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$.

Note that restriction (I2) is necessary because Brownian motion has unbounded variation. If incentives were provided in any non-tangential direction, the continuation value would escape $\mathcal{W}$ immediately; see Figure 5 for a visualization.

When information is Brownian, only local information about the boundary is necessary to describe the equilibrium payoff set. Crucial in this regard is that Brownian information arrives continuously, i.e., only the informational restrictions (I1) and (I2) are observed. These restrictions depend on the geometry of the equilibrium payoff set only through the normal vector N_w at w, which gives rise to an explicit description of the boundary through an ordinary differential equation in the state (w, N_w). When information arrives discontinuously as well, such a local description is no longer possible as (I3) is a global restriction involving the precise shape of the equilibrium payoff set. A local description of $\mathcal{E}(r)$ using restrictions (I1)–(I3) thus involves $\mathcal{E}(r)$ itself, creating a non-trivial fixed-point problem. We solve it with an iteration over the arrival times of infrequent events, where (I3) is relaxed so that continuation payoffs after an event come from a fixed payoff set $\mathcal{W}$.

Definition 5.1. Fix a payoff set $\mathcal{W} \subseteq \mathbb{R}^2$. Let σ_n denote the time of the n^{th} event.

(i) We say that a solution (W, A, β, δ, M) to (2) is $\mathcal{W}$-*enforceable* if $W_{\sigma_1} \in \mathcal{W}$ a.s. and $(\beta_t(\omega), \delta_t(\omega))$ enforces $A_t(\omega)$ for almost every (ω, t) with $0 \leq t < \sigma_1(\omega)$.[8]

(ii) A payoff set $\mathcal{X}$ is $\mathcal{W}$-*relaxed self-generating* if for every $w \in \mathcal{X}$, there exists a $\mathcal{W}$-enforceable solution (W, A, β, δ, M) to (2) with $W_0 = w$ a.s. and $W_t(\omega) \in \mathcal{X}$ for almost every (ω, t) with $0 \leq t < \sigma_1(\omega)$.

(iii) We denote by $\mathcal{B}_r(\mathcal{W})$ the largest bounded $\mathcal{W}$-relaxed self-generating set.[9]

The operator $\mathcal{B}_r$ is a continuous-time extension of the standard set operator in Abreu, Pearce and Stacchetti [3]. Payoff pairs in $\mathcal{B}_r(\mathcal{W})$ can be attained by an enforceable strategy profile with continuation promise W_{σ_1} at time σ_1 in $\mathcal{W}$. The additional requirement due to the continuous-time setting is that before the state transition, the continuation value remain in $\mathcal{B}_r(\mathcal{W})$. Similarly to [3], the operator $\mathcal{B}_r$ is closely related to the concept of self-generation.

Lemma 5.2. *Let $\mathcal{W} \subseteq \mathcal{V}$. If $\mathcal{W}$ is self-generating, then $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$. If $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$, then $\mathcal{B}_r(\mathcal{W})$ is self-generating.*

Since payoff pairs in $\mathcal{B}_r(\mathcal{W})$ can be attained with a continuation promise in $\mathcal{W}$ at time σ_1, the condition $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$ allows us to attain W_{σ_1} with an enforceable strategy profile that remains in $\mathcal{B}_r(\mathcal{W})$ until σ_2, and so on. Because Poisson processes have only countably many jumps, repeating this concatenation argument countably many times yields solutions that remain in $\mathcal{B}_r(\mathcal{W})$ forever. The proof in Appendix A.2 additionally deals with some measurability issues of the concatenation. We obtain the following algorithm to approximate $\mathcal{E}(r)$.

Proposition 5.3. *Let $\mathcal{W}_0 = \mathcal{V}^*$ and $\mathcal{W}_n = \mathcal{B}_r(\mathcal{W}_{n-1})$ for $n \geq 1$. Then $(\mathcal{W}_n)_{n \geq 0}$ is decreasing in the set-inclusion sense with $\bigcap_{n \geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*[10]

This algorithm is similar to the algorithm in Abreu, Pearce and Stacchetti [3]. However, unlike its discrete-time counterpart, we will show in the next section that informational restrictions (I1)–(I3) give rise to an explicit characterization of the boundary of the resulting set in each step. The characterization of the boundary of

[8] The condition $\mathcal{W}_{\sigma_1} \in \mathcal{W}$ a.s. is equivalent to requiring that (I3) holds for a.e. (ω, t) with $0 \leq t < \sigma_1(\omega)$. Indeed, jump times of Poisson processes are totally inaccessible, which means that it is impossible to anticipate a discrete event—the information is truly abrupt. It is thus necessary that $W + r\delta(y) \in \mathcal{E}(r)$ holds $P \otimes$ *Lebesgue*-a.e. to ensure that $W_{\sigma_1} \in \mathcal{W}$ a.s.

[9] Since the union of two $\mathcal{W}$-relaxed self-generating sets is again $\mathcal{W}$-relaxed self-generating, $\mathcal{B}_r(\mathcal{W})$ is well defined as the union of all bounded $\mathcal{W}$-relaxed self-generating payoff sets.

[10] If there are no pure-strategy PPE, then the algorithm described in Proposition 5.3 must converge to the empty set: if it converged to a non-empty fixed point of $\mathcal{B}_r$, that fixed point would be self-generating by Lemma 5.2, contradicting the fact that there are no pure-strategy PPE.

$\mathcal{B}_r(\mathcal{W})$ differs depending on whether or not incentives from the continuous component of the public signal are required. If they are, the boundary is characterized by a second-order ordinary differential equation, called the *optimality equation*. If incentives from the continuous component are not required, the boundary is characterized by a first-order differential equation, which we call the *abrupt-information optimality equation*. We conclude this section by stating conditions, under which the optimality equation is locally Lipschitz continuous so that it admits a classical solution.[11]

Definition 5.4. For any $a \in \mathcal{A}$, let $M^i(a)$ denote the $d \times (|\mathcal{A}^i| - 1)$-dimensional matrix, whose column vectors are given by $\mu(\tilde{a}^i, a^{-i}) - \mu(a)$, for $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$. An action profile a is *pairwise identifiable* if $\operatorname{span} M^1(a) \cap \operatorname{span} M^2(a) = \{0\}$.

Assumption 2. Every enforceable action profile is pairwise identifiable.

Similarly to Sannikov [22], pairwise identifiability guarantees that the optimality equation is locally Lipschitz continuous in non-coordinate directions N. In the presence of discontinuous information arrival, we also need a condition that guarantees local Lipschitz continuity in the payoff pair w. Let Ψ_a denote the set of all incentives δ, for which there exists β such that (β, δ) enforces a.

Assumption 3. For $i = 1, 2$, the minmax profile $\underline{a}_i$ against player i and the global maximizer $\bar{a}_i$ of g^i are enforceable and $0 \in \operatorname{int} \Psi_a$ for every enforceable action profile a.[12] Moreover, at least one of the following conditions is satisfied:

(i) Static best replies to $\underline{a}_i^{-i}$ and $\bar{a}_i^{-i}$ are unique.

(ii) $\operatorname{span} M^1(a)$ and $\operatorname{span} M^2(a)$ are orthogonal for $a \in \{\underline{a}_1, \bar{a}_1, \underline{a}_2, \bar{a}_2\}$.

Assumption 3 provides a condition, under which local changes in jump incentives can be compensated with incentives from the continuous component of the public signal. Together with Assumption 2, jump incentives can then be compensated with tangential transfers in any non-coordinate direction. This is sufficient to guarantee local Lipschitz continuity of the optimality equation in non-coordinate directions. Such a local exchange of incentives in coordinate directions $\pm e_i$ is not possible, however, unless the enforced action profile involves a best reply by player i. Assumptions 3.(i) and 3.(ii) make sure that $\partial\mathcal{B}_r(\mathcal{W})$ is well-behaved in coordinate directions.

[11] If the public signal has no continuous component or if the continuous component is completely uninformative, then no part of $\partial\mathcal{B}_r(\mathcal{W})$ is described by the optimality equation, hence no additional assumptions are necessary. If the continuous component is informative, Assumptions 2 and 3 guarantee local Lipschitz continuity of the optimality equation almost everywhere. Without those assumptions, $\partial\mathcal{B}_r(\mathcal{W})$ would be characterized by a differential inclusion instead.

[12] A sufficient condition is an *individual full rank* condition for each action profile: if $M^i(a)$ has rank $|\mathcal{A}^i| - 1$, then (3) can be solved with equality for any δ. If each $M^i(a)$ has individual full rank, then the characterization of Theorem 6.10 is valid for any set $\mathcal{B}_r(\mathcal{W})$ even if $\mathcal{B}_r(\mathcal{W}) \not\subseteq \mathcal{W}$.

6 Characterization of equilibrium payoffs

In Section 4, we motivated the construction of equilibrium profiles as $\mathcal{W}$-enforceable solutions to (2). Due to the iterative procedure in Proposition 5.3, it is sufficient to construct such solutions up until the occurrence of the first infrequent event. The obvious difficulty in the characterization of $\mathcal{B}_r(\mathcal{W})$ is that we do not know a priori in which set the continuation value has to remain until the arrival of the first event. However, since any $\mathcal{W}$-relaxed self generating payoff sets is contained in $\mathcal{B}_r(\mathcal{W})$, we can start by constructing the simplest of such sets: singletons.

6.1 Stationary payoffs

Definition 6.1. A payoff pair w is $\mathcal{W}$-*stationary* if there exist a and δ_0 such that $(0, \delta_0)$ enforces a, $w = g(a) + \delta_0\lambda(a)$, and $w + r\delta_0(y) \in \mathcal{W}$ for every $y \in Y$.

At a stationary payoff, intertemporal incentives are provided only through punishments/rewards after the occurrence of infrequent events. Since the occurrence of events is independent from the history of the public signal, non-occurrence of an event does not make its occurrence more likely or more imminent. Thus, if players go unpunished/unrewarded for an infinitesimal period of time $\mathrm{d}t$, their continuation value increases/decreases by $r\delta_0\lambda(a)\,\mathrm{d}t$. At a stationary payoff, this change in continuation value is precisely offset by the impact $r(w-g(a))\,\mathrm{d}t$ from over-/underconsumption of the instantaneous payoff rate relative to the continuation value. It follows from (2) that the continuation value of the locally constant strategy profile $A \equiv a$ enforced by $(\beta, \delta) \equiv (0, \delta_0)$ remains in w until the occurrence of the first infrequent event, at which point the continuation value jumps to $\mathcal{W}$. We immediately obtain the following result for the set $\mathcal{S}_r(\mathcal{W})$ of all $\mathcal{W}$-stationary payoffs.

Lemma 6.2. *For any $\mathcal{W}$-stationary payoff w, the singleton $\{w\}$ is $\mathcal{W}$-relaxed self-generating. In particular, $\mathcal{S}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$.*

Remark 6.1. Note that any static Nash payoff pair in $\mathcal{W}$ is $\mathcal{W}$-stationary without the provision of incentives. In particular, $\mathcal{V}^N \cap \mathcal{W} \subseteq \mathcal{S}_r(\mathcal{W})$.

The set of stationary payoffs can computed fairly easily as intersections, scalings, and projections of convex sets. We illustrate it here in the partnership example of Section 3 as it allows a nice geometric construction: When there is only one type of events, the stationarity condition implies $w + r\delta_0 = g(a) + \delta_0\lambda(a) + r\delta_0 \in \mathcal{W}$. The set of all payoffs that can be reached from a stationary payoff can be computed as

$$\big(g(a) + \Psi_a^0(\lambda(a) + r)\big) \cap \mathcal{W},$$

where $\Psi_a^0 := \{\delta \mid (0, \delta) \text{ enforces } a\}$. In order to get the stationary payoffs themselves, we simply shrink this set towards $g(a)$ by a factor $\frac{\lambda(a)}{\lambda(a)+r}$; see the left panel Figure 6.

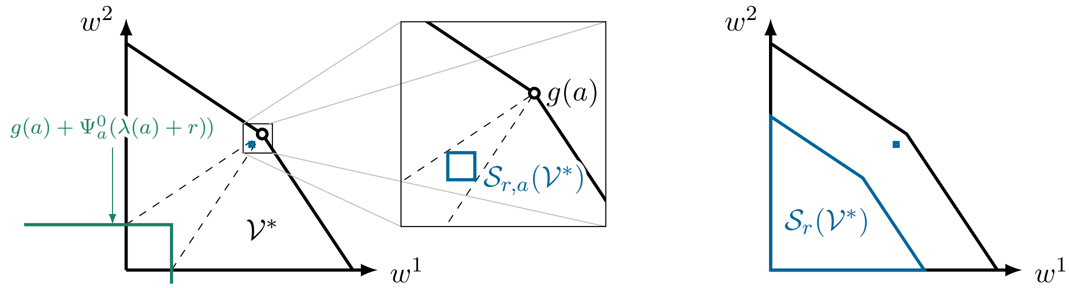

Figure 6: The left panel shows the construction of payoffs $\mathcal{S}_{r,a}(\mathcal{V}^*)$ that can be attained with local play of $a = (1, 1)$ in the partnership example. The right panel shows all $\mathcal{V}^*$-stationary payoffs.

Repeating this construction for all action profiles a yields the set of stationary payoff. The right panel of Figure 6 shows the set $\mathcal{S}_r(\mathcal{V}^*)$ of all $\mathcal{V}^*$-stationary payoffs in the partnership game. If there is more than one type of events, an a analogue construction is carried out in the incentive space rather than in the payoff space; see Section 8.1.

6.2 Incentives provided through Brownian information

Outside the set of stationary payoffs, the continuation value of an enforceable strategy profile does not remain locally constant. Thus, informational restrictions (I1)–(I3) must hold at the boundary of $\mathcal{B}_r(\mathcal{W})$. Since the location of the boundary is a priori unknown, we abstract away from (I1)–(I3) with the following definition.

Definition 6.3. For a payoff pair $w \in \mathbb{R}^2$, a direction $N \in S^1$, discount rate $r > 0$, and a payoff set $\mathcal{W}$, we say that a continuation promise (β, δ) from the set

$$\Xi_a(w, N, r, \mathcal{W}) := \left\{ (\beta, \delta) \;\middle|\; \begin{array}{l} (\beta, \delta) \text{ enforces } a, N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0, \\ N^\top \beta = 0, \text{ and } w + r\delta(y) \in \mathcal{W} \text{ for every } y \in Y \end{array} \right\}$$

restricted-enforces a. An action profile $a \in \mathcal{A}$ is *restricted-enforceable* for $(w, N, r, \mathcal{W})$ if the set $\Xi_a(w, N, r, \mathcal{W})$ of all such continuation promises is non-empty.

We first illustrate how play of a restricted-enforceable strategy profile gives rise to an expression for the curvature of the boundary of $\mathcal{B}_r(\mathcal{W})$. While the presence of abrupt information creates many technical challenges that we address in Appendices B and C, this step of the argument is similar to Sannikov [22] because the infinite-variation property of Brownian information causes the curvature.

Consider, for a moment, that a $\mathcal{W}$-enforceable solution W to (2) remains on a continuously differentiable curve $\mathcal{C}$ with curvature κ before the first infrequent event occurs.[13] Then Brownian information can be used to provide incentives only via

[13] 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.

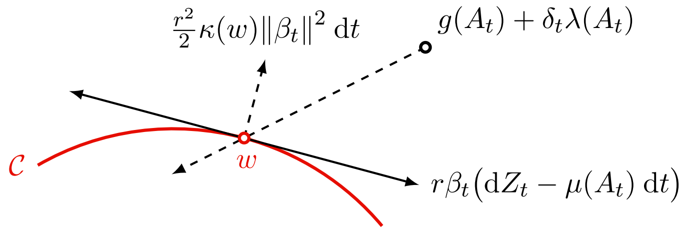

Figure 7: The infinitesimally quick oscillation in the continuation value due to unbounded variation of value transfers causes the continuation value to drift away from the curve in orthogonal direction.

tangential value transfers as illustrated in Figure 7. 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 magnitude of the outward drift is given by Itō's formula and equals $\frac{r^2}{2}\kappa(w)\|\beta_t\|^2$. For the continuation value to remain on the curve, this outward drift is precisely counteracted by the inward-pointing drift so that at w,

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

Finally, 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 6.4. *Let $\mathcal{C}$ be a continuously differentiable curve oriented by $w \mapsto N_w$ with:*

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

$$\kappa(w) = \frac{2N_w{}^\top\Big(g\big(a^*(w)\big) + \delta^*(w)\lambda\big(a^*(w)\big) - w\Big)}{r\|\beta^*(w)\|^2}. \tag{4}$$

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

Then $\mathcal{C} \subseteq \mathcal{B}_r(\mathcal{W})$ and the solution W to (2) 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 an event occurs.

As motivated before the lemma, such a solution W to (2) stays on the curve $\mathcal{C}$ with curvature (4) until an endpoint of $\mathcal{C}$ is reached or an infrequent event occurs. Since $\big(\beta, \delta\big) \in \Xi_A\big(W, N_W, r, \mathcal{W}\big)$, the continuation promise (β, δ) enforces strategy profile A with punishments/rewards that come from the set $\mathcal{W}$ after the occurrence an infrequent event. By definition of $\mathcal{B}_r(\mathcal{W})$, an end point v of the curve can be attained by a $\mathcal{W}$-enforceable solution W' to (2) that remains in $\mathcal{B}_r(\mathcal{W})$ until the arrival of the first event, at which point the continuation value jumps to $\mathcal{W}$. By concatenating the two solutions, we see that $\mathcal{C} \cup \mathcal{B}_r(\mathcal{W})$ is relaxed self-generating. Because $\mathcal{B}_r(\mathcal{W})$ is the largest bounded relaxed self-generating set, it follows that $\mathcal{C}$ must be contained in $\mathcal{B}_r(\mathcal{W})$. Details of this proof are contained in Appendix C.

The above construction works only if Brownian information is used to provide incentives, that is, if $\beta \neq 0$. Where it is used, Brownian information makes the boundary of $\mathcal{B}_r(\mathcal{W})$ smooth and the characterization via its curvature will follow from an elaborate application of Lemma 6.4. Let us begin by formalizing where incentives must be provided at least partially through the Brownian information. We do so by defining its complement, where Brownian information is not needed.

Definition 6.5. Denote by $\Psi_a^0(w, r, \mathcal{W})$ the set of all incentives $\delta \in \Psi_a^0$, for which $w + r\delta(y) \in \mathcal{W}$ for every event $y \in Y$. Define

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

Due to informational restriction (I1), whether or not Brownian information is needed at a boundary point depends not only on the location of the payoff pair but also on the direction of the outward normal vector. For any convex set $\mathcal{X}$, 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 $\mathcal{B}_r(\mathcal{W})$ inside and outside of $\Gamma(r, \mathcal{W})$ when referring to boundary points $w \in \partial\mathcal{B}_r(\mathcal{W})$, for which $(w, N) \in \Gamma(r, \mathcal{W})$ for some and no outward normal N, respectively.

Lemma 6.6. *If $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$, then outside of $\Gamma(r, \mathcal{W})$ the boundary $\partial\mathcal{B}_r(\mathcal{W})$ is continuously differentiable such that for almost every $(w, N_w) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \setminus \Gamma(r, \mathcal{W})$, the curvature $\kappa(w)$ is given by the optimality equation*

$$\kappa(w) = \max_{a \in \mathcal{A}} \max_{(\beta,\delta) \in \Xi_a(w, N_w, r, \mathcal{W})} \frac{2{N_w}^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\beta\|^2}, \tag{5}$$

where we set $\kappa(w) = 0$ if the maxima are taken over empty sets.[14]

Looking at the expression in the right-hand side of (5), we notice that the curvature of the boundary arises from (4) by taking the maximum over all restricted-enforceable action profiles and all restricted-enforcing continuation promises. The reason can be summarized as follows. Suppose that a solution $\mathcal{C}$ to (5) starting at (w, N_w) falls into the interior of $\mathcal{B}_r(\mathcal{W})$, then a solution $\mathcal{C}'$ with initial conditions

[14] As we show in Appendix B, $\kappa(w) = 0$ only if the continuous component of the public signal is completely uninformative.

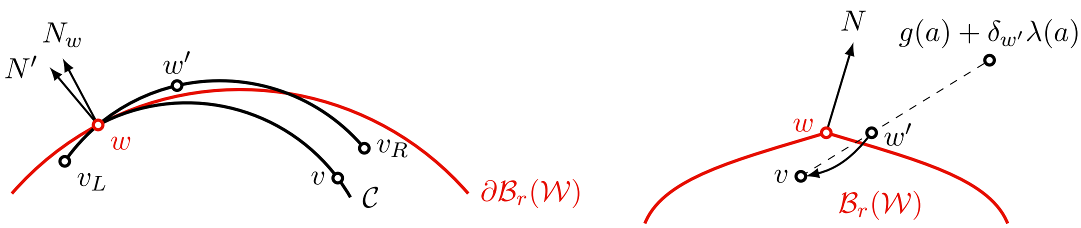

Figure 8: If a solution $\mathcal{C}$ to (5) starting at (w, N_w) falls into the interior of $\mathcal{B}_r(\mathcal{W})$, then a solution $\mathcal{C}'$ with initial conditions (w, N') for a slight rotation N' of N_w would leave and re-enter $\mathcal{B}_r(\mathcal{W})$. Lemma 6.4 then implies that $w' \in \mathcal{B}_r(\mathcal{W})$, which is a contradiction. The right panel illustrates that if both conditions of Lemma 6.8 hold simultaneously at $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma(r, \mathcal{W})$, then they also hold at (w', N) for w' outside of $\mathcal{B}_r(\mathcal{W})$, leading to a contradiction.

(w, N') for a slight rotation N' of N_w would leave and re-enter $\mathcal{B}_r(\mathcal{W})$; see Figure 8. That is, it would attain a payoff pair w' strictly outside of $\mathcal{B}_r(\mathcal{W})$. But Lemma 6.4 implies that $w' \in \mathcal{B}_r(\mathcal{W})$, which is a contradiction. This argument requires continuity of solutions to (5) in initial conditions, which we address in Appendix B. On the other hand, a solution $\mathcal{C}$ to (5) cannot escape $\mathcal{B}_r(\mathcal{W})$ 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. The details of this proof are contained in Appendix C.

6.3 Incentives provided through abrupt information exclusively

Without the use of Brownian information, players are punished/rewarded only at the occurrence of discrete events with continuation value in $\mathcal{W}$. This is similar to the decomposition of payoffs in discrete time and we shall adopt the same terminology.

Definition 6.7. Consider $w \in \partial\mathcal{B}_r(\mathcal{W})$ and a set of action profiles $\mathcal{A}_w \subseteq \mathcal{A}$.

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

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

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

Referring to the definition of $\Psi^0_a(w, r, \mathcal{W})$ in Definition 6.5, we note that the defining characteristics of a decomposable payoff pair are the following: incentives are provided through the abrupt information only, the drift rate points towards the interior of $\mathcal{B}_r(\mathcal{W})$, and the continuation payoff after an infrequent event is in $\mathcal{W}$. The following lemma establishes that at most one of these defining conditions holds strictly.

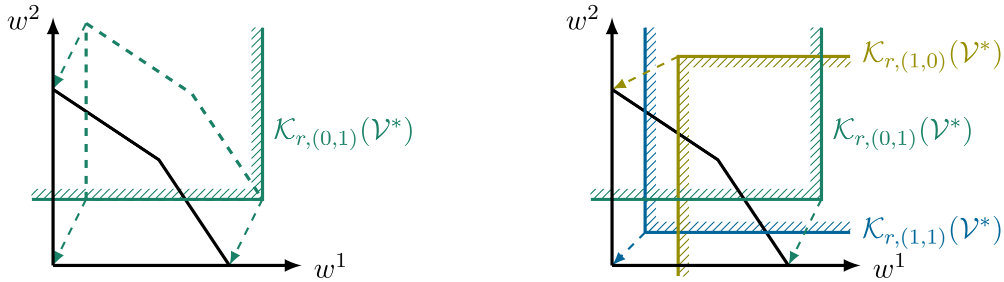

Figure 9: In the partnership game, mutual effort is enforceable by punishing both players by at least $r/8\gamma$. One-sided effort is enforceable by punishing the working player by at least $r/4\gamma$ and punishing the shirking player by at most $r/8\gamma$. The arrows show these extremal punishments for $r = 5\gamma$. The sets $\mathcal{K}_{r,a}(\mathcal{V}^*)$ are infinite extensions of translates of $\mathcal{V}^*$ by those extremal punishments.

Lemma 6.8. *For any* $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma(r, \mathcal{W})$*, it is impossible that there exist* $a \in \mathcal{A}$ *and* $\delta \in \Psi_a^0(r, w, \mathcal{W})$ *such that both of the following conditions hold:*

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

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

Proof. Fix $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma(r, \mathcal{W})$ and suppose that there exist a and δ that satisfy Conditions (i) and (ii). For any w' sufficiently close to w, Conditions (i) and (ii) are satisfied by a and $\delta_{w'}$ with $\delta_{w'}(y) := \delta(y) + \frac{w-w'}{r}$ for every event y. In particular, $\delta_{w'} \in \Psi_a^0(w', r, \mathcal{W})$. Choose $w' \notin \operatorname{cl} \mathcal{B}_r(\mathcal{W})$ close enough to w such that the straight line segment through $x := g(a) + \delta_{w'}\lambda(a)$ and w' contains a point $v \in \operatorname{int} \mathcal{B}_r(\mathcal{W})$; see Figure 9. Let W be a solution to (2) with $W_0 = w'$, $A \equiv a$, $\beta \equiv 0$, $\delta \equiv \delta_{w'}$, and $\mathrm{d}M_t = (v - w)(\mathrm{d}J'_t - \lambda' \,\mathrm{d}t)$, where J' is a Poisson process independent of the public signal with intensity $\lambda' = \|v - x\|/\|v - w\|$. By construction, W remains in w' until it jumps either to v at the first jump time of J' or to $\mathcal{W}$ when an event y occurs since $\delta_{w'} \in \Psi_a^0(w', r, \mathcal{W})$. Concatenating (W, A, β, δ, M) with a $\mathcal{W}$-enforceable solution to (2) that attains $v \in \mathcal{B}_r(\mathcal{W})$ shows that $\{w'\} \cup \mathcal{B}_r(\mathcal{W})$ is relaxed self-generating and hence $w' \subseteq \mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. This is a contradiction. Lemma A.2 in the appendix contains the details to the concatenation procedure. □

Lemma 6.8 tells us that there are two kinds of boundary payoffs in $\Gamma(r, \mathcal{W})$. On the one hand, there are boundary payoffs that are not strictly decomposable, hence (w, N) has to satisfy the abrupt-information optimality equation

$$N^\top w = \max_{a \in \mathcal{A}} \max_{\delta \in \Psi_a^0(w, r, \mathcal{W})} N^\top\big(g(a) + \delta\lambda(a)\big). \tag{6}$$

Note that (6) is in some sense the limiting ODE of (5) as (w, N) approaches $\Gamma(r, \mathcal{W})$, corresponding to the case where both numerator and denominator of (5) are 0. Only for a solution to (6), does the expression for the curvature in (5) not explode for (w, N) in $\Gamma(r, \mathcal{W})$. 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 $\mathcal{W}$ after at least one type of event $y \in Y$. Thus such a payoff pair w must lie on the boundary of the set

$$\mathcal{K}_{r,a}(\mathcal{W}) := \{w \mid \exists\, \delta \in \Psi_a^0(w,r,\mathcal{W})\}.$$

In games with a single type of events, the set $\mathcal{K}_{r,a}(\mathcal{W})$ is an infinite extension of a translate of $\mathcal{W}$ by the extremal punishments/rewards necessary to enforce a. This is visualized for the partnership example of Section 3 in the right panel of Figure 9. We discuss the construction of the sets $\mathcal{K}_{r,a}(\mathcal{W})$ in the general case in Section 8.2.

The implication of Lemma 6.8 is particularly pungent for the decomposition of corners of $\mathcal{B}_r(\mathcal{W})$, i.e., boundary points with more than one outward normal vector. Condition (ii) can be violated for all outward normal vectors only if $w = g(a)+\delta\lambda(a)$, i.e., if w is stationary. Thus, corners of $\mathcal{B}_r(\mathcal{W})$ are either stationary or in the set $\mathcal{K}_r(\mathcal{W}) := \bigcup_{a\in\mathcal{A}} \partial\mathcal{K}_{r,a}(\mathcal{W})$. This discussion is formalized in the following proposition.

Proposition 6.9. *For any $\mathcal{W}$ with $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$ and any $(w,N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma(r,\mathcal{W})$ with non-stationary w, exactly one of the following conditions holds:*

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

$$\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})\big), \tag{7}$$

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

(ii) (w,N) satisfies (6) and either of the following conditions hold:

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

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

where a_* is the action profile attaining the maximum in (6).

Proposition 6.9 implies that within $\Gamma(r,\mathcal{W})$, the boundary of $\mathcal{B}_r(\mathcal{W})$ can have three kinds of continuously differentiable or smooth line segments and three kinds of corners. Consider a smooth segment $\mathcal{C} \subseteq \partial\mathcal{B}_r(\mathcal{W})$ with $\mathcal{N}_{\mathcal{C}} \subseteq \Gamma(r,\mathcal{W})$. 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 $\partial\mathcal{K}_{r,a}(\mathcal{W})$. If no action profile strictly decomposes w, then the drift rate of any enforceable strategy profile attaining w must be parallel to the boundary. According to Proposition 6.9, $\mathcal{C}$ is thus either stationary or a solution to (6) 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 $\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})$ satisfying (7). Note that contrary to smooth line segments, corners are not necessarily decomposable by a single action profile; see the left panel of Figure 10. If w cannot be strictly decomposed, then it is either stationary or the starting point of a continuously differentiable solution to (6) as in case (ii.b). Figure 10 illustrates these different boundary segments.

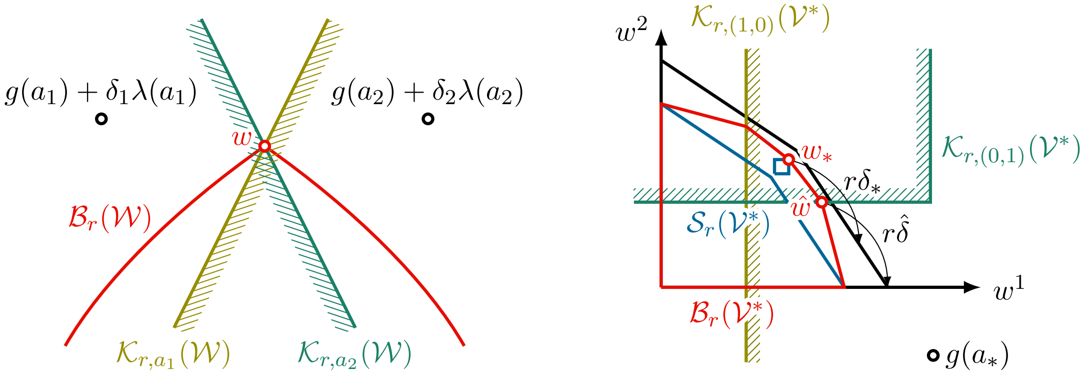

Figure 10: The left panel illustrates case (i) of Proposition 6.9: w is a corner of $\mathcal{B}_r(\mathcal{W})$ that is minimally decomposed by $\mathcal{A}_w = \{a_1, a_2\}$ with strictly inward-pointing drift. While w may not be a corner of either $\mathcal{K}_{r,a_1}(\mathcal{W})$ or $\mathcal{K}_{r,a_2}(\mathcal{W})$, it is a corner of $\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})$ satisfying (7). The right panel shows $\mathcal{B}_r(\mathcal{V}^*)$ for $r = 5$ in the partnership game when $\gamma = 1$, i.e., only the discrete events are informative. Between w_* and $\hat{w}$, the boundary is a solution to (6), satisfying case (ii.a) of Proposition 6.9 for maximizing action profile $a_* = (0, 1)$. Since player 2 has to be punished by at least $5/4$ payoff units in order to play a_*, extremal incentives δ_* and $\hat{\delta}$ maximize the reward for player 1 subject to $r\delta^2 = -5/4$. Note that $\delta_*^1 > \hat{\delta}^1$, hence $g(a_*)+\delta_*\lambda(a_*)$ is further "east" than $g(a_*)+\hat{\delta}\lambda(a_*)$, showing that this segment is strictly curved. Finally, $\hat{w}$ satisfies case (ii.b) of Proposition 6.9.

6.4 Characterization of $\mathcal{E}(r)$

Aggregating the results from this section, the following is a complete characterization of $\mathcal{B}_r(\mathcal{W})$. The proof of Lemma 6.6 is contained in Appendix C, building on the regularity of the optimality equation shown in Appendix B. The proofs of Proposition 6.9 and Theorem 6.10 are contained in Appendices D and E, respectively. Auxiliary results and technical bounds are deferred to Appendix F.

Theorem 6.10. *Let $\mathcal{W} \subseteq \mathcal{V}^*$ be convex and compact with $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$. Then $\mathcal{B}_r(\mathcal{W})$ is the largest closed convex subset $\mathcal{X}$ of $\mathcal{V}^*$ that contains $\mathcal{S}_r(\mathcal{W})$ such that:*

(i) Outside of $\mathcal{S}_r(\mathcal{W}) \cup \mathcal{K}_r(\mathcal{W})$, the boundary $\partial\mathcal{X}$ is continuously differentiable and $(w, N_w) \in \mathcal{N}_\mathcal{X}$ solves (6) inside of $\Gamma(r, \mathcal{W})$ and (5) outside of $\Gamma(r, \mathcal{W})$,

(ii) Every corner $w \in \partial\mathcal{X}$ is either stationary, minimally decomposable by a maximizer a_* of (6) for (w, N) with $N \in \operatorname{ext} \mathcal{N}_w(\mathcal{X})$ and $w \in \partial\mathcal{K}_{r,a_*}(\mathcal{W})$, or minimally and strictly decomposable by some $\mathcal{A}_w$ with $\mathcal{N}_w(\mathcal{X}) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})\big)$.

Even though the curvature is characterized only at almost every point on the boundary, a solution is unique with the additional requirement that it be continuously differentiable. This implies that $\partial\mathcal{B}_r(\mathcal{W})$ is twice continuously differentiable almost everywhere, which is important for the numerical solution of (5) as numerical procedures rely on discretizations. We will elaborate on the numerical implementation in Section 8. Since $\mathcal{B}_r$ preserves compactness due to Theorem 6.10, it follows from Proposition 5.3 that $\mathcal{E}(r)$ is compact. An application of Theorem 6.10 for $\mathcal{W} = \mathcal{E}(r)$ thus provides a fixed-point characterization of $\mathcal{E}(r)$ since $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$.

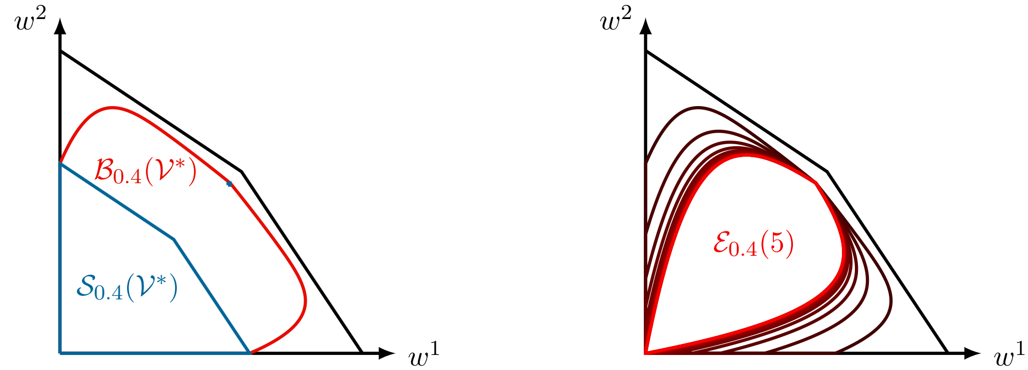

Figure 11: The left panel shows $\mathcal{B}_r(\mathcal{V}^*)$ in the partnership example for $\gamma = 0.4$. The right panel shows the convergence of the algorithm in Proposition 5.3.

7 Discussion

In this section, we discuss the use of public randomization, where and how incentives are provided via abrupt information, as well as how our result relates to Sannikov and Skrzypacz [24] in greater detail.

7.1 Public randomization

By giving players access to a public randomization device, the analysis is simplified for two reasons. First, public randomization allows us to conclude early on that the equilibrium payoff set is convex. Second, if players have a very rich public randomization device available—sufficiently rich for an exact law of large numbers to hold—then players can randomize between pure action profiles quickly enough as if they played a correlated pure action profile in every instant; see Sun [26] for a mathematical description of such a randomization device. This is necessary to attain corners of $\mathcal{B}_r(\mathcal{W})$ that are not decomposable by a single action profile as illustrated in the left panel of Figure 10: if players cannot randomize between a_1 and a_2 sufficiently quickly, then play of either action profile at w would cause the continuation value to escape $\mathcal{B}_r(\mathcal{W})$ immediately. A sufficiently rich randomization device for an exact law of large numbers to hold is also needed to attain boundary payoffs w, where the maximizer in (6) changes and the induced drift rate on both sides of w points towards w. Very similarly to the left panel of Figure 10, play of any pure action profile at w would cause the continuation value to escape $\mathcal{B}_r(\mathcal{W})$ immediately. With sufficiently rich randomization however, players can randomize among the maximizing action profiles to the left and right of w so that the drift rate is 0 and the continuation value is locally constant.

Despite these two important uses of public randomization in the derivation of our result, it can easily be verified whether the equilibrium payoff set $\mathcal{E}(r)$ with public randomization coincides with the equilibrium payoff set $\mathcal{E}_*(r)$ without public randomization. Since $\mathcal{E}_*(r) \subseteq \mathcal{E}(r)$, it is sufficient to find conditions such that any payoff in

$\mathcal{E}(r)$ can be attained without public randomization. Consider first the case, in which there is no continuous component of the public signal (or the continuous component is completely uninformative as in the partnership example when $\gamma = 1$). Then boundary payoffs of $\partial\mathcal{E}(r)$ can be attained without public randomization if and only if:

(i) $(w, N) \in \Gamma(r, \mathcal{W})$ for every boundary payoff w and every normal vector N,

(ii) Every corner is minimally decomposed by a single action profile,

(iii) If the maximizer in (6) changes at a payoff pair w where the boundary is smooth, then the drift rate induced by the maximizing action profile on at least one side of w points away from w.

Note that if Condition (i) is not satisfied, then the boundary of $\mathcal{E}(r)$ is a trivial solution to (5) outside of $\Gamma(r, \mathcal{W})$, i.e., it is a straight line segment. Since no action profile is enforceable with inward-pointing drift by definition of $\Gamma(r, \mathcal{W})$, public randomization must be used at those points. If, additionally, there exists at least one static Nash profile a_N, then any payoff in the interior of $\mathcal{E}(r)$ can be attained by play of a_N until the continuation value reaches the boundary, hence $\mathcal{E}(r) = \mathcal{E}_*(r)$.

If the continuous component of the public signal is informative, then Assumptions 2 and 3 are satisfied. Since any action profile is enforceable without incentives via abrupt information by Assumption 3, any payoff pair in the interior can be attained by an enforceable strategy profile with a continuous continuation value until the boundary is reached. Moreover, since the boundary is strictly curved outside of $\Gamma(r, \mathcal{W})$, any boundary payoff outside of $\Gamma(r, \mathcal{W})$ can be attained without public randomization by Lemma 6.4. Thus, if the continuous component of the public signal is informative, Conditions (ii) and (iii) are necessary and sufficient for $\mathcal{E}(r) = \mathcal{E}_*(r)$.

Note that it is possible that $\mathcal{B}_r(\mathcal{W}_n)$ violates some of the above conditions for every set $\mathcal{W}_n$ in the approximating sequence $(\mathcal{W}_n)_{n\geq 0}$ of Proposition 5.3, yet the limiting equilibrium payoff set satisfies all of them. This is the case in the partnership example for $\gamma = 1$. For every $n > 0$, the set $\mathcal{W}_n$ contains some payoff pair w_n on the x-axis with $w_n^1 > 0$ such that $\partial\mathcal{W}_n$ between w_n and $\partial\mathcal{K}_{r,(1,0)}(\mathcal{W}_{n-1})$ is a straight line segment outside of $\Gamma(r, \mathcal{W})$; see the right panel of Figure 10. Public randomization is thus necessary to attain those payoff pairs. In the limit as $n \to \infty$, however, w_n converges to the static Nash payoff and (6) is satisfied by the static Nash profile on all of $\mathcal{E}(r)$. For all monitoring structures of the partnership game considered in the following section, the above conditions are satisfied and hence $\mathcal{E}(r) = \mathcal{E}_*(r)$.

7.2 Incentives provided via abrupt information

The optimality equation (5) trades off incentives provided via the two channels of information. Because the boundary is strictly curved where it satisfies (5), incentives provided via abrupt information necessarily decrease the numerator of (5) since the continuation value after a discrete event has to lie strictly below the tangent. However, the use of such incentives may be efficient if it sufficiently reduces the tangential transfers that are required, which decreases the denominator of (5). In this section, we analyze the equilibrium payoff set when only the abrupt information is informative as well as the trade-off between the two types of information for three different monitoring structures of the partnership example.

The first monitoring structure is as in Section 3. If the continuous information is completely uninformative, i.e., $\gamma = 1$, then the equilibrium payoff set is entirely contained on the positive diagonal between the static Nash payoff and the highest symmetric stationary payoff; see Figure 3. Even though $\mathcal{B}_r(\mathcal{V}^*)$ contains asymmetric payoffs, at which one-sided effort is $\mathcal{V}^*$-restricted enforceable as illustrated in Figure 10, those disappear as we apply $\mathcal{B}_r$ repeatedly. To understand why that is the case, observe first that $a_* = (0, 1)$ can only contribute to $\mathcal{B}_r(\mathcal{W}_n)$ to the bottom right of the diagonal due to the inward-drift condition. Moreover, because the continuation value after a demand shock has to lie in $\mathcal{W}_n$, any boundary payoff of $\mathcal{B}_r(\mathcal{W}_n)$, in which a_* is played, has to lie in $\mathcal{K}_{r,a_*}(\mathcal{W}_n)$. To enforce a_*, player 2 has to be punished by at least $r/4$ payoff units. To attain payoff pairs as far below the diagonal as possible, it is necessary to punish player 1 as well. However, if those punishments are too severe, player 1 will seek to reduce the number of demand shocks by deviating and exerting effort. Since a deviation by player 2 has a smaller impact on the arrival rate of demand shocks than a deviation by player 1, it follows from (3) that player 1 can be punished at most half as severely as player 2, with extremal punishments given by $r\delta_* = -(r/8, r/4)^\top$. As illustrated in Figure 9, the south-easternmost point of $\mathcal{K}_{r,a_*}(\mathcal{W}_n)$ is thus a translate of the south-easternmost point of $\mathcal{W}_n$ by $-r\delta_*$. As a consequence, $\mathcal{K}_{r,a_*}(\mathcal{W}_n)$ moves $r/\sqrt{8}$ payoff units closer to the diagonal in every iteration of algorithm until, eventually, no payoffs below the diagonal are left, for which the continuation value after a discrete event lies in $\mathcal{W}_n$.

Let us consider a second monitoring structure, in which the demand shocks arrive according to $\lambda_\gamma(a) = \gamma(33 - 16(a^1 + a^2))$ and μ is defined as in Section 3. Demand shocks arrive with the same frequency in both monitoring structures when at least one player exerts effort. However, demand shocks arrive more frequently in monitoring structure 2 when neither player exerts effort so that deviations from a_* by player 2 have a larger impact on demand shocks. The extremal punishments $r\delta_*$ that enforce a_* are now $-(r/8, r/16)^\top$ so that $\mathcal{K}_{r,a_*}(\mathcal{W})$ need not lie further northwest than $\mathcal{W}$. The left panel of Figure 12 shows that $\mathcal{E}_1(5)$ indeed contains asymmetric payoffs under this monitoring structure. The right panel of Figure 12 shows that for $\gamma = 0.8$, the boundary of the equilibrium payoff set changes differentiably from a solution to (5) to a solution to (6) at payoffs in $\partial\mathcal{K}_{5,(0,1)}\big(\mathcal{E}_{0.8}(5)\big)$ and $\partial\mathcal{K}_{5,(1,0)}\big(\mathcal{E}_{0.8}(5)\big)$.

In both monitoring structures, the matrices $M^i(a)$ for $i = 1, 2$ have rank $|\mathcal{A}^i| - 1$ for any action profile a, that is, any action profile has *individual full rank*. Thus, for any incentives δ provided via abrupt information, there exists β such that the

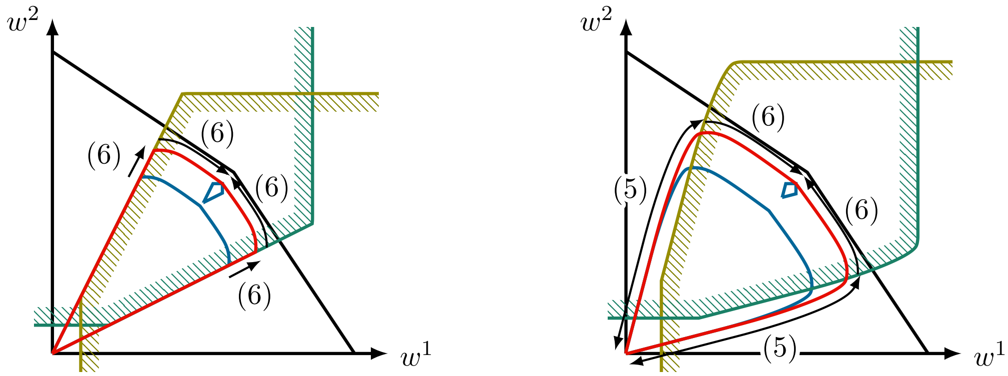

Figure 12: The equilibrium payoff set $\mathcal{E}_\gamma(r)$, the set of stationary equilibrium payoffs $\mathcal{S}_r(\mathcal{E}_\gamma(r))$, as well as $\mathcal{K}_{r,(0,1)}(\mathcal{E}_\gamma(r))$ and $\mathcal{K}_{r,(1,0)}(\mathcal{E}_\gamma(r))$ in monitoring structure 2 are shown for $r = 5$ and $\gamma = 1$ in the left and $\gamma = 0.8$ in the right panel. On any segment of the boundary that satisfies (5), the continuation value of any equilibrium has infinite variation and it may reach either end point of that segment. On segments that satisfy (6), the continuation value of any equilibrium moves deterministically in the direction of the arrows until a demand shock occurs.

enforceability constraints (3) for both players can be solved with equality. Pairwise identifiability additionally implies that there exists such β with $N^\top\beta = 0$ for any non-coordinate direction. One may wonder whether an individual full rank condition or a pairwise identifiability condition is equally powerful if it is imposed for matrices $\Lambda^i(a)$ consisting of column vectors $\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)$ for $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$. From a theoretical perspective, the answer to that is negative because only bounded amounts of value can be transferred or destroyed upon the arrival of a discrete event. While such a condition for $\Lambda^i(a)$ allows us to find sufficient rewards/punishments δ with $N^\top\delta(y)$ for any choice of β, there is no guarantee that the resulting continuation value $w + r\delta(y)$ after an event y lies in the desired set $\mathcal{W}$. Besides, contrary to (I2), informational restriction (I3) does not require the provision of tangential rewards/punishments, hence a pairwise identifiability condition for infrequent events is not as potent as it is for Brownian information. Nevertheless, the ability for player-specific punishments upon the arrival of a discrete event can increase the equilibrium payoff set quite drastically as we illustrate with the following example.

Consider a third monitoring structure, in which there are two types of demand shocks y_1, y_2 with instantaneous intensities $\lambda_\gamma(y_i, a) = \gamma\big(16.5 - 16a^i\big)$ for $i = 1, 2$. The frequency that any event occurs is the same as in the second monitoring structure, but deviations by the two players can be statistically distinguished. Mutual effort is still enforced by punishing the likely deviator i with $5/(8\gamma)$ payoff units after any demand shock y_i occurs. However, since rewards/punishments after event y_i do not affect $-i$'s incentives, it is possible to reward player $-i$ after event y_i so that less total value is destroyed upon the arrival of a demand shock. As a consequence, the highest symmetric equilibrium payoff is much closer to the efficient frontier than when there is only one type of demand shock; see Figure 13. Moreover, there are

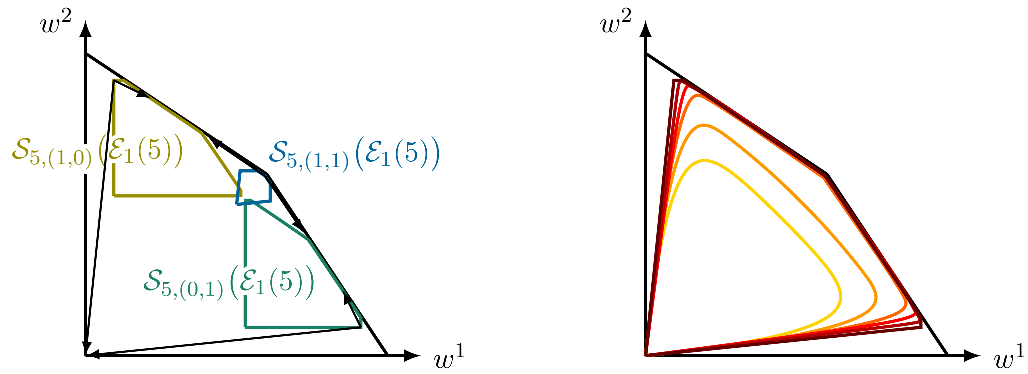

Figure 13: The left panel shows the stationary equilibrium payoffs that can be attained by local play of the three non-static Nash profiles in the partnership game with monitoring structure 3 and $\gamma = 1$. The arrows indicate the rewards/punishments after the two different types of demand shocks for three extremal equilibrium payoff pairs. The right panel shows the equilibrium payoff sets $\mathcal{E}_\gamma(5)$ for γ between 0 (yellow) and 1 (dark red) in intervals of size 0.2.

stationary equilibrium payoffs that support one-sided effort due to the possibility of attaching different continuation values after the two types of demand shocks. In the case $\gamma = 1$, when only the demand shocks are informative, the equilibrium payoff set is the convex hull of stationary payoffs: outside of $\mathcal{S}_r(\mathcal{E}_1(r))$, the boundary is a trivial solution to (6), i.e., those are straight lines through stationary payoffs.

Figure 14 illustrates the optimal incentives provided via demand shocks in the three different monitoring structures. We see that rewards/punishments $\delta(y)$ after event y are provided close to tangentially in order to minimize the expected value that is burnt through these incentives. Whether close to tangential δ provides strong incentives for players not to deviate depends on the monitoring structure. Consider a boundary payoff w, where an asymmetric action profile a is played in equilibrium. Let $\beta(w)$ and $\delta(w)$ denote the maximizers in (5) at w and define by $\rho(w) := \delta^i(w)\Lambda^i(a)/(\delta^i(w)\Lambda^i(a) + \beta^i(w)M^i(a))$ the ratio of incentives that are provided to player i that is exerting effort in a. Figure 14 shows ρ for the three monitoring structures. In monitoring structures 1 and 2, jump incentives are provided almost tangentially, hence ρ is largest where the curvature is small because this simply allows $|\delta^i(y)|$ to be large. Incentives provided in this way are stronger in monitoring structure 2 because mutual shirking is easier to distinguish from one-sided effort than it is in monitoring structure 1. In monitoring structure 3, mutual effort can be enforced more efficiently because player-specific punishments are available. At boundary payoffs w, where one-sided effort by player i is optimal, incentives $\delta(y_i)$ are provided close to tangentially except where the curvature is very large. Where the curvature is large, only small tangential transfers can be provided due to (5), hence incentives via abrupt information are simply necessary even if they burn a lot of value. Moreover, since changes in continuation payoff after an event y_{-i} have no impact on i's incentives, at those points we have $\delta(w; y_{-i}) = 0$ to minimize the value burnt.

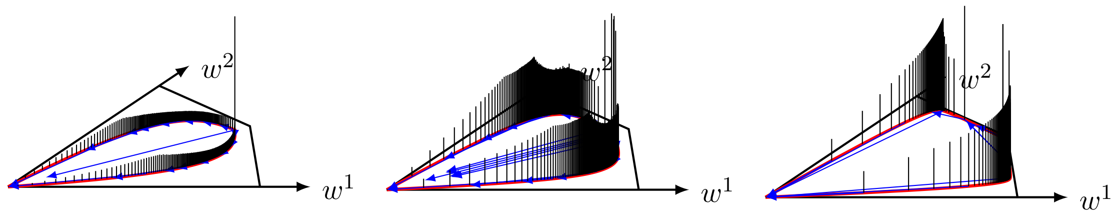

Figure 14: The blue arrows indicate the equilibrium rewards/punishments provided upon the realization of a demand shock in monitoring structure 1 (left), 2 (middle), and 3 (right) for $r = 5$ and $\gamma = 0.4$. The vertical axis shows the fraction ρ of total incentives that are provided via abrupt information. For most payoffs outside of $\Gamma(r, \mathcal{W})$, it is in the range of 8%–15% in monitoring structure 1, 34%–54% in monitoring structure 2, and 49%–82% in monitoring structure 3.

7.3 Relation to Sannikov and Skrzypacz [24]

Because of the similarity of our model with Sannikov and Skrzypacz [24], it is a natural question to ask how the results of the two papers relate. In their paper, Sannikov and Skrzypacz establish a payoff bound for discrete-time games with a sufficiently short time period. They apply the techniques from Abreu, Pearce, and Stacchetti [3] together with the informational restrictions from the continuous-time limiting game to describe a linear program that results in an equilibrium payoff bound M. Since the same restrictions on the use of information apply to our model, M is also an upper bound for the equilibrium payoff set in our model.[15] Note, however, that the algorithm in Proposition 5.3 cannot be started with M as it is unknown whether or not $\mathcal{B}_r(M)$ is contained in M. Identically to the algorithm in Abreu, Pearce, and Stacchetti [3], if $\mathcal{B}_r(M) \not\subseteq M$, an iterated application of $\mathcal{B}_r$ may not lead to a decreasing sequence of payoff sets and may converge to a different limit. This paper moves beyond the payoff bound and gives a way to compute $\mathcal{E}(r)$ for any value of the discount rate r. Moreover, the techniques in this paper allow us to elicit the equilibrium strategies that attain payoff pairs on the efficient frontier of the equilibrium payoff set, which is not possible in a discrete-time setting of this generality.

8 Computation

In this section, we illustrate how to implement Theorem 6.10 numerically. We discuss how $\mathcal{S}_r(\mathcal{W})$ and $\mathcal{K}_r(\mathcal{W})$ can be implemented when there is more than one type of events in Sections 8.1 and 8.2 before illustrating how (5) and (6) can be implemented numerically in Section 8.3. Finally, we provide an improvement over the algorithm in Proposition 5.3 for numerical implementation in Section 8.4.

[15] This follows from the fact that $\mathcal{E}(r)$ is closed by Theorem 6.10. Any payoff pair on the boundary can thus be attained in equilibrium, which implies that the informational restrictions in their paper are satisfied. Therefore, the boundary and hence the equilibrium payoff set are contained in M.

8.1 Computing $\mathcal{S}_r(\mathcal{W})$ for arbitrary sets $\mathcal{W}$

Denote by $\mathcal{S}_{r,a}(\mathcal{W})$ the set of stationary payoffs, where players locally play action profile a. Let $e_y \in \mathbb{R}^{|Y|}$ denote the unit vector corresponding to event y and define $f_y(\delta) := \delta\big(\lambda(a) + re_y\big)$. For a stationary payoff $w = g(a) + \delta\lambda(a)$, the continuation value after event y lies in the set $\mathcal{W}$ if and only if $g(a) + f_y(\delta) = w + r\delta(y) \in \mathcal{W}$. Incentives at stationary payoffs in $\mathcal{S}_{r,a}(\mathcal{W})$ thus have to come from the set $\mathcal{X}_a(\mathcal{W}) := \Psi_a^0 \cap \bigcap_{y\in Y} f_y^{-1}\big(\mathcal{W} - g(a)\big)$, where f_y^{-1} denotes the inverse image under f_y and we denote by $\mathcal{W} \pm g(a)$ the translate of $\mathcal{W}$ by $\pm g(a)$. It is now straightforward that

$$\mathcal{S}_{r,a}(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \big(g(a) + \mathcal{X}_a(\mathcal{W})\lambda(a)\big), \qquad \mathcal{S}_r(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \mathcal{S}_{r,a}(\mathcal{W}),$$

where $\mathcal{X}_a(\mathcal{W})\lambda(a) := \{w \in \mathbb{R}^2 \mid \exists\, \delta \in \mathcal{X}_a(\mathcal{W}) \text{ with } \delta\lambda(a) = w\}$ denotes the projection of $\mathcal{X}_a(\mathcal{W})$ onto $\mathbb{R}^2$ in the direction $\lambda(a)$. Note that Ψ_a^0 is a convex polyhedron characterized by the affine inequalities in (3). For a discretization $\mathcal{W}'$ of $\mathcal{W}$ with extremal points $x_1, \ldots, x_n$ and corresponding normal vectors $N_1, \ldots, N_n$, the set $\mathcal{X}_a(\mathcal{W}')$ is a convex polyhedron again, characterized by the affine inequalities in (3) and

$$N_j^\top \delta(\lambda(a) + re_y) \leq N_j^\top (x_j - g(a)), \quad j = 1, \ldots, n,\ y \in Y.$$

One can thus compute extremal points $z_1, \ldots, z_{n+m}$ of the set $\mathcal{X}_a(\mathcal{W}')\lambda(a)$ by maximizing $N_j^\top \delta\lambda(a)$ for $j = 1, \ldots, n+m$ over $\delta \in \mathcal{X}_a(\mathcal{W}')$, where $N_{n+1}, \ldots, N_{n+m}$ are the normal vectors of Ψ_a^0. This linear program can be computed efficiently numerically.

8.2 Computing $\mathcal{K}_r(\mathcal{W})$ for arbitrary sets $\mathcal{W}$

We start by reducing the dimensionality of the necessary computations involved in finding $\mathcal{K}_{r,a}(\mathcal{W})$ for fixed a. Let $\mathcal{Y}(a) = \big(Y_1(a), \ldots, Y_m(a)\big)$ denote the finest partition of Y such that any deviation from a affects only the intensities of events $y \in Y_j(a)$ for a single j. In terms of enforceability constraints in (3), this means that each inequality restricts $\delta^i(y)$ only for $y \in Y_j(a)$ for a single j. Let $\Psi_a^i(Y_j)$ denote the solution set to those inequalities and let Ψ_a^i denote the set of all δ that provide sufficient incentives for player i to play a. Then $\Psi_a^i = \bigtimes_{j=1}^m \Psi_a^i(Y_j(a))$ for $i = 1, 2$. Since boundary points of $\mathcal{K}_{r,a}(\mathcal{W})$ are translates of $\partial\mathcal{W}$ by extremal incentives in Ψ_a^i, one can now write

$$\mathcal{K}_{r,a}(\mathcal{W}) = \bigcap_{j=1}^m \mathcal{K}_{r,a}^j(\mathcal{W}) \qquad \mathcal{K}_{r,a}^j(\mathcal{W}) := \bigcup_{\delta\in\partial\Psi_a(Y_j)} \bigcap_{y\in Y_j} \big(\mathcal{W} - r\delta(y)\big). \tag{8}$$

Computing $\mathcal{K}_{r,a}^j(\mathcal{W})$ may be a daunting task since $\Psi_a^i(Y_j)$ is potentially high-dimensional. The computation is simple, however, in two cases that are of frequent interest.

Consider first the case where $Y_j(a) = \{y\}$ is a singleton. Then $\Psi_a(Y_j)$ is a (possibly unbounded) rectangle in $\mathbb{R}^2$, whose boundaries determine the extremal rewards/punishments upon the arrival of event of type y that enforce a. It is easy to

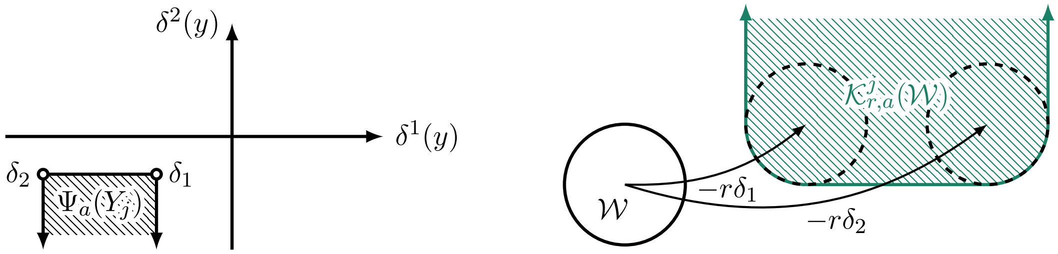

Figure 15: Construction of $\mathcal{K}^j_{r,a}(\mathcal{W})$ if $Y_j = \{y\}$ is a singleton. The arrows indicate that $\Psi_a(Y_j)$ and $\mathcal{K}^j_{r,a}(\mathcal{W})$ continue indefinitely in these directions, i.e., two "corners" of $\Psi_a(Y_j)$ are degenerate.

see that $\mathcal{K}^j_{r,a}(\mathcal{W})$ is then simply the convex hull of $\mathcal{W} - r\delta_k$ for $k = 1, \ldots, 4$, where δ_k denote the four (possibly degenerate) corners of $\Psi_a(Y_j)$; see Figure 15.

Consider next the case, where the intensities of events in $Y_j(a)$ for any j are affected only by deviations of a single player i. Punishments/rewards for player $-i$ after an event in $Y_j(a)$ do not affect $-i$'s incentives, hence $\mathcal{K}^j_{r,a}(\mathcal{W})$ is the infinite band that contains all payoff pairs v that satisfy

$$\min_{w\in\mathcal{W}} w^i - r\bar{\delta}^i \le v^i \le \max_{w\in\mathcal{W}} w^i - r\underline{\delta}^i \tag{9}$$

where $\underline{\delta}^i := \min_{\delta\in\Psi^i_a} \max_{y\in Y_i} \delta^i(y)$ and $\bar{\delta}^i := \max_{\delta\in\Psi^i_a} \min_{y\in Y_i} \delta^i(y)$ are the extremal punishments/rewards for player i that can be carried out after the arrival of any event in Y_j such that i is incentivized to play a^i against a^{-i}. This case contains games with a *product structure*, where the intensity of each event is determined by the actions of a single player, such as monitoring structure 3 in Section 7.2. In games with a product structure, (8) implies that $\mathcal{K}_{r,a}(\mathcal{W})$ is a rectangle for every $a \in \mathcal{A}$, whose the boundaries can be computed efficiently with (9).

While the computation of $\mathcal{K}_{r,a}(\mathcal{W})$ is cumbersome in the most general case, the verification of whether a given payoff pair w lies in $\mathcal{K}_{r,a}(\mathcal{W})$ is easy: w is contained in $\mathcal{K}_{r,a}(\mathcal{W})$ if and only if

$$\Psi_a \cap \left(\frac{\mathcal{W} - w}{r}\right)^{|Y|} \neq \emptyset.$$

For an implementation with a polygon approximation $\mathcal{W}'$ of $\mathcal{W}$ with extremal points $x_1, \ldots, x_n$ and corresponding normal vectors $N_1, \ldots, N_n$, this corresponds to verifying whether or not the system of affine inequalities in (3) and

$$rN_j^\top \delta(y) \le N_j^\top (x_j - w), \quad j = 1, \ldots, n,\ y \in Y$$

is feasible. As illustrated below, this is sufficient to compute the boundary of $\mathcal{B}_r(\mathcal{W})$.

8.3 Computing $\mathcal{B}_r(\mathcal{W})$ for arbitrary sets $\mathcal{W}$

The first step is to compute the set of stationary payoffs. The second step is to find the largest solution to (5) and (6) that contains $\mathcal{S}_r(\mathcal{W})$. Since $\mathcal{B}_r(\mathcal{W})$ is convex, we parametrize it via tangential angle θ. Let $w(\theta)$ denote the set of payoff pairs in $\mathcal{B}_r(\mathcal{W})$ with normal vector $N(\theta) = \big(\cos(\theta), \sin(\theta)\big)^\top$. If the continuous component of the public signal is not completely uninformative, then the boundary is strictly curved outside of $\Gamma(r, \mathcal{W})$, hence $w(\theta)$ is a singleton. Thus, one can solve

$$\frac{\mathrm{d}w(\theta)}{\mathrm{d}\theta} = \frac{T(\theta)}{\kappa(\theta)}$$

numerically, where $T(\theta) = \big(-\sin(\theta), \cos(\theta)\big)^\top$ and $\kappa(\theta) = \kappa\big(w(\theta)\big)$ is given by the optimality equation (5). If the continuous signal is completely uninformative, solutions to (5) are simply straight line segments in the direction of $T(\theta)$. If $(w(\theta), N(\theta))$ is in $\Gamma(r, \mathcal{W})$, we distinguish two cases:

(i) If there exist no a and $\delta \in \Psi_a^0(w(\theta), r, \mathcal{W})$ with $N(\theta)^\top\big(g(a)+\delta\lambda(a)-w(\theta)\big) > 0$, then we solve (6) numerically as follows: $w(\theta_{k+1})$ is the payoff w furthest from $w(\theta_k)$ in the direction of $T(\theta_k)$, for which $\max_{a,\delta} N(\theta_{k+1})^\top\big(g(a)+\delta\lambda(a)-w\big) = 0$.

(ii) If there exist a and $\delta \in \Psi_a^0(w(\theta), r, \mathcal{W})$ with $N(\theta)^\top\big(g(a) + \delta\lambda(a) - w(\theta)\big) > 0$, then $w(\theta)$ is strictly decomposable by some set of action profiles $\mathcal{A}_\theta$ such that $N(\theta)$ is a normal vector to $\mathcal{K}_{r,\mathcal{A}_\theta}(\mathcal{W})$. It follows from Propotisition 6.9 that $w(\theta)$ is locally constant until $N(\theta)$ is an extremal normal vector of $\mathcal{K}_{r,\mathcal{A}_\theta}(\mathcal{W})$. If $\max_{a,\delta} N(\theta)^\top\big(g(a)+\delta\lambda(a)-w(\theta)\big) = 0$ for the extremal normal vector $N(\theta)$, we resume solving (6), otherwise $\partial\mathcal{B}_r(\mathcal{W})$ locally follows $\partial\mathcal{K}_{r,a}(\mathcal{W})$ of the maximizing action profile a. That is, $w(\theta)$ maximizes $N(\theta)^\top w$ among all $w \in \mathcal{K}_{r,a}(\mathcal{W})$, for which there exists a and $\delta \in \Psi_a^0(w, r, \mathcal{W})$ with $N(\theta)^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$.

While it is helpful to know $\mathcal{K}_{r,a}(\mathcal{W})$ explicitly for case (ii), one can also use the verification procedure at the end of Section 8.2. Specifically, $N(\theta)$ is a normal vector to $\mathcal{K}_{r,\mathcal{A}_\theta}(\mathcal{W})$ at $w(\theta)$ if and only if $w(\theta) \in \mathcal{K}_{r,a}(\mathcal{W})$ for every $a \in \mathcal{A}_\theta$ and $w(\theta)+\varepsilon N'$ lies outside of $\mathcal{K}_{r,a}(\mathcal{W})$ for any $a \in \mathcal{A}_\theta$, any sufficiently small $\varepsilon > 0$, and sufficiently slight outward rotations N' of $\pm T(\theta)$. To find the payoff pair $w \in \mathcal{K}_{r,a}(\mathcal{W})$ that maximizes $N(\theta_{k+1})^\top w$, one can perform a binary search between $w(\theta_k)$ and the intersection point $\hat{w}$ of $\partial\mathcal{K}_{r,a}(\mathcal{W})$ and the straight line through $w(\theta_k)$ in direction $T(\theta_{k+1})$ as follows. Let $\tilde{w}$ denote the projection of the mid point between $w(\theta_k)$ and $\hat{w}$ onto $\partial\mathcal{K}_{r,a}(\mathcal{W})$ in direction $N(\theta_{k+1})$. If $N(\theta_{k+1})$ is a normal vector of $\mathcal{K}_{r,a}(\mathcal{W})\big)$ at $\tilde{w}$, then $w(\theta_{k+1})$ is the payoff pair on the tangent through $\tilde{w}$ closest to $w(\theta_k)$. If $N(\theta_{k+1})$ is not a normal vector to $\mathcal{K}_{r,a}(\mathcal{W})$ at $\tilde{w}$, then we continue the binary search between $\tilde{w}$ and $\hat{w}$ if $\tilde{w} + \varepsilon T(\theta_{k+1}) \in \mathcal{K}_{r,a}(\mathcal{W})$ and between $\tilde{w}$ and $w(\theta_k)$ if $\tilde{w} - \varepsilon T(\theta_{k+1}) \in \mathcal{K}_{r,a}(\mathcal{W})$.

We illustrate this procedure by showing how $\mathcal{B}_r(\mathcal{V}^*)$ can computed in the partnership example for the second monitoring structure with $r = 5$ and $\gamma = 0.8$. We first compute the set of stationary payoffs as described in Section 6.1. Since $\mathcal{S}_r(\mathcal{V}^*)$ overlaps with $\partial\mathcal{V}^*$, it is necessary that $\mathcal{S}_r(\mathcal{V}^*) \cap \partial\mathcal{V}^*$ lies on the boundary of $\mathcal{B}_r(\mathcal{V}^*)$.

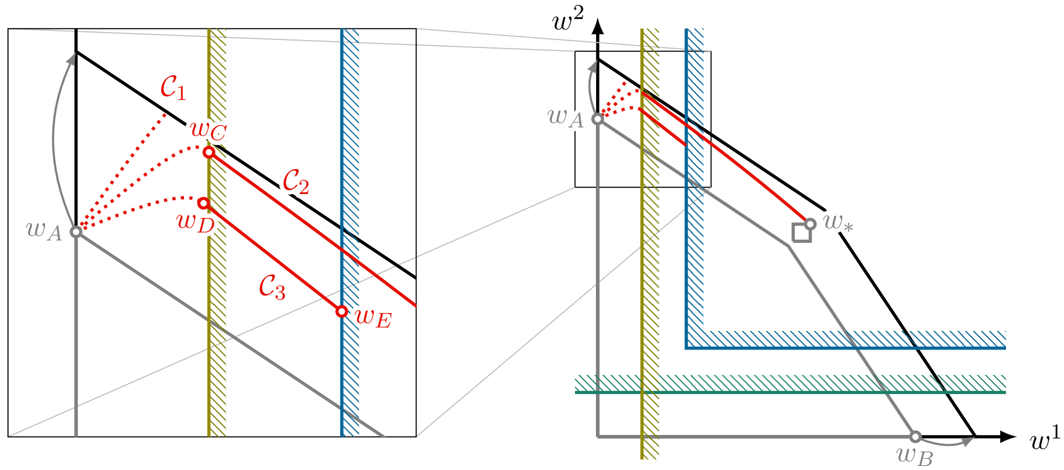

Figure 16: Shown are the set $\mathcal{S}_r(\mathcal{V}^*)$ of stationary payoffs, the sets $\mathcal{K}_{r,(1,1)}(\mathcal{V}^*)$, $\mathcal{K}_{r,(0,1)}(\mathcal{V}^*)$, and $\mathcal{K}_{r,(1,0)}(\mathcal{V}^*)$, where incentives are binding and extremal, as well as three solutions $\mathcal{C}_i$ to (5) (dotted) and (6) (solid) for $i = 1, 2, 3$ with initial conditions $(w_A, N(\theta_i))$ and $\theta_1 > \theta_2 > \theta_3$. The light gray arrows indicate the rewards/punishments that attain stationary payoff pairs w_A and w_B.

Payoff pairs w_A and w_B in Figure 16 thus serve as starting points for solving (5) and (6).[16] We search for the maximal angle θ at w_A, for which a solution connects to either w_B or w_*. Solution $\mathcal{C}_1$ for initial condition $(w_A, N(\theta_1))$ escapes $\mathcal{V}^*$. It follows that θ_1 is too large. Solution $\mathcal{C}_3$ for initial condition $(w_A, N(\theta_3))$ reaches a payoff pair w_D, for which action profile $(0,1)$ satisfies (6) at (w_D, N_{w_D}). We extend $\mathcal{C}_3$ by solving (6) with initial conditions (w_D, N_{w_D}). The curve eventually reaches a payoff pair $w_E \in \partial\mathcal{K}_{r,(1,1)}(\mathcal{V}^*)$, for which $\max_\delta {N_{w_E}}^\top\big(g(1,1) + \delta\lambda(1,1) - w_E\big) > 0$. Since (7) is violated, $\mathcal{C}_3$ cannot be the boundary of $\mathcal{B}_r(\mathcal{W})$, hence θ_3 is too small. Solution $\mathcal{C}_2$ for initial condition $(w_A, N(\theta_2))$ changes from a solution to (5) to a solution to (6) at w_C and it connects to w_*. After verifying that any solution with initial conditions $(w_A, N(\theta))$ for $\theta > \theta_2$ escapes $\mathcal{V}^*$, we conclude that $\mathcal{C}_2$ coincides with $\partial\mathcal{B}_r(\mathcal{V}^*)$.[17]

8.4 Refinement of the algorithm for $\mathcal{E}(r)$

Observe that the set of stationary payoffs shrinks with every iteration: extremal payoff pairs in $\mathcal{S}_r(\mathcal{W}_n) \cap \partial\mathcal{V}^*$ are not contained in the next step of the iteration if they are decomposed using outward jumps, i.e., jumps $\delta(y)$ with $N^\top\delta(y) > 0$ for some event y. Examples are the payoff pairs w_A and w_B in Figure 16. Those are decomposed by the static Nash profile and incentives that give player 2 and 1, respectively, his highest feasible and individually rational payoff upon the occurrence of a demand shock. In the limit as $\mathcal{B}_r$ is applied infinitely often, incentives using outward jumps cannot

[16] If no extremal payoff pair from $\mathcal{B}_r(\mathcal{W})$ is known, one can use an iterative procedure as described in Section 8 of Sannikov [22].

[17] Since the game is symmetric, one can also conclude that w_* is a boundary payoff of $\mathcal{B}_r(\mathcal{V}^*)$ by verifying that a solution with initial condition $(w_*, \pi/4)$ escapes $\mathcal{V}^*$.

support equilibrium behavior since $\mathcal{E}(r)$ is convex. This motivates the definition of a refined algorithm that excludes the use of such incentives from the beginning.

We begin by defining the set of payoff pairs that can be decomposed using inward jumps. For any action profile $a \in \mathcal{A}$ and any direction $N \in S^1$, let

$$H_a(N) := \left\{ w \in \mathbb{R}^2 \;\middle|\; \begin{array}{l} \exists\, \delta \in \Psi_a^0 \text{ with } N^\top \delta(y) \leq 0 \text{ for every } y \in Y \\ \text{and } N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0 \end{array} \right\}.$$

be the half-space of payoff pairs that can be decomposed with inward-jumps with respect to the direction N. Similarly to Fudenberg and Levine [9], $g(a)$ cannot be in the interior of $H_a(N)$. Let $D_a := \{N \in S^1 \mid H_a(N) \neq \emptyset\}$ denote the set of directions, with respect to which a can be decomposed using inward jumps only. The set $\mathcal{Q}_a := \bigcap_{N \in D_a} H_a(N)$ is a bound for all payoff pairs that are decomposable by a using only inward-pointing jumps. Since a static Nash profile $a \in \mathcal{A}^N$ is decomposable without any jumps at all, $g(a) \in \partial H_a(N)$ for all $N \in S^1$ and hence $\mathcal{Q}_a = \{g(a)\}$. The following lemma follows straight from the definition of $\mathcal{Q}_a$.

Lemma 8.1. *Suppose that $\mathcal{A}_w \subseteq \mathcal{A}$ decomposes $w \in \partial\mathcal{B}_r(\mathcal{W})$. For any $a \in \mathcal{A}_w$, either $w \in \mathcal{Q}_a$ or any δ with (a, δ) decomposing w satisfies $N^\top\delta(y) > 0$ for some $y \in Y$ and some $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$.*

We obtain the following fixed-point characterization of $\mathcal{E}(r)$ as an immediate consequence to Theorem 6.10 and Lemma 8.1.

Corollary 8.2. *$\mathcal{E}(r)$ is the largest closed convex subset $\mathcal{X}$ of $\mathcal{V}^*$ that satisfies Conditions (i) and (ii) of Theorem 6.10 for $\mathcal{W} = \mathcal{X}$ such that every corner w of $\mathcal{X}$ lies in $\mathcal{S}_{r,a}(\mathcal{W}) \cap \mathcal{Q}_a$ or in $\mathcal{K}_{r,a}(\mathcal{W}) \cap \mathcal{Q}_a$ for every a in $\mathcal{A}_w$ that minimally decomposes w.*

The refinement in Corollary 8.2 allows us to exclude many points that would have to be considered as potential corners and extremal points of $\mathcal{E}(r)$. The following algorithm clarifies that we can exclude all of these points straight from the beginning, leading to a faster computation of $\mathcal{E}(r)$.

Definition 8.3. For a compact and convex set $\mathcal{W} \subseteq \mathcal{V}^*$, let $\tilde{\mathcal{B}}_r(\mathcal{W})$ denote the largest closed, convex subset $\mathcal{X}$ of $\mathcal{V}^*$ such that:

(i) Outside of $\mathcal{D} := \bigcup_{a \in \mathcal{A}}\big(\mathcal{S}_{r,a}(\mathcal{W}) \cup \partial\mathcal{K}_{r,a}(\mathcal{W})\big) \cap \mathcal{Q}_a$, the boundary $\partial\mathcal{X}$ is a solution to (6) within $\Gamma(r, \mathcal{W})$ and to (5) outside of $\Gamma(r, \mathcal{W})$, where admissible incentives in (5) and (6) additionally have to satisfy $N^\top\delta(y) \leq 0$ for every $y \in Y$.

(ii) Within $\mathcal{D}$, the boundary $\partial\mathcal{X}$ satisfies Condition (ii) of Theorem 6.10.

Proposition 8.4. *Let $\mathcal{W}_0 = \mathcal{V}^*$ and $\mathcal{W}_n := \tilde{\mathcal{B}}_r(\mathcal{W}_{n-1})$ for $n \geq 1$. Then $(\mathcal{W}_n)_{n \geq 0}$ is decreasing in the set-inclusion sense with $\bigcap_{n \geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*

By comparing Figures 11 and 17, we see that the algorithm in Proposition 8.4 converges much faster than the algorithm in Proposition 5.3.

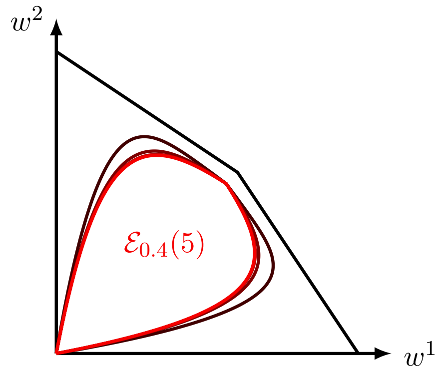

Figure 17: Convergence of the algorithm in Proposition 8.4.

9 Conclusion

This paper characterizes the PPE payoff set for two-player games with imperfect public observation, in which information arrives through a mixture of a drifted Brownian motion and Poisson processes. The boundary is characterized by a second-order ODE where intertemporal incentives via Brownian information are required and by a first-order ODE where intertemporal incentives are provided by abrupt information exclusively. In the presence of abrupt information, the equilibrium payoff set may have corners at so-called stationary payoffs, which demand locally constant play in equilibrium, or at payoffs with binding and extremal incentives. The methods from this paper allows us to recover the generically unique action profiles that have to be played at extremal equilibrium payoffs as well as the unique equilibrium incentives.

Equilibria are constructed iteratively between the arrival times of the Poisson events. The main tool used in the construction is the notion of relaxed self-generating payoff sets. It extends the set-valued operator of Abreu, Pearce, and Stacchetti [3] to a continuous-time setting. This concept is useful beyond the theory of repeated games for any subsequent research on continuous-time games that involve abrupt information arrival such as, for example, stochastic games with a finite state space. Indeed, the state process in such a game is precisely described by a set of Poisson processes, whose intensities govern the rate at which the underlying state changes.

Because of the quantitative nature of the result, the impact of information on equilibrium payoffs can be measured precisely, paving the way for future research on information revelation: a company may choose to publicly disclose certain information (make it continuously observable) or keep the information from the public until the media finds out and reports on it (abrupt information). Because continuous and discontinuous information have fundamentally different impacts on equilibrium payoffs, a strategic company may prefer one over the other and act accordingly.

References

[1] D. Abreu, P. Milgrom, and D. Pearce: Information and timing in repeated partnerships, *Econometrica*, **59** (1991), 1713–1733

[2] D. Abreu, D. Pearce, and E. Stacchetti: Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory*, **39** (1986), 251–269

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

[4] D. Abreu and Y. Sannikov: An algorithm for two-player repeated games with perfect monitoring, *Theoretical Economics*, **9** (2014), 313–338

[5] R. Anderson: Quick-response equilibrium. Working Papers in Economic Theory and Econometrics # IP-323, Center for Research in Management, University of California, Berkeley (1984)

[6] A.V. Arutyunov and V. Obukhovskii: *Convex and Set-Valued Analysis, Selected Topics*, De Gruyter, 2017

[7] J.-P. Aubin and H. Frankowska: *Set-Valued Analysis*, Springer, Berlin, 1990

[8] B. Bernard and C. Frei: The folk theorem with imperfect public information in continuous time, *Theoretical Economics*, **11** (2016), 411–453

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

[10] D. Fudenberg and D. K. Levine: Continuous time limits of repeated games with imperfect public monitoring, *Review of Economic Dynamics*, **10** (2007), 173–192

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

[12] E. J. Green and R. H. Porter: Noncooperative collusion under imperfect price information, *Econometrica*, **52** (1984), 87–100

[13] T. Hashimoto: Corrigendum to 'Games with imperfectly observable actions in continuous time', *Econometrica*, **78** (2010), 1155–1159

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

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

[17] Ioannis Karatzas and Steven Shreve: *Methods of Mathematical Finance*, Applications of Mathematics 39, Springer, New York, 1998

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

[19] G.J. Mailath and L. Samuelson: *Repeated Games and Reputations*, Oxford University Press, 2006

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

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

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

[23] Y. Sannikov and A. Skrzypacz: Impossibility of collusion under imperfect monitoring with flexible production, *American Economic Review*, **97** (2007), 1794–1823

[24] Y. Sannikov and A. Skrzypacz: The role of information in repeated games with frequent actions, *Econometrica*, **78** (2010), 847–882

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

[26] Y. Sun: The exact law of large numbers via Fubini extension and characterization of insurable risks, *Journal of Economic Theory*, **126** (2006), 31–69

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

A Proofs of auxiliary results in the main text

This appendix contains the proofs of the results in Sections 4, 5, and 8.

A.1 Dynamics of the continuation value and continuation promises

The proofs of the results in Section 4 follow the proofs of the analogous results in Bernard and Frei [8] rather closely, extending them to information structures that also contain abrupt information. We begin with the stochastic differential representation of the continuation value in Lemma 4.1. Intuitively, we wish to apply a martingale representation result to the continuation value at time ∞ under the family Q^A of probability measures induced by play of strategy profile A. However, since the limiting measure Q^A_∞ is not equivalent to P on $\mathcal{F}_\infty$, we instead use the martingale representation at finite times and show that the representation is time consistent.[18]

[18] The existence of such a limiting measure Q^A_∞, which coincides with Q^A_t on $\mathcal{F}_t$ for every $t \geq 0$, is given by Proposition 1.7.4 in Karatzas and Shreve [17].

Proof of Lemma 4.1. Fix a strategy profile A and let $W := W(A)$. We will show that W^i satisfies (2) for $i = 1, 2$. Observe first that W 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 $w^i_T := W^i_T - r\int_0^T \big(W^i_t - g^i(A_t)\big)\,\mathrm{d}t$. Because $(J^y)_{y\in Y}$ are pairwise orthogonal and orthogonal to Z, the stable subspace generated by Z and $(J^y)_{y\in Y}$ is the space of all stochastic integrals with respect to these processes by Theorem IV.36 in Protter [21]. Therefore, we obtain the unique martingale representation property for a square-integrable martingale by Corollary 1 to Theorem IV.37 in [21]. That is, for a bounded $\mathcal{F}_T$-measurable random variable w^i_T, there exist an $\mathcal{F}_0$-measurable c^i_T, predictable and square-integrable processes $(\beta^i_{t,T})_{0\le t\le T}$ and $\big(\delta^i_{t,T}(y)\big)_{0\le t\le T}$ for $y \in Y$ with $\mathbb{E}_{Q^A_T}\Big[\int_0^T \big|\beta^i_{t,T}\big|^2\,\mathrm{d}t\Big] < \infty$ and $\mathbb{E}_{Q^A_T}\Big[\int_0^T \big|\delta^i_{t,T}(y)\big|^2 \lambda(y\,|A_t)\,\mathrm{d}t\Big] < \infty$ and a Q^A_T-martingale M^i orthogonal to Z and all processes $(J^y)_{y\in Y}$ with $M^i_0 = 0$ such that

$$w^i_T = c^i_T + \int_0^T r\beta^i_{t,T}\,(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t) + \sum_{y\in Y}\int_0^T r\delta^i_{t,T}(y)\,(\mathrm{d}J^y_t - \lambda(y\,|A_t)\,\mathrm{d}t) + M^i_{T,T}.$$

To prove that (2) holds, we need to show that c^i_T, $\beta^i_{t,T}$, $\delta^i_{t,T}(y)$ and $M^i_{t,T}$ do not depend on T. It follows from (1) and Fubini's theorem that

$$w^i_T = W^i_T + r\int_0^T g^i(A_t)\,\mathrm{d}t - r\int_0^\infty\int_0^{s\wedge T} r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q^A_s}\big[g^i(A_s)\;\big|\;\mathcal{F}_t\big]\,\mathrm{d}t\,\mathrm{d}s. \qquad (10)$$

Let $\tilde{T} \le T$ and take conditional expectations on $\mathcal{F}_{\tilde{T}}$ under Q^A_T of (10) to deduce that

$$\begin{aligned}
\mathbb{E}_{Q^A_T}\big[w^i_T\,\big|\,\mathcal{F}_{\tilde{T}}\big] - w^i_{\tilde{T}} &= \mathbb{E}_{Q^A_T}\big[W^i_T\,\big|\,\mathcal{F}_{\tilde{T}}\big] - W^i_{\tilde{T}} + r\int_{\tilde{T}}^T \mathbb{E}_{Q^A_t}\big[g^i(A_t)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}t \\
&\quad - r\int_{\tilde{T}}^\infty\int_{\tilde{T}}^{s\wedge T} r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q^A_s}\big[g^i(A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}t\,\mathrm{d}s \\
&= \mathbb{E}_{Q^A_T}\big[W^i_T\,\big|\,\mathcal{F}_{\tilde{T}}\big] - W^i_{\tilde{T}} - \int_T^\infty r\mathrm{e}^{-r(s-T)}\mathbb{E}_{Q^A_s}\big[g^i(A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}s \\
&\quad + \int_{\tilde{T}}^\infty r\mathrm{e}^{-r(s-\tilde{T})}\mathbb{E}_{Q^A_s}\big[g^i(A_s)\,\big|\,\mathcal{F}_{\tilde{T}}\big]\,\mathrm{d}s \\
&= 0.
\end{aligned}$$

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

$$w^i_{\tilde{T}} = W^i_0 + r\int_0^{\tilde{T}} \beta^i_{t,T}\big(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t\big) + \sum_{y\in Y} r\int_0^{\tilde{T}} \delta^i_{t,T}(y)\big(\mathrm{d}J^y_t - \lambda(y|A_t)\,\mathrm{d}t\big) + M^i_{\tilde{T},T},$$

which yields $\beta^i_{\cdot,T} = \beta^i_{\cdot,\tilde{T}}$ and $\delta^i_{\cdot,T}(y) = \delta^i_{\cdot,\tilde{T}}(y)$ for every $y \in Y$ 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, $\big(\delta^i(y)\big)_{y\in Y}$, and M^i, hence W^i satisfies (2).

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

$$\begin{aligned} \mathrm{d}(\mathrm{e}^{-rt}W^i_t) = -r\mathrm{e}^{-rt}g^i(A_t)\,\mathrm{d}t + r\mathrm{e}^{-rt}\beta^i_t\left(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t\right) \\ + r\mathrm{e}^{-rt}\sum_{y\in Y}\delta^i_t(y)\big(\mathrm{d}J^y_t - \lambda(y|A_t)\,\mathrm{d}t\big) + \mathrm{e}^{-rt}\,\mathrm{d}M^i_t. \end{aligned} \tag{11}$$

Since M^i is strongly orthogonal to Z and $(J^y)_{y\in Y}$, it is also strongly orthogonal to the density process given in Footnote 4. Therefore, M^i is a martingale also under the family of probability measures Q^A. Integrating (11) from t to T and taking Q^A_T-conditional expectations on $\mathcal{F}_t$ thus yields

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

Since W is bounded, the second summand converges to zero a.s. as T tends to ∞, hence W^i_t is indeed the continuation value of A. □

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 $W(A)$ by (2), integrating (11) from t to u yields

$$\begin{aligned} W^i_t(A) = -\int_t^u \mathrm{e}^{-r(s-t)}\Big(\beta^i_s\left(\mathrm{d}Z_s - \mu(A_s)\,\mathrm{d}s\right) - g^i(A_s)\,\mathrm{d}s - \mathrm{d}M^i_s\Big) \\ - \sum_{y\in Y}\int_t^u \mathrm{e}^{-r(s-t)}\delta^i_s(y)\left(\mathrm{d}J^y_s - \lambda(y|A_s)\,\mathrm{d}s\right) + \mathrm{e}^{-r(u-t)}W^i_u(A). \end{aligned}$$

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

$$\begin{aligned} W^i_t(\tilde{A}) &= \lim_{u\to\infty}\mathbb{E}_{Q^{\tilde{A}}_u}\left[\int_t^u r\mathrm{e}^{-r(s-t)}g^i(\tilde{A}_s)\,\mathrm{d}s\,\middle|\,\mathcal{F}_t\right] \\ &= W^i_t(A) + \lim_{u\to\infty}\mathbb{E}_{Q^{\tilde{A}}_u}\bigg[\int_t^u r\mathrm{e}^{-r(s-t)}\Big(\big(g^i(\tilde{A}_s) - g^i(A_s)\big)\,\mathrm{d}s \\ &\qquad + \beta^i_s\left(\mathrm{d}Z_s - \mu(A_s)\,\mathrm{d}s\right) + \sum_{y\in Y}\delta^i_s(y)\left(\mathrm{d}J^y_s - \lambda(y|A_s)\,\mathrm{d}s\right)\Big)\bigg|\,\mathcal{F}_t\bigg] \text{ a.s.} \end{aligned}$$

Because the processes β and $\delta(y)$ for $y \in Y$ are constructed using a martingale representation result for the bounded random variable w_T^i in (10), the processes $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\beta_s^i \left(\mathrm{d}Z_s - \mu(A_s)\,\mathrm{d}s\right)$ and $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\delta_s^i(y)\left(\mathrm{d}J_s^y - \lambda(y|A_s)\,\mathrm{d}s\right)$ are bounded mean oscillation (BMO) martingales under the probability measure Q_u^A up to any time $u > t$. Since Assumption 1 implies that the jumps of $\left(\lambda(y|A_{s-})-1\right)\Delta J_s^y$ in the density process in Footnote 4 are bounded from below by $-1+\varepsilon$ for any $y \in Y$, it follows from Remark 3.3 and Theorem 3.6 in Kazamaki [18] that $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\beta_s^i \left(\mathrm{d}Z_s - \mu(\tilde{A}_s)\,\mathrm{d}s\right)$ and $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\delta_s^i(y)\left(\mathrm{d}J_s^y - \lambda(y|\tilde{A}_s)\,\mathrm{d}s\right)$ are BMO-martingales under $Q_u^{\tilde{A}}$. Together with Fubini's theorem, this implies

$$W_t^i(\tilde{A}) - W_t^i(A) = \int_t^\infty \mathrm{e}^{-r(s-t)}\mathbb{E}_{Q_s^{\tilde{A}}}\Big[g^i(\tilde{A}_s) - g^i(A_s) + \beta_s^i\left(\mu(\tilde{A}_s) - \mu(A_s)\right) + \delta_s^i\left(\lambda(\tilde{A}_s) - \lambda(A_s)\right)\Big|\,\mathcal{F}_t\Big]\,\mathrm{d}s \quad \text{a.s.} \tag{12}$$

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

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

on the set Ξ. 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 $\tilde{A} = (\tilde{A}^i, A^{-i})$, the expectation in (12) is strictly positive for $t = 0$, which is a contradiction to the assumption that A is a PPE. □

Proof of Lemma 4.5. The proof of Lemma 2 in Bernard and Frei [8] works for signals given by any Lévy process. It thus directly applies here. □

A.2 Convergence of the algorithm

The main idea behind the proof of Lemma 5.2 is the following. Any payoff pair in $\mathcal{B}_r(\mathcal{W})$ can be attained by an en enforceable solution to (2) that remains in $\mathcal{B}_r(\mathcal{W})$ until time σ_1 and $W_{\sigma_1} \in \mathcal{W}$. Since $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$, we would like to attain W_σ by another enforceable solution W to (2) until the arrival of the next event and then concatenate the two solutions. Such a concatenation, however, is subject to some subtle measurability issues that we will address here.

Without restrictions on β and δ, solutions to (2) are weak solutions, that is, the Brownian motion, the Poisson processes, and the entire probability space are part of the solution. We refer to $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$ as the *stochastic framework* for the solution of (2). Thus, more formally, the payoff set $\mathcal{B}_r(\mathcal{W})$ is defined as

$$\mathcal{B}_r(\mathcal{W}) = \left\{ w \in \mathcal{V} \;\middle|\; \begin{array}{l} \text{There exists a stochastic framework } (\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y}) \\ \text{containing a solution } (W, A, \beta, \delta, M) \text{ to (2) on } [\![0,\sigma]\!] \text{ with} \\ W_0 = w \text{ and } W_\sigma \in \mathcal{W} \ P\text{-a.s., and on } [\![0,\sigma)\!), W \in \mathcal{B}_r(\mathcal{W}) \text{ and} \\ (\beta,\delta) \text{ enforces } A, \text{where } \sigma \text{ is the first jump time of } (J^y)_{y\in Y}. \end{array} \right\}.^{19,20}$$

It is, therefore, not a priori clear that we can find a continuation solution for each realization of W_σ that lives in the same probability space as $\mathcal{W}_\sigma$. The following result establishes that this is possible: because the probability spaces for different realizations of W_τ share the same path space, we can construct a regular conditional probability that allows us to aggregate the different probability spaces.

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

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

(b) There exists a solution (W, A, β, δ, M) to (2) in the given stochastic framework such that $W_0 = W^*$ P-a.s., $W_\sigma \in \mathcal{W}$ P-a.s., and on $[\![0,\sigma)\!)$, we have $W \in \mathcal{B}_r(\mathcal{W})$ and (β,δ) enforces A, where σ is the first jump time of $(J^y)_{y\in Y}$.

Proof. Let first $W^* \in \mathcal{B}_r(\mathcal{W})$ a.s. Although a solutions to (2) 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 $\mathcal{A}^{[0,\infty)} \times \mathcal{D}^{d+|\mathcal{Y}|}$, where $\mathcal{D}^{d+|\mathcal{Y}|}$ is the space of càdlàg functions from $[0,\infty)$ into $\mathbb{R}^{d+|\mathcal{Y}|}$. By Theorem A.2.2 in Kallenberg [16], there exists a metric on $\mathcal{D}^{d+|\mathcal{Y}|}$ that induces the Skorohod topology, under which $\mathcal{D}^{d+|\mathcal{Y}|}$ is complete and separable. Since $\mathcal{A}^{[0,\infty)}$ is compact by Tychonoff's theorem (see Theorem 37.3 in Munkres [20]), it follows that $\Omega = \mathcal{V}\times\mathcal{A}^{[0,\infty)}\times\mathcal{D}^{d+|\mathcal{Y}|}$ is complete and separable. Thus, by Theorem V.3.19 in Karatzas and Shreve [17], 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^*.

[19]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.

[20]We say that a stochastic process X satisfies a certain property on $\Xi \subseteq \Omega\times[0,\infty)$ if $X_t(\omega)$ satisfies that property for $P \otimes \textit{Lebesgue}$-almost every $(\omega,t) \in \Xi$.

For each $w \in \mathcal{B}_r(\mathcal{W})$ there exists a solution $(W^w, A^w, \beta^w, \delta^w, M^w)$ to (2) such that $W_0^w = w$ 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 (2) such that $W_0 = W^*$, (β, δ) enforces A, $W + r\delta(y) \in \mathcal{W}$ a.e. for every $y \in Y$, and $W \in \mathcal{B}_r(\mathcal{W})$. It follows from the properties of a regular conditional probability that

$$P(\Xi) = \int_{\mathcal{B}_r(\mathcal{W})} P(\Xi \,|\, W^* = w)\, \mathrm{d}\nu(w) = 1. \tag{13}$$

To show the converse, let (W, A, β, δ, M) be a solution to (2) such that $W_0 = W^*$ P-a.s., (β, δ) enforces A, $W + r\delta(y) \in \mathcal{W}$ a.e. for every $y \in Y$, and $W \in \mathcal{B}_r(\mathcal{W})$. Define $(W^w, A^w, \beta^w, \delta^w, M^w)$ by setting $W^w := W1_{\{W^*=w\}}$ and similarly for the other processes. Let $\Xi(w)$ denote the set on which (β^w, δ^w) enforces A^w, $W^w + r\delta^w(y) \in \mathcal{W}$ a.e. for every $y \in Y$, and $W^w \in \mathcal{B}_r(\mathcal{W})$. By (13), the set of w, for which $P_w(\Xi(w)) < 1$ has ν-measure 0. Therefore, $W^* \in \mathcal{B}_r(\mathcal{W})$ holds P-almost surely. □

Equipped with Lemma A.1, we are ready to concatenate enforceable solutions to (2). The concatenation procedure is an essential tool also in the proofs of other results, hence we isolate it in a separate lemma for ease of reference.

Lemma A.2. *Fix two sets $\mathcal{W}, \mathcal{X} \subseteq \mathbb{R}^2$ such that for every $w \in \mathcal{X}$ there exists a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y}\big)$ containing a solution (W, A, β, δ, M) to (2) on a stochastic interval $[\![0, \tau]\!]$ for an $\mathbb{F}$-stopping time τ such that*

(i) on $[\![0, \tau \wedge \sigma)\!)$, we have $W \in \mathcal{X}$, (β, δ) enforces A, and $W + r\delta(y) \in \mathcal{W} \;\; \forall\, y \in Y$,

(ii) on $\{\tau \leq \sigma\}$, we have $W_\tau \in \mathcal{B}_r(\mathcal{W})$,

where σ denotes the first jump time of any of the processes $(J^y)_{y\in Y}$. Let $\hat\sigma$ denote the first jump time of any of the processes $(J^y)_{y\in Y}$ strictly after τ. Then for every $w \in \mathcal{X}$, there exists a solution $(\hat W, \hat A, \hat\beta, \hat\delta, \hat M)$ to (2) on $[\![0, \sigma]\!]$ in the same stochastic framework, which coincides with (W, A, β, δ, M) on $[\![0, \tau\wedge\sigma)\!)$, such that $\big(\hat\beta, \hat\delta\big)$ enforces $\hat A$ on $[\![0, \hat\sigma)\!)$, $\hat W \in \mathcal{B}_r(\mathcal{W})$ on $[\![\tau, \hat\sigma)\!)$, and $\hat W_{\hat\sigma} \in \mathcal{W}$ P-a.s. In particular, $\mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$.

Proof. Fix a stochastic framework $\big(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y}\big)$, an $\mathbb{F}$-stopping time τ, and a solution (W, A, β, δ, M) to (2) on $[\![0, \tau]\!]$ with $W_0 = w$ for arbitrary $w \in \mathcal{X}$. Set

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

and denote by $\tilde{\mathbb{F}} := \big(\tilde{\mathcal{F}}_t\big)_{t\geq 0}$ the filtration 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 $\tilde J^y$ for any $y \in Y$ is an $\tilde{\mathbb{F}}$-Poisson process. Therefore, $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde Z, (\tilde J^y)_{y\in Y})$ is a stochastic framework with $W_\tau \in \tilde{\mathcal{F}}_0$. Since

$\mathcal{W}_\tau \in \mathcal{B}_r(\mathcal{W})$ on $\{\tau < \sigma\}$ by assumption, the equivalence in Lemma A.1 implies the existence of a solution $\big(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}, \tilde{M}\big)$ to (2) on $[\![0, \tilde{\sigma}]\!]$ in the stochastic framework $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde{Z}, (\tilde{J}^y)_{y\in Y})$ 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 $\tilde{A}$ and $\tilde{W} \in \mathcal{B}_r(\mathcal{W})$, where $\tilde{\sigma}$ is the first time any of the processes $(\tilde{J}^y)_{y\in Y}$ jump. We define the concatenated processes $\hat{W}$, $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, and $\hat{M}$ by setting

$$\hat{W} := \big(W 1_{[\![0,\tau)\!)} + \tilde{W}_{\cdot-\tau} 1_{[\![\tau,\infty)\!)}\big) 1_{\{\tau\leq\sigma\}} + W 1_{\{\tau>\sigma\}}$$

and similarly for $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, and $\hat{M}$. Since $\tilde{\sigma} = \hat{\sigma}$ a.s., the concatenation $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ is a solution to (2) on $[\![0, \hat{\sigma}]\!]$ in the concatenated stochastic framework defined by

$$\hat{Z} := \big(Z 1_{[\![0,\tau)\!)} + (\tilde{Z} + Z_\tau) 1_{[\![\tau,\infty)\!)}\big) 1_{\{\tau\leq\sigma\}} + Z 1_{\{\tau>\sigma\}}$$

and similarly for $\hat{J}^y$ for every $y \in Y$. Observe that $\hat{Z} = Z$ and $\hat{J}^y = J^y$ for all events y, hence the concatenated stochastic framework is identical to the original framework. Thus, $\big(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}\big)$ is a solution to (2) on $[\![0, \hat{\sigma}]\!]$ in $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$.

By construction, $\big(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}\big)$ coincides with (W, A, β, δ, M) on $[\![0, \tau\wedge\sigma)\!)$, $\big(\hat{\beta}, \hat{\delta}\big)$ enforces $\hat{A}$ on $[\![0, \hat{\sigma})\!)$, and $\hat{W} \in \mathcal{B}_r(\mathcal{W})$ on $[\![\tau, \hat{\sigma})\!)$. It follows from the choice of $\tilde{W}$ that $\hat{W}_\sigma \in \mathcal{W}$ on $\{\tau < \sigma\}$. Since $W + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ on $[\![0, \tau \wedge \sigma)\!)$ and J^y is orthogonal to M, it follows that $\hat{W}_\sigma = \hat{W}_{\sigma-} + r\hat{\delta}_{\sigma-}(y) \in \mathcal{W}$ also on $\{\tau \geq \sigma\}$. This shows that any $w \in \mathcal{X}$ can be attained by an enforceable solution to (2) that remains in $\mathcal{X} \cup \mathcal{B}_r(\mathcal{W})$ until time $\hat{\sigma}$, at which point the solution jumps to $\mathcal{W}$. Therefore, $\mathcal{X}\cup\mathcal{B}_r(\mathcal{W})$ is relaxed self-generating, hence $\mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. □

Proof of Lemma 5.2. We first show that $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$ implies that $\mathcal{B}_r(\mathcal{W})$ is self-generating. Fix any payoff pair $w \in \mathcal{B}_r(\mathcal{W})$. By definition, there exists a solution (W, A, β, δ, M) to (2) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$ such that $W_0 = w$ P-a.s., $W_{\sigma_1} \in \mathcal{W}$ P-a.s., and on $[\![0, \sigma_1)\!)$, we have $W \in \mathcal{B}_r(\mathcal{W})$ and (β, δ) enforcing A. Since $W_{\sigma_1} \in \mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$, we can apply Lemma A.2 to $\tau = \sigma_1$ to obtain a solution $\big(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}\big)$ to (2) on the same stochastic framework such that $\hat{W}_0 = w$ P-a.s., $\hat{W}_{\sigma_2} \in \mathcal{W}$ P-a.s., and on $[\![0, \sigma_2)\!)$, $(\hat{\beta}, \hat{\delta})$ enforces $\hat{A}$ and $\hat{W} \in \mathcal{B}_r(\mathcal{W})$. With an iteration of this concatenation procedure, we can construct solutions to (2) on $[\![0, \sigma_n)\!)$ that remains in $\mathcal{B}_r(\mathcal{W})$ 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 (2) that remains in $\mathcal{B}_r(\mathcal{W})$ forever, showing that $\mathcal{B}_r(\mathcal{W})$ is self-generating.

For the converse, suppose that $\mathcal{W}$ is self-generating. By definition, for any $w \in \mathcal{W}$, there exists an enforceable solution W to (2) that is in $\mathcal{W}$ almost everywhere. In particular, $W_{\sigma_1} \in \mathcal{W}$ a.s. It follows that $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. □

Before we prove the convergence of the algorithm in Proposition 5.3 to $\mathcal{E}(r)$, we need one more auxiliary result, stating that $\mathcal{B}_r(\mathcal{W})$ is monotone in $\mathcal{W}$.

Lemma A.3. *Let $\mathcal{W} \subseteq \mathcal{W}'$. Then $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W}')$.*

Proof. Fix any payoff pair $w \in \mathcal{B}_r(\mathcal{W})$. By definition of $\mathcal{B}_r(\mathcal{W})$, there exists a solution (W, A, β, δ, M) to (2) in a stochastic framework such that $W_0 = w$ a.s., $W_{\sigma_1} \in \mathcal{W}$ a.s., and on $[\![0, \sigma_1)\!)$, (β, δ) enforces A and $W \in \mathcal{B}_r(\mathcal{W})$. Since this implies that also $W_{\sigma_1} \in \mathcal{W}'$ a.s., it follows that $\mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}'$-relaxed self-generating. In particular, $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W}')$ by maximality of $\mathcal{B}_r(\mathcal{W}')$. □

We are now ready to prove Propositions 5.3 and 8.4.

Proof of Proposition 5.3. Lemma 5.2 implies that $\mathcal{E}(r) \subseteq \mathcal{B}_r\big(\mathcal{E}(r)\big)$ and that $\mathcal{B}_r\big(\mathcal{E}(r)\big)$ is self-generating. Since $\mathcal{E}(r)$ is the largest bounded self-generating set, it follows that $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$. Next, we show that $\mathcal{B}_r(\mathcal{V}^*) \subseteq \mathcal{V}^*$, i.e., each player i attains at least his minmax payoff in $\mathcal{B}_r(\mathcal{V}^*)$. Suppose towards a contradiction that this is not the case, i.e., there exists $w \in \mathcal{B}_r(\mathcal{V}^*)$ with $w^i < \underline{v}^i$ for some i. Since $w \in \mathcal{B}_r(\mathcal{V}^*)$, there exists a solution (W, A, β, δ, M) to (2) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$ with $W_0 = w$ a.s., $W_{\sigma_1} \in \mathcal{V}^*$ a.s., such that on $[\![0, \sigma_1)\!)$, we have $W \in \mathcal{B}_r(\mathcal{V}^*)$ and (β, δ) enforces A. By Lemma 4.1, we can write

$$w^i = \int_0^{\sigma_1} r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^A}\big[g^i(A_s)\big]\,\mathrm{d}s + \mathrm{e}^{-r\sigma_1}\mathbb{E}_{Q_{\sigma_1}^A}\big[W_{\sigma_1}^i\big]. \tag{14}$$

Since $w^i < \underline{v}^i$ but $W_{\sigma_1}^i \geq \underline{v}^i$ a.s., it follows that the first term in (14) is smaller than $(1-\mathrm{e}^{-r\sigma_1})\underline{v}^i$. Define now the strategy $\tilde{A}^i = \arg\max_{a^i \in \mathcal{A}^i} g^i(\,\cdot\,, A^{-i})1_{[\![0,\sigma_1)\!)} + A^i 1_{[\![\sigma_1,\infty)\!)}$. Because $\tilde{A}^i$ is i's myopic best reply to A^{-i} on $[\![0, \sigma_1)\!)$, it follows that $g^i(\tilde{A}^i, A^{-i}) \geq \underline{v}^i$ on that same stochastic interval. We deduce that

$$\int_0^{\sigma_1} r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^{\tilde{A}^i,A^{-i}}}\Big[g^i(\tilde{A}_s^i, A_s^{-i})\Big]\,\mathrm{d}s \geq (1-\mathrm{e}^{-r\sigma_1})\underline{v}^i > \int_0^{\sigma_1} r\mathrm{e}^{-rs}\mathbb{E}_{Q_s^A}\big[g^i(A_s)\big]\,\mathrm{d}s.$$

This contradicts the fact that (β, δ) enforces A on $[\![0, \sigma_1)\!)$.[21] We have thus shown that $\mathcal{W}_1 = \mathcal{B}_r(\mathcal{V}^*)$ is contained in $\mathcal{W}_0 = \mathcal{V}^*$. Since $\mathcal{E}(r) \subseteq \mathcal{W}_0$, monotonicity of $\mathcal{B}_r$ implies that $\mathcal{E}(r) \subseteq \mathcal{W}_1 \subseteq \mathcal{W}_0$. An iterated application of Lemma A.3 thus shows that $(\mathcal{W}_n)_{n\geq 0}$ is decreasing in the set-inclusion sense and that it is bounded from below by $\mathcal{E}(r)$. It must converge to a limit $\mathcal{W}_\infty \supseteq \mathcal{E}(r)$ which satisfies $\mathcal{W}_\infty = \mathcal{B}_r(\mathcal{W}_\infty)$. The limit set $\mathcal{W}_\infty$ is thus self-generating and hence $\mathcal{W}_\infty = \mathcal{E}(r)$ by Lemma 4.5. □

Proof of Proposition 8.4. Observe first that the operators $\mathcal{S}_r$ and $\mathcal{K}_{r,a}$ are monotone, that is, $\mathcal{S}_r(\mathcal{W}) \subseteq \mathcal{S}_r(\mathcal{W}')$ and $\mathcal{K}_{r,a}(\mathcal{W}) \subseteq \mathcal{K}_{r,a}(\mathcal{W}')$ for two payoff sets $\mathcal{W} \subseteq \mathcal{W}'$. Because (5) is solved over a larger set of controls in $\tilde{\mathcal{B}}_r(\mathcal{W}')$ than in $\tilde{\mathcal{B}}_r(\mathcal{W})$, it follows that $\tilde{\mathcal{B}}_r(\mathcal{W}') \supseteq \tilde{\mathcal{B}}_r(\mathcal{W})$, i.e., $\tilde{\mathcal{B}}_r$ is monotone as well. The same reasoning shows that $\tilde{\mathcal{B}}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$ for a fixed payoff set $\mathcal{W}$, hence $\mathcal{W}_1 = \tilde{\mathcal{B}}_r(\mathcal{W}_0) \subseteq \mathcal{B}_r(\mathcal{W}_0) \subseteq \mathcal{W}_0$. Together with monotonicity of $\tilde{\mathcal{B}}_r$, this shows that $(\mathcal{W}_n)_{n\geq 0}$ is decreasing in the set-inclusion sense. Therefore, $(\mathcal{W}_n)_{n\geq 0}$ converges to some limit $\mathcal{W}_\infty$.

[21] This follows along the same lines as the proof of Lemma 4.3 by restricting attention to $[\![0, \sigma_1)\!)$.

Corollary 8.2 asserts that $\tilde{\mathcal{B}}_r(\mathcal{E}(r)) = \mathcal{B}_r(\mathcal{E}(r)) = \mathcal{E}(r)$. Monotonicity of $\tilde{\mathcal{B}}_r$ thus shows that $\mathcal{E}(r) \subseteq \mathcal{W}_n$ for any n, hence the limit $\mathcal{W}_\infty$ contains $\mathcal{E}(r)$. It remains to show that $\mathcal{W}_\infty$ is not larger than $\mathcal{E}(r)$. To that effect, let $(\mathcal{W}'_n)_{n\geq 0}$ denote the sequence of iterated applications of $\mathcal{B}_r$ to $\mathcal{V}^*$. Since $\tilde{\mathcal{B}}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$ for any set $\mathcal{W}$, it follows that $\mathcal{W}_1 \subseteq \mathcal{W}'_1$ and hence $\mathcal{W}_{n+1} \subseteq \tilde{\mathcal{B}}_r(\mathcal{W}'_n) \subseteq \mathcal{B}_r(\mathcal{W}'_n) = \mathcal{W}'_{n+1}$ by induction. Thus, $\mathcal{E}(r) \subseteq \mathcal{W}_n \subseteq \mathcal{W}'_n$ for any n and hence $\mathcal{W}_n \to \mathcal{E}(r)$ as $n \to \infty$. □

B Regularity of the optimality equation

The purpose of this appendix is to prove that the optimality equation (5) is locally Lipschitz continuous at almost every point so that, locally, it admits a unique solution. The main results of this appendix, Proposition B.6 and Lemma B.7, will be invoked in Appendix C to show that outside of $\Gamma(r, \mathcal{W})$, the boundary $\partial\mathcal{B}_r(\mathcal{W})$ is the unique solution to the optimality equation. Throughout the entire appendix, we consider the discount rate $r > 0$ and a closed, convex set $\mathcal{W} \subseteq \mathcal{V}$ with non-empty interior as fixed. Consider the optimality equation first for a fixed action profile a:

$$\kappa_a(w, N) = \max_{(\beta,\delta)\in\Xi_a(w,N,r,\mathcal{W})} \frac{2N^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\beta\|^2}. \tag{15}$$

The general idea is to show that κ_a can be suitably rewritten as the maximum of a locally Lipschitz continuous function over a locally Lipschitz continuous set of incentives. We refer to Aubin and Frankowska [7] for a detailed overview of set-valued maps and their properties and state here only the most central property.

Definition B.1. A set-valued map $G : x \mapsto G(x) \subseteq \mathbb{R}^k$ is said to be *Lipschitz continuous* if there exists a constant K such that $G(x) \subseteq G(\tilde{x}) + K\|x - \tilde{x}\| B_1(0)$ for any x and $\tilde{x}$, where $B_1(0)$ denotes the closed unit ball in $\mathbb{R}^k$ centered at the origin.

Lemma B.2. *Let $f(x, y)$ be a single-valued Lipschitz-continuous function and let $G(x)$ be a set-valued (locally) Lipschitz-continuous map. Then $h(x) = \max_{y\in G(x)} f(x, y)$ is (locally) Lipschitz continuous.*

Proof. For any x, let U be a neighbourhood of x such that G is Lipschitz continuous on U with Lipschitz constant K_G. Let $x_1, x_2 \in U$ and suppose without loss of generality that $h(x_1) \geq h(x_2)$. Let K_f be the Lipschitz constant of f. Then $f(x_1, y) \leq f(x_2, y) + K_f\|x_2 - x_1\|$ for any y, hence

$$\begin{aligned} h(x_1) - h(x_2) &\leq K_f\|x_2 - x_1\| + \max_{y\in G(x_1)} f(x_2, y) - \max_{y\in G(x_2)} f(x_2, y) \\ &\leq K_f\|x_2 - x_1\| + \max_{y\in G(x_2)+K_G\|x_2-x_1\|B_1(0)} f(x_2, y) - \max_{y\in G(x_2)} f(x_2, y) \\ &\leq K_f\|x_2 - x_1\| + K_f K_G\|x_2 - x_1\|. \end{aligned}$$

□

Let us begin by reducing the two-variable optimization problem to a one-variable optimization by expressing incentives β in terms of δ and the normal vector N. In the remainder of the paper, we denote by $T(N)$ the vector obtained from rotating N by 90 degrees in the clockwise direction. We will often omit the argument and simply use T, T', $\tilde{T}$, etc. to denote $T(N)$, $T(N')$, $T(\tilde{N})$ and so on. Let $\Phi_a(N,\delta)$ denote the set of all vectors $\phi \in \mathbb{R}^d$ such that $(T\phi, \delta)$ enforces a and denote by $\phi(a,N,\delta)$ the vector of smallest length in $\Phi_a(N,\delta)$.

Lemma B.3. *For any $a \in \mathcal{A}$, the map $(N,\delta) \mapsto \phi(a,N,\delta)$ is locally Lipschitz continuous for $N \notin \{\pm e_1, \pm e_2\}$ and $\delta \in \Psi_a$.*

Proof. Let $G^i(a)$ denote the row vector with entries $g^i(\tilde{a}^i, a^{-i}) - g^i(a)$ and let $\Lambda^i(a)$ denote the matrix with column vectors $\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)$. For non-coordinate N, $\Phi_a(N,\delta)$ is the solution to $\phi M^i(a) \leq -\frac{1}{T^i}\big(G^i(a) + \delta^i \Lambda^i(a)\big)$. Since $M^i(a)$ does not depend on N or δ and the right-hand side is locally Lipschitz continuous in N, δ, the set $\Phi_a(N,\delta)$ is a constant-rank polyhedron with locally Lipschitz continuous right-hand side. It follows from the main result of Yen [27] that the projection of 0 onto $\Phi_a(N,\delta)$ is locally Lipschitz continuous in (N,δ). □

Lemma B.3 simplifies the constraints in the maximization in (15) by reducing it to a maximization over δ only. We will prove regularity of a slightly more general form of the optimality equation suitable for the proofs in Appendix C. For any $h \geq 0$, let $\mathcal{W}_h := \{v \in \mathbb{R}^2 \mid \min_{w\in\mathcal{W}} \|v - w\| \leq h\}$ denote the set of all payoff pairs within distance h from $\mathcal{W}$. Call a set-valued map $\mathcal{L} : \mathbb{R}^2 \to \mathbb{R}^{2\times|Y|}$ a *Lipschitz expansion* if

$$\mathcal{L}(w) = \big\{\delta \in \mathbb{R}^{2\times|Y|} \;\big|\; w + r\delta(y) \in \mathcal{W}_{h(w)} \text{ for all } y \in Y\big\}$$

for a non-negative Lipschitz-continuous function $h : \mathbb{R}^2 \to \mathbb{R}$ with Lipschitz constant K_h. Observe that a Lipschitz expansion is Lipschitz continuous in the sense of set-valued maps with Lipschitz constant $\sqrt{|Y|}(K_h + 1)$. We denote by $\mathcal{L}_0$ the trivial Lipschitz expansion with $h \equiv 0$. Set $\Psi_a(w,N,\mathcal{L}) := \{\delta \in \Psi_a \cap \mathcal{L}(w) \mid \Phi_a(N,\delta) \neq \emptyset\}$ and consider the optimality equation in the form

$$\kappa_{a,\mathcal{L}}(w,N) = \max_{\delta\in\Psi_a(w,N,\mathcal{L})} \frac{2N^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\phi(a,N,\delta)\|^2} \vee 0. \tag{16}$$

Observe that ODE (16) for $\mathcal{L}_0$ reduces to (15) if the expression in the right-hand side of (15) is understood to be 0 when $\Psi_a(w,N,\mathcal{L}) = \emptyset$. Denote by

$$E_a(\mathcal{L}) := \big\{(w,N) \in \mathbb{R}^2 \times S^1 \;\big|\; \Psi_a(w,N,\mathcal{L}) \neq \emptyset\big\}$$

the effective domain of of $(w,N) \mapsto \Psi_a(w,N,\mathcal{L})$. For the remainder of the analysis, it will be convenient to denote by $\mathcal{P} := \mathbb{R}^2 \times \{\pm e_1, \pm e_2\}$ the set of all *poles*, i.e., all payoff-direction pairs with coordinate directions.

Lemma B.4. *Suppose that Assumptions 2 and 3 are satisfied. Then there exists* $\varepsilon > 0$ *such that* $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \subseteq E_a(\mathcal{L})$ *for any action profile* a *with* $E_a(\mathcal{L}) \neq \emptyset$ *and any Lipschitz expansion* $\mathcal{L}$*. Moreover,* $(w, N) \mapsto \Psi_a(w, N, \mathcal{L})$ *is locally Lipschitz continuous on* $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\})$*.*

Proof. Any enforceable action profile a is pairwise identifiable by Assumption 2. Fix any $\delta \in \Psi_a$ and any non-coordinate N. It follows from Lemma 2 in Sannikov [22] applied to payoff function $\tilde{g}(a) = g(a) + \delta\lambda(a)$ that there exists β with $N^\top \beta = 0$ such that (β, δ) enforces a. In particular, $\Phi_a(N, \delta) \neq \emptyset$ and hence $\Psi_a(w, N, \mathcal{L}) = \Psi_a \cap \mathcal{L}(w)$ for any non-coordinate N. This implies that $N \mapsto \Psi_a(w, N, \mathcal{L})$ is locally Lipschitz continuous and that $E_a(\mathcal{L})$ contains $\{w \mid \Psi_a \cap \mathcal{L}(w) \neq \emptyset\} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$.

By Assumption 3, there exists ε such that $\delta \in \Psi_a$ for any δ with $\|\delta(y)\| \leq \varepsilon$ and any enforceable action profile a. Fix any $w \in \mathcal{W}_{r\varepsilon/2}$ and let $v \in \mathcal{W} \cap B_{3r\varepsilon/4}(w)$ with $B_{r\varepsilon/4}(v) \subseteq \mathcal{W}$. In particular,

$$\frac{v - w}{r}\mathbf{1} + \big(B_{\varepsilon/4}(v)\big)^{|Y|} \subseteq \Psi_a \cap \mathcal{L}_0(w) \subseteq \Psi_a \cap \mathcal{L}(w)$$

shows that $\Psi_a \cap \mathcal{L}(w)$ has non-empty interior, where $\mathbf{1}$ is the $|Y|$-dimensional row vector $\mathbf{1} = (1, \ldots, 1)$. This shows that $E_a(\mathcal{L})$ contains $\mathcal{W}_{r\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. Moreover, since Ψ_a and $\mathcal{L}(w)$ are both closed and convex, Lemma F.1 implies that $w \mapsto \Psi_a \cap \mathcal{L}(w)$ is locally Lipschitz continuous at $w \in \mathcal{W}_{r\varepsilon/2}$. □

The last ingredient to establish local Lipschitz continuity of $\kappa_{a,\mathcal{L}}$ is to characterize the set of payoff-direction pairs, for which the right-hand side of (16) is unbounded. This is the case on the set $\Gamma_a(\mathcal{L})$ of all (w, N), for which there exists a $\delta \in \Psi_a(w, N, \mathcal{L})$, for which $\phi(a, N, \delta) = 0$ and $N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$. Set $\Gamma(\mathcal{L}) := \bigcup_{a \in \mathcal{A}} \Gamma_a(\mathcal{L})$.

Remark B.1. Note that $\Gamma(\mathcal{L}_0) = \Gamma(r, \mathcal{W})$, where $\Gamma(r, \mathcal{W})$ is defined in Definition 6.5.

Lemma B.5. *Let* $\mathcal{L}$ *be a Lipschitz expansion. Then* $\Gamma_a(\mathcal{L})$ *is closed for any* $a \in \mathcal{A}$*. In particular,* $\Gamma(\mathcal{L})$ *is closed as the finite union of closed sets.*

Proof. Let $\Psi_a^0(w, N)$ denote the set of all $\delta \in \Psi_a$, for which $(0, \delta)$ enforces a and $N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$. Clearly, $(w, N) \mapsto \Psi_a^0(w, N)$ is closed-valued and upper semi-continuous, hence so is $(w, N) \mapsto \Psi_a^0(w, N) \cap \mathcal{L}(w)$. Consider a sequence $(w_n, N_n)_{n \geq 0}$ in $\Gamma_a(\mathcal{L})$ converging to (w, N). There exists a sequence $(\delta_n)_{n \geq 0}$ with $\delta_n \in \Psi_a^0(w_n, N_n) \cap \mathcal{L}(w_n)$. Because $\mathcal{L}(w_n)$ is uniformly bounded by $\mathcal{L}(\mathcal{V})$, the sequence $(\delta_n)_{n \geq 0}$ is uniformly bounded as well. Therefore, $(\delta_n)_{n \geq 0}$ converges along a subsequence $(n_k)_{k \geq 0}$ to some finite limit δ. Since $\Psi_a^0 \cap \mathcal{L}$ is closed-valued and upper semi-continuous, it follows that $\delta \in \Psi_a^0(w, N) \cap \mathcal{L}(w)$, hence $(w, N) \in \Gamma_a(\mathcal{L})$. □

We are now ready to establish local Lipschitz continuity of $\kappa_{a,\mathcal{L}}$.

Proposition B.6. *Suppose that Assumptions 2 and 3 hold. For any convex* $\mathcal{W}$ *with* $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$ *and any Lipschitz expansion* $\mathcal{L}$*, there exists* $\varepsilon > 0$ *such that*

$$\kappa_{\mathcal{L}}(w, N) = \max_{a\in\mathcal{A}} \kappa_{a,\mathcal{L}}(w, N) \tag{17}$$

is locally Lipschitz continuous on $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \setminus \Gamma(\mathcal{L})$.

Proof. Let $\varepsilon > 0$ be given by Lemma B.4 so that $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \subseteq E_a(\mathcal{L})$ for any action profile a with $E_a(\mathcal{L}) \neq \emptyset$ and that $\Psi_a(w, N, \mathcal{L})$ is locally Lipschitz continuous on $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. It follows from Lemmas B.2 that $\kappa_{a,\mathcal{L}}$ is locally Lipschitz continuous at any $(w, N) \in \mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \setminus \Gamma(\mathcal{L})$ for any action profile a. Therefore, $\kappa_{\mathcal{L}}$ is locally Lipschitz continuous except at $\bigcup_{a\in\mathcal{A}} \partial E_a(\mathcal{L})$. Since the intersection of $\bigcup_{a\in\mathcal{A}} \partial E_a(\mathcal{L})$ with $\mathcal{W}_{\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$ is empty by Lemma B.4, $\kappa_{\mathcal{L}}$ is locally Lipschitz continuous on $\mathcal{W}_{\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. □

Finally, to guarantee that $\partial\mathcal{B}_r(\mathcal{W})$ is well-behaved at coordinate directions, we will need the following Lemma.

Lemma B.7. *Suppose that Assumptions 2 and 3 are satisfied. There exists* $\varepsilon > 0$ *such that for any* $w \in \mathcal{W}_\varepsilon \cap \operatorname{int} \mathcal{V}^*$ *and any* $N \in \{\pm e_1, \pm e_2\}$*, there exists* $a \in \mathcal{A}$ *such that* $\kappa_a(w, N) > 0$ *in a neighborhood of* $\mathcal{U}$ *of* (w, N) *with* $\mathcal{U} \subseteq E_a(\mathcal{L}_0)$.

Proof. By Assumption 3, the minmax profile $\underline{a}_i$ against player i and the global maximizer $\bar{a}_i$ of g^i are enforceable and there exists $\varepsilon > 0$ sufficiently small such that $\Delta_\varepsilon := \{\delta \mid \|\delta(y)\| \leq \varepsilon \text{ for every } y \in Y\}$ is contained in $\Psi_{a(N)}$ for every $N \in \{\pm e_1, \pm e_2\}$, where we denote $a(e_i) = \bar{a}_i$, $a(-e_i) = \underline{a}_i$. Let $G^i(a)$ and $\Lambda^i(a)$ be defined as in the proof of Lemma B.3. If Assumption 3.(ii).(a) is satisfied, we additionally choose ε small enough such that $G^i(a(N)) + \delta^i\Lambda^i(a(N)) < 0$ for player i with static best reply in $a(N)$ and $\delta \in \Delta_\varepsilon$. Let $k \leq 1$ be such that for any $w \in \mathcal{W}_{rk\varepsilon} \cap \partial\mathcal{V}^*$ with outward normal vector N, there exists $\delta \in \Delta_\varepsilon$ with $\delta^i = 0$ such that $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$. For any $w \in \mathcal{W}_{rk\varepsilon} \cap \operatorname{int} \mathcal{V}^*$, choose $\delta \in \Delta_\varepsilon$ with $w + r\delta(y) \in \mathcal{W}$ for every event $y \in Y$ such that $N^\top\big(g(a(N)) + \delta\lambda(a(N)) - w\big) > 0$. Since $\delta \in \Psi_{a(N)}$, there exists β such that (β, δ) enforces $a(N)$. If Assumption 3.(ii).(a) holds, the choice of ε implies that $\big(\frac{\tilde{T}}{\tilde{T}^{-i}}\beta^{-i}, \delta\big)$ enforces $a(N)$ for $\tilde{T}$ orthogonal to $\tilde{N}$ sufficiently close to N. In particular, $(w, N) \in \operatorname{int} E_{a(N)}(\mathcal{L}_0)$ and $\kappa_a(w, N) > 0$ since the numerator is strictly positive. If Assumption 3.(ii).(b) holds, then player i may have more than one static best reply to $a^{-i}(N)$. Let $a_+(N)$ and $a_-(N)$ denote action profiles a with $a^{-i} = a^{-i}(N)$ that maximize and minimize $\sum_y \lambda(y \mid a)$, respectively, among all static best replies to $a^{-i}(N)$. For any $w \in \mathcal{W}_{rk\varepsilon}$, choose now $\delta \in \Delta_\varepsilon$ with $w + r\delta(y) = v$ for some $v \in \mathcal{W}$ with $N^\top\big(g(a(N)) + \delta\lambda(a(N)) - w\big) > 0$. If $v^i \geq w^i$, then $G^i(a_+(N)) + \delta^i\Lambda^i(a_+(N)) \leq 0$ and if $v^i \leq w^i$, then $G^i(a_-(N)) + \delta^i\Lambda^i(a_-(N)) \leq 0$. Since $\delta \in \Psi_{a(N)}$, there exists β such that (β, δ) enforces $a(N)$. Let β' denote the projection of β^{-i} onto $\operatorname{span} M^{-i}(a)$. By Assumption 3.(ii).(b), $\frac{\tilde{T}}{\tilde{T}^{-i}}\beta' M^i(a) = 0$, hence $\big(\frac{\tilde{T}}{\tilde{T}^{-i}}\beta', \delta\big)$ enforces $a(N)$ for $\tilde{T}$ orthogonal to $\tilde{N}$ sufficiently close to N. □

C Proof of Lemma 6.6

In this appendix, we show that the boundary of $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$ outside of $\Gamma(r, \mathcal{W})$ satisfies the optimality equation. Because the continuous part of the signal is what creates the curvature, these steps are similar in ideas to Sannikov [22]. We begin with the proof of Lemma 6.4, showing how to construct Markovian strategy profiles whose continuation values remain on a given curve.

Proof of Lemma 6.4. Let $\mathcal{C}$ be a curve as indicated in Lemma 6.4. Fix w in the relative interior of $\mathcal{C}$ and choose $\eta > 0$ small enough such that ${N_w}^\top N_v > 0$ for all $v \in \mathcal{C} \cap B_\eta(w)$, where $B_\eta(w)$ denotes the closed ball around w with radius η. On $B_\eta(w)$, $\mathcal{C}$ admits a local parametrization f in the direction N_w. For any $v \in B_\varepsilon(w)$, define the orthogonal projection $\hat{v} = {T_w}^\top v$ onto the tangent and denote by $\pi(v) = \big(\hat{v}, f(\hat{v})\big)$ the projection of $v \in B_\eta(w)$ onto $\mathcal{C}$ in the direction N_w. Let $\big(W, A, \beta, \delta, Z, (J^y)_{y\in Y}, M\big)$ be a weak solution to (2) starting at $W_0 = w$ with $A = a^*\big(\pi(W_-)\big)$, $\beta = \beta^*\big(\pi(W_-)\big)$, $\delta = \delta^*\big(\pi(W_-)\big)$, and $M \equiv 0$ on $[\![0, \tau)\!)$, where we set $\tau := \sigma_1 \wedge \inf\{t \geq 0 \mid W_t \notin B_\eta(w)\}$ for the first jump time σ_1 of any of the processes $(J^y)_{y\in Y}$. Since π is measurable, the processes A, β and δ are all predictable.

We measure the distance of W to $\mathcal{C}$ by $D_t = N^\top W_t - f\big(\hat{W}_t\big)$. Note that f is differentiable by assumption and $\big(-f'(\hat{W}_t), 1\big) = \ell_t N_t$, where $\ell_t := \big\|\big(-f'(\hat{W}_t), 1\big)\big\|$. Since f is locally convex it is second order differentiable at almost every point by Alexandrov's Theorem. In particular, f' has Radon-Nikodým derivative $f''(\hat{W}_t) = -\kappa(\pi(W_t))\ell_t^3$. It follows from the Meyer-Itō formula (see Theorem 19.5 in Kallenberg [16]) that

$$\begin{aligned} \mathrm{d}D_t &= r\ell_t {N_t}^\top\big(W_t - g(A_t) - \delta_t\lambda(A_t)\big)\,\mathrm{d}t + r\ell_t {N_t}^\top\beta_t\left(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t\right) \\ &\quad + r\ell_t \sum_{y\in Y} {N_{t-}}^\top\delta^*\big(\pi(W_t); y\big)\,\mathrm{d}J_t^y - \frac{1}{2}f''(\hat{W}_{t-})\,\mathrm{d}[\hat{W}]_t, \end{aligned}$$

where we abbreviated $N_t = N_{\pi(W_t)}$ and $T_t = T_{\pi(W_t)}$. The volatility term is zero since $N^\top\beta = 0$. Note that on $[\![0, \sigma_1)\!)$, $\Delta J^y \equiv 0$ for any $y \in Y$ implies that $[\hat{W}] = \langle\hat{W}\rangle$. Using (4) and the fact that ${N_w}^\top N_t = {T_w}^\top T_t = \ell_t^{-1}$, we obtain that on $[\![0, \tau)\!)$,

$$\mathrm{d}D_t = r\ell_t {N_t}^\top\big(W_t - g(A_t) - \delta_t\lambda(A_t)\big)\,\mathrm{d}t + \frac{r^2}{2}\kappa\big(\pi(W_t)\big)\,\ell_t^3\,\big|{T_w}^\top T_t\big|^2|\beta_t|^2\,\mathrm{d}t = rD_t\,\mathrm{d}t,$$

where we used ${N_t}^\top\big(W_t - \pi(W_t)\big) = {N_t}^\top N_w D_t = \ell_t^{-1}D_t$ in the second equality. It follows that $D_t = D_0\mathrm{e}^{rt} = 0$ since $D_0 = 0$. On $\{\tau < \sigma_1\}$ we can repeat this procedure and concatenate the solutions as in the proof of Lemma A.2 to obtain a solution to (2) that remains on $\mathcal{C}$ until either an event y occurs or an end point of $\mathcal{C}$ is reached.

Let ρ denote the hitting time of an end point of $\mathcal{C}$. Note that we have constructed a solution to (2) such that on $[\![0, \rho \wedge \sigma_1)\!)$, we have $(\beta, \delta) \in \Xi_A(\pi(W), N_W, r, \mathcal{W})$ and $W \in \mathcal{C}$. The latter implies that $\pi(W) = W$ and hence $(\beta, \delta) \in \Xi_A(W, N_W, r, \mathcal{W})$.

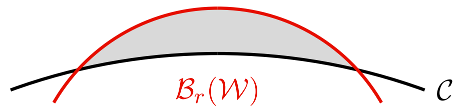

Figure 18: A solution $\mathcal{C}$ to (5) that cuts through $\mathcal{B}_r(\mathcal{W})$.

It follows that $W_{\sigma_1} \in \mathcal{W}$ on the event $\{\sigma_1 \leq \rho\}$. On the event $\{\rho < \sigma_1\}$, we apply Lemma A.2 for stopping time ρ to extend (W, A, β, δ, M) to a $\mathcal{W}$-enforceable solution to (2) such that $W \in \mathcal{B}_r(\mathcal{W})$ on $[\![\rho, \sigma_1]\!)$. Since $\mathcal{B}_r(\mathcal{W})$ is relaxed self-generating, it follows that any payoff pair in $\mathcal{C} \cup \mathcal{B}_r(\mathcal{W})$ can be attained by a $\mathcal{W}$-enforceable solution W to (2) that remains in $\mathcal{C} \cup \mathcal{B}_r(\mathcal{W})$ until time σ_1. Therefore, $\mathcal{C} \cup \mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}$-relaxed self-generating, hence $\mathcal{C} \cup \mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. □

Corollary C.1. *Let $\mathcal{C}$ be a C^1 solution to* (5) *with positive curvature throughout such that $\mathcal{C}$ is closed or its endpoints are in $\mathcal{B}_r(\mathcal{W})$. Then $\mathcal{C} \subseteq \mathcal{B}_r(\mathcal{W})$.*

Proof. For any $w \in \mathcal{C}$, let $a^*(w)$, $\beta^*(w)$, and $\delta^*(w)$ denote the maximizers in (5). Since $\mathcal{C}$ is assumed to have positive curvature throughout, the maximization in (5) is not taken over empty sets. By Lemmas B.4 and B.7, the maximizers are attained. The result thus follows from Lemma 6.4. □

The following result is known as the "escaping lemma," stating that it is impossible for a solution $\mathcal{C}$ to the optimality equation to cut through $\mathcal{B}_r(\mathcal{W})$ as indicated in Figure 18. The main idea is that the optimality equation maximizes the curvature over all enforceable action profiles and all restricted-enforceable incentives. Thus, the continuation value of no enforceable strategy profile that attains a payoff pair above $\mathcal{C}$—in the shaded area of Figure 18—can cross $\mathcal{C}$ unless the continuation value jumps across $\mathcal{C}$ due to the occurrence of a discrete event. Because arrival times of Poisson jumps are arbitrarily large with positive probability and Brownian motion has inifinite variation, the continuation value escapes $\partial\mathcal{B}_r(\mathcal{W})$ to the sides with positive probability before a discrete event occurs.

The presence of abrupt information complicates the argument in two ways. First, simply because Brownian incentives are needed to provide incentives on the curve does not imply that they are needed in the shaded area as well. Second, the curvature in the optimality equation is the maximum over all restricted-enforceable incentives at that point on the curve. Incentives provided in an enforceable strategy profile with continuation value in the shaded area may not come from the same set of incentives because the set of admissible jump incentives depends on the location of the payoff pair. Luckily, both problems can be solved by considering a Lipschitz expansion over jump incentives that guarantees that the set of admissible jump incentives on the curve contains the set of admissible jumps incentives in the shaded area.

Lemma C.2. *Let $w \in \partial\mathcal{B}_r(\mathcal{W})$ with outward normal N'. Let $\pi : U \to \partial\mathcal{B}_r(\mathcal{W})$ be the projection of a neighborhood U of w onto $\partial\mathcal{B}_r(\mathcal{W})$ in the direction of N' and set*

$$\mathcal{L}_k(v) := \left\{\delta \in \mathbb{R}^{2\times m} \;\middle|\; v + r\delta(y) \in \mathcal{W}_{k\|\pi(v)-v\|} \text{ for all } y \in Y\right\} \tag{18}$$

Let $\mathcal{C}$ be a C^1-solution to (17) for $\mathcal{L}_k$ with $k \geq 1$, oriented by $v \mapsto N_v$ with end points $v_L, v_R \in U$. It is impossible that the following properties hold simultaneously:

(i) $v_L + \varepsilon N' \notin \mathcal{B}_r(\mathcal{W})$ *and* $v_L + \varepsilon N' \notin \mathcal{B}_r(\mathcal{W})$ *for any* $\varepsilon > 0$,

(ii) there exists $v_0 \in \mathcal{C}$ *such that* $v_0 + \eta N' \in \mathcal{B}_r(\mathcal{W})$ *for some* $\eta > 0$,

(iii) $\inf_{v\in\mathcal{C}} {N_v}^\top N' > 0$,

(iv) $\inf_{v\in\mathcal{C}}|N_v^i| > 0$ *for* $i = 1, 2$,

(v) $\mathcal{N}_\mathcal{C} \cap \Gamma(\mathcal{L}_k) = \emptyset$,

Proof. Suppose that there exists such a curve $\mathcal{C}$ for some $k \geq 1$. Since $\mathcal{L}_k$ is a Lipschitz expansion, it follows from Conditions (iv), (v) and Proposition B.6 that $\mathcal{C}$ is twice differentiable. By Condition (iii), there exists a local parametrization f of $\mathcal{C}$ in the direction N'. Let T' denote the counterclockwise rotation of N' by 90° and denote by $\hat{v} = {T'}^\top v$ the projection onto the tangent for any $v \in U$. Let $\hat{\pi} : U \to \mathcal{C}$ denote the projection of U onto $\mathcal{C}$ in direction N', i.e., $\hat{\pi}(v) = \big(\hat{v}, f(\hat{v})\big)$. By Condition (ii), there exists a solution $\big(W, A, \beta, \delta, M, Z, (J^y)_{y\in Y}\big)$ to (2) starting at $W_0 = v_0 + \eta N'$ such that on $[\![0, \sigma_1)\!)$, we have $W \in \mathcal{B}_r(\mathcal{W})$ and (β, δ) enforces A with $\delta \in \mathcal{L}_0(W)$. Let $\tau_1 := \inf\{t \geq 0 \mid W_t \notin U\}$ denote the time when W leaves the neighborhood U.

For the sake of brevity, let $N := N_{\hat{\pi}(W)}$ and $T := T_{\hat{\pi}(W)}$ denote the normal and tangent vectors to $\mathcal{C}$ at $\hat{\pi}(W)$. Observe that since $W \in U$ on $[\![0, \tau_1)\!)$, the projection $\hat{\pi}(W)$ and hence N and T are well defined on $[\![0, \tau_1)\!)$. We measure the distance of W to $\mathcal{C}$ by $D := {N'}^\top W - f\big(\hat{W}\big)$. Denote $\ell := 1/(T^\top T')$ and $\gamma := \ell N^\top T'$ for the sake of brevity and observe that $\bar{\gamma} := \sup_{w\in\mathcal{C}} {N_w}^\top T'/({T_w}^\top T') < \infty$ by Condition (iii). Then, similarly as in Footnote 3 of Hashimoto [13], it follows from Itō's formula that

$$D_t \geq D_0 + \int_0^t \zeta_s \,\mathrm{d}s + \int_0^t \xi_s\big(\mathrm{d}Z_s - \mu(A_s)\,\mathrm{d}s\big) + \sum_{y\in Y}\int_0^t \rho_s(y)\mathrm{d}J_s^y + \tilde{M}_t,$$

where $\xi_t = r\ell_t {N_t}^\top \beta_t$, $\rho_t(y) = r\ell_{t-}{N_{t-}}^\top \delta_t(y)$ for every $y \in Y$,

$$\begin{aligned}\zeta_t &= r\ell_t\Big({N_t}^\top\big(W_t - g(A_t) - \delta_t\lambda(A_t)\big) + \frac{r}{2}\kappa\big(\hat{\pi}(W_t)\big)\big\|{T_t}^\top\beta_t + \gamma_t {N_t}^\top\beta_t\big\|^2\Big)\\ &= rD_t + r\ell_t\Big({N_t}^\top\big(\hat{\pi}(W_t) - g(A_t) - \delta_t\lambda(A_t)\big) + \frac{r}{2}\kappa\big(\hat{\pi}(W_t)\big)\big\|{T_t}^\top\beta_t + \gamma_t {N_t}^\top\beta_t\big\|^2\Big),\end{aligned}$$

and $\tilde{M}_t = \int_0^t r\ell_{s-} N_{s-}{}^\top \mathrm{d}M_s$. Let $\tau_2 := \inf\{t \geq 0 \mid D_t \leq 0\}$ denote the first time when W crosses $\mathcal{C}$. Observe that $\tau_2 \leq \tau_1$ a.s. by Condition (i). We proceed to show that W escapes $\mathcal{V}$ and hence $\mathcal{B}_r(\mathcal{W})$ with positive probability before τ_2, a contradiction.

Specifically, we show that there exists an equivalent probability measure R, under which the drift rate of D is bounded from below by rD. Then, D becomes arbitrarily large with positive R-probability, and hence positive Q^A-probability. Because it may take arbitrarily long until an infrequent occurs, it follows that W escapes the bounded set $\mathcal{V}$ with positive probability. Let $\Xi_1 \subseteq [\![0, \tau_2)\!)$ denote the set of all (ω, t), for which $N_t(\omega)^\top\big(\hat{\pi}(W_t(\omega)) - g(A_t(\omega)) - \delta_t(\omega)\lambda(A_t(\omega))\big) \geq 0$. Since $\zeta_t(\omega) \geq rD_t(\omega)$ for $(\omega, t) \in \Xi_1$, there is no need to change the probability measure on Ξ_1. Let $\Xi_1^c = [\![0, \tau_2)\!) \setminus \Xi_1$ denote the complement of Ξ_1 relative to $[\![0, \tau_2)\!)$. Because W has not crossed $\mathcal{C}$ yet, it follows that $\|\hat{\pi}(W_t(\omega)) - W_t(\omega)\| \leq \|\hat{\pi}(W_t(\omega)) - \pi(W_t(\omega))\|$ for $(\omega, t) \in [\![0, \tau_2)\!)$ and hence $\delta_t(\omega) \in \Psi_{A_t(\omega)} \cap \mathcal{L}_k(\hat{\pi}(W_t(\omega)))$. This implies that $\beta_t(\omega) \neq 0$ for $(\omega, t) \in \Xi_1^c$ as otherwise $(\hat{\pi}(W_t(\omega)), N_t(\omega)) \in \Gamma(\mathcal{L}_k)$, contradicting Condition (v). It follows from Condition (iv) and Lemma B.4 that $\big(\hat{\pi}(W), N\big) \in E_A(\mathcal{L}_k)$ on $[\![0, \tau_2)\!)$. Since $\kappa\big(\hat{\pi}(W)\big)$ solves (17), it follows that

$$\zeta \geq rD - r\ell N^\top\big(g(A) + \delta\lambda(A) - \hat{\pi}(W)\big)\left(1 - \frac{\big(\|T^\top\beta\| - \gamma\|N^\top\beta\|\big)^2}{\big\|\phi\big(A, N, \delta\big)\big\|^2}\right).$$

Denote $\Lambda := \max_{a\in\mathcal{A}} \sum_{y\in Y} \lambda(y|a)$ and observe that $N^\top\big(g(A) + \delta\lambda(A) - \hat{\pi}(W)\big)$ is uniformly bounded above by the constant $K_1 := \operatorname{diam}\mathcal{V} + \sup(\mathcal{W} - \mathcal{V})\Lambda < \infty$. Due to Lemma F.2, there exist constants $K_2, \bar{\Psi}$ such that

$$\zeta_t(\omega) \geq rD_t(\omega) - r\ell_t(\omega)K_1\frac{2K_2 + 2\bar{\gamma}}{\bar{\Psi}}\big\|N_t(\omega)^\top\beta_t(\omega)\big\|$$

for $(\omega, t) \in \Xi_1^c$. Let $T := \min\big\{t \geq 0 \;\big|\; D_0(1 + rt)/2 \geq \sup_{w\in\mathcal{V}} N^\top w - f(\hat{w})\big\}$ and observe that T is deterministic. We define a density process L on $[0, T]$ by setting

$$\frac{\mathrm{d}L_t}{L_t} = \big(\psi_t - \mu(A_t)\big)\,\mathrm{d}Z_t + \sum_{y\in Y}\left(\frac{1}{\lambda(y|A_{t-})} - 1\right)\mathrm{d}J_t^y,$$

where

$$\psi_t = K_1\frac{2K_2 + 2\bar{\gamma}}{\bar{\Psi}}\frac{\xi_t}{\|\xi_t\|}1_{\Xi_1^c}.$$

Because $\int_0^T \|\psi_t\|^2\,\mathrm{d}t < \infty$ Q_T^A-a.s., it follows from Girsanov's theorem that L defines a probability measure R equivalent to Q_T^A on $\mathcal{F}_T$ such that $\mathrm{d}Z'_t = \mathrm{d}Z_t - \psi_t\,\mathrm{d}t$ is an R-Brownian motion on $[0, T]$, such that J^y has intensity 1 for every $y \in Y$, and $\tilde{M}_t$ is an R-martingale because it is orthogonal to L. Then

$$D_t \geq D_0 + \int_0^t rD_s\,\mathrm{d}s + \int_0^t \xi_s{}^\top\,\mathrm{d}Z'_s + \sum_{y\in Y}\int_0^t \rho_s(y)\,\mathrm{d}J_s^y + \tilde{M}_t. \tag{19}$$

Since W is bounded, $\int_0^t \xi_s^\top \, \mathrm{d}Z_s$ is a $BMO(Q_T^A)$-martingale. Therefore, $\int_0^t \xi_s^\top \, \mathrm{d}Z'_s$ is a $BMO(R)$-martingale by Theorem 3.6 in Kazamaki [18]. Define the stopping time $\tau_3 := \inf\{t \geq 0 \mid D_t \leq D_0(1+rt)/2\}$ and observe that $\tau_3 \leq \tau_2 \wedge T$. It follows from (19) that

$$D_{\tau_3} - \frac{D_0}{2}(1 + r\tau_3) \geq \frac{D_0}{2} + F_{\tau_3} + \sum_{y \in Y} \int_0^{\tau_3} \rho_s(y) \, \mathrm{d}J_s^y,$$

where $F_t = \int_0^t \xi_s \, \mathrm{d}Z'_s + \tilde{M}_t$ is an R-martingale starting at 0. Define the R-martingale $G_t := \mathrm{e}^{|Y|t} 1_{\{t<\sigma\}}$ and observe that G is orthogonal to F. Because $\tau_3 \leq T$ a.s.,

$$\begin{aligned} 0 &\geq \mathbb{E}_R\left[\left(D_{\tau_3} - \frac{D_0}{2}(1+r\tau_3)\right)1_{\{T<\sigma\}}\right] \geq \mathbb{E}_R\left[\frac{D_0}{2}1_{\{T<\sigma\}} + F_{\tau_3}1_{\{T<\sigma\}}\right] \\ &= \frac{D_0}{2}R(T<\sigma) + \mathrm{e}^{-|Y|T}\mathbb{E}_R[F_{\tau_3}G_T] > 0, \end{aligned}$$

where the last inequality follows from the optional stopping theorem and because R is equivalent to Q_T^A on $\mathcal{F}_T$. This is a contradiction. □

A first application of Lemma C.2 shows that the boundary of $\mathcal{B}_r(\mathcal{W})$ is smooth outside of $\Gamma(r, \mathcal{W})$. If it did have a corner, we could construct a solution to (17) that cuts through that corner, which is an impossibility by Lemma C.2.

Lemma C.3. *For any corner w of $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$, we have $(w, N) \in \Gamma(r, \mathcal{W})$ for every outward normal vector $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$.*

Proof. Fix a corner w of $\mathcal{B}_r(\mathcal{W})$ and set

$$\mathcal{N}_0 := \operatorname{int} \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \setminus \{\pm e_1, \pm e_2\}.$$

Fix $N' \in \mathcal{N}_0$ and let π denote the projection of a neighborhood U of w onto $\partial\mathcal{B}_r(\mathcal{W})$ in the direction of N'. For $\eta > 0$ and $k \geq 1$, let $\mathcal{C}_{\eta,k}$ be a C^1 solution to (15) for $\mathcal{L}_k$ defined in (18) starting at (w_η, N'), where $w_\eta := w - \eta N'$. Since N' is in the interior of $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$, the curve $\mathcal{C}_{\eta,k}$ satisfies Conditions (i)–(iv) of Lemma C.2 for any $\eta > 0$ sufficiently small. Suppose towards a contradiction that $(w, N) \notin \Gamma(r, \mathcal{W})$ and hence $(w, N) \notin \Gamma(\mathcal{L}_k)$. Since $\Gamma(\mathcal{L}_k)$ is closed by Lemma B.5, there exists a neighborhood $\mathcal{U}$ of (w, N) with $\mathcal{U} \cap \Gamma(\mathcal{L}_k) = \emptyset$. Therefore, for η sufficiently small such that $\mathcal{N}_{\mathcal{C}_{\eta,k}} \subseteq \mathcal{U}$, Condition (v) of Lemma C.2 is satisfied. Such a curve cannot exist by Lemma C.2, which must be in contradiction to $(w, N') \in \Gamma(r, \mathcal{W})$. Finally, $(w, N) \in \Gamma(r, \mathcal{W})$ for any $N \in \operatorname{cl} \mathcal{N}_0 = \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ since $\Gamma(r, \mathcal{W}) = \Gamma(\mathcal{L}_0)$ is closed. □

Lemma C.4. *Suppose that $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$. The boundary $\partial\mathcal{B}_r(\mathcal{W})$ coincides with a solution to* (5) *in a neighborhood of any $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \setminus (\Gamma(r, \mathcal{W}) \cup \mathcal{P})$. Moreover, $\mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \mathcal{P} \setminus \Gamma(r, \mathcal{W})$ consists of at most* 4 *isolated points.*

Proof. Fix $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \setminus (\Gamma(r, \mathcal{W}) \cup \mathcal{P})$. Due to Proposition B.6, the optimality equation (5) is locally Lipschitz continuous at (w, N). It follows that (5) admits a unique C^2 solution that is continuous in initial values that lie in a sufficiently small neighborhood of (w, N). Let $\mathcal{C}$ denote the solution to (5) starting at (w, N) and suppose towards a contradiction that $\mathcal{C}$ escapes $\operatorname{cl} \mathcal{B}_r(\mathcal{W})$ in a neighborhood U of w. Without loss of generality, we may assume that $\mathcal{C}$ escapes $\operatorname{cl} \mathcal{B}_r(\mathcal{W})$ to the right of w. Let π denote the projection of U onto $\partial \mathcal{B}_r(\mathcal{W})$ in direction N. Define

$$\mathcal{W}(v) := \begin{cases} \mathcal{W}_{\|\pi(v)-v\|} & \text{if } v \in \operatorname{int} \mathcal{B}_r(\mathcal{W}), \\ \mathcal{W} & \text{otherwise}, \end{cases}$$

and $\mathcal{L}(v) := \{\delta \in \mathbb{R}^{2\times m} \mid v + r\delta(y) \in \mathcal{W}(v) \text{ for all } y \in Y\}$. Then, $\mathcal{C}$ is also the unique solution to (17) starting in (w, N) for Lipschitz expansion $\mathcal{L}$. Since $(w, N) \notin \Gamma(r, \mathcal{W})$, the boundary $\partial \mathcal{B}_r(\mathcal{W})$ is C^1 at w by Lemma C.3. Therefore, a solution $\mathcal{C}'$ to (17) with initial conditions (w, N') for arbitrarily slight rotations N' of N cuts through $\mathcal{B}_r(\mathcal{W})$ as illustrated in Figure 18. Since $(w, N) \notin \Gamma(r, \mathcal{W}) \subseteq \Gamma(\mathcal{L})$ and $\Gamma(\mathcal{L})$ is closed, such a curve $\mathcal{C}'$ satisfies the conditions of Lemma C.2 for sufficiently slight rotations N'. Such a curve $\mathcal{C}'$ cannot exist by Lemma C.2, hence $\mathcal{C}$ cannot escape $\operatorname{cl} \mathcal{B}_r(\mathcal{W})$.

Suppose next that $\mathcal{C}$ falls into the interior of $\mathcal{B}_r(\mathcal{W})$ in a neighborhood of (w, N), that is, there exists $v \in \mathcal{C} \cap \operatorname{int} \mathcal{B}_r(\mathcal{W})$ arbitrarily close to w. By convexity of $\mathcal{B}_r(\mathcal{W})$, this is not possible if $\mathcal{C}$ is a trivial solution to (5), hence $\mathcal{C}$ is a solution with positive curvature. We may assume without loss of generality that this happens to the right of w as illustrated in the left panel of Figure 8. Let $\varepsilon > 0$ be sufficiently small such that $B_\varepsilon(v) \subseteq \operatorname{int} \mathcal{B}_r(\mathcal{W})$. By continuity in initial conditions, the solution $\mathcal{C}'$ to (5) with initial condition (w, N') for sufficiently slight counterclockwise rotations N' must contain $v_R \in B_\varepsilon(v)$. Moreover, $\mathcal{C}'$ enters $\mathcal{B}_r(\mathcal{W})$ to the left of w so that it contains a point $v_L \in \operatorname{int} \mathcal{B}_r(\mathcal{W})$. Finally, $\mathcal{C}'$ contains a point $w' \notin \operatorname{cl} \mathcal{B}_r(\mathcal{W})$ between w and v_R as illustrated in the left panel of Figure 8. Since $\mathcal{C}'$ satisfies the conditions of Corollary C.1, it follows that $w' \in \mathcal{B}_r(\mathcal{W})$, a contradiction.

For the second statement, suppose towards a contradiction that $\partial \mathcal{B}_r(\mathcal{W})$ contains a straight line segment L of positive length with coordinate outward normal N with $L \times \{N\} \cap \Gamma(r, \mathcal{W}) = \emptyset$. Let w be in the relative interior of L. Due to Lemma B.7, there exists a with $(w, N) \in \operatorname{int} E_a(\mathcal{L}_0)$ and $\kappa_a > 0$ in a neighborhood of (w, N). Let $\mathcal{C}_\eta$ be a strictly curved solution to (5) starting at $(w + \eta N, N)$. For η sufficiently small, $\mathcal{N}_{\mathcal{C}_\eta} \subseteq E_a(\mathcal{L}_0)$ and $\mathcal{C}_\eta$ contains end points in $\operatorname{int} B_r(\mathcal{W})$. Corollary C.1 thus implies that $w + \eta N \in \mathcal{B}_r(\mathcal{W})$, which is a contradiction. □

Proof of Lemma 6.6. Lemma C.3 implies that $\partial \mathcal{B}_r(\mathcal{W})$ is continuously differentiable outside of $\Gamma(r, \mathcal{W})$. In particular, $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) = \{N_w\}$ is a singleton. It follows from Lemma C.4 that $\partial \mathcal{B}_r(\mathcal{W})$ solves (5) at almost every $(w, N_w) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \setminus \Gamma(r, \mathcal{W})$. □

D Proof of Proposition 6.9

In this appendix, we characterize all boundary points of $\mathcal{B}_r(\mathcal{W})$ in $\Gamma(r,\mathcal{W})$, where Brownian information is not needed to provide incentives. We begin with the following auxiliary result, whose proof relies on an exact law of large numbers.

Lemma D.1. *Consider a payoff pair w, a set of action profiles $\mathcal{A}_w$, a set of weights $(\nu_a)_{a\in\mathcal{A}_w}$, and $(\delta_a)_{a\in\mathcal{A}_w}$ with $\delta_a \in \Psi_a^0(w,r,\mathcal{W})$ for each $a\in\mathcal{A}_w$. If w coincides with*

$$x := \sum_{a\in\mathcal{A}_w} \nu_a\big(g(a)+\delta_a\lambda(a)\big), \tag{20}$$

or w lies between x and some $v\in\mathcal{B}_r(\mathcal{W})$ on a straight line segment, then $w\in\mathcal{B}_r(\mathcal{W})$.

Proof. Suppose that such $w, v, \mathcal{A}_w, (\delta_a)_{a\in\mathcal{A}_w}$, and $(\nu_a)_{a\in\mathcal{A}_w}$ exist. Consider a solution (W,A,δ,β,M) to (2) on a stochastic framework $(\Omega,\mathcal{F},\mathbb{F},P,Z,(J^y)_{y\in Y})$ satsifying:

(i) $(\Omega,\mathcal{F},\mathbb{F},P)$ is rich enough to admit partitions $\Xi_t = (\Xi_t(a))_{a\in\mathcal{A}_w}$ of Ω with $P(\Xi_t(a)) = \nu_a$ for each $a\in\mathcal{A}_w$ and each $t>0$ such that $\Xi = (\Xi_t)_{t\geq 0}$ is independent of Z, $(J^y)_{y\in Y}$ and Ξ satisfies an exact law of large numbers; see Sun [26],

(ii) $W_0 = w$,

(iii) before the first jump time σ of any of the processes $(J^y)_{y\in Y}$, we have

$$A = \sum_{a\in\mathcal{A}_w} a 1_{\Xi(a)}, \qquad \delta = \sum_{a\in\mathcal{A}_w} \delta_a 1_{\Xi(a)},$$

$\beta\equiv 0$, and $\mathrm{d}M_t = (v-w)(\mathrm{d}J'_t - \lambda'\,\mathrm{d}t)$, where J' is a Poisson process independent of $(J^y)_{y\in Y}$ with intensity $\lambda' = \|x-w\|/\|v-w\|$.

It follows from an exact law of large numbers (see Proposition 2.5 in Sun [26]) that

$$\mathrm{d}W_t = \sum_{y\in Y} \delta_w(y)\,\mathrm{d}J_t^y + (v-w)\,\mathrm{d}J'_t.$$

Thus, W stays at w until either an event y or a jump in J' occurs. Moreover, on the set $[\![0,\sigma)\!)$, $(0,\delta)$ enforces A with $W + r\delta(y)\in\mathcal{W}$ for every $y\in Y$. Let τ denote the first jump time of J'. On the event $\{\tau<\sigma\}$, the process W jumps from w to v, that is, $W_\tau = v\in\mathcal{B}_r(\mathcal{W})$. Thus, Lemma A.2 applies and yields $w\in\mathcal{B}_r(\mathcal{W})$. □

Lemma D.2. *Let $(w,N)\in\mathcal{N}_{\mathcal{B}_r(\mathcal{W})}\cap\Gamma_a(\mathcal{W})$ with $w\notin\mathcal{S}_r(\mathcal{W})\cup\partial\mathcal{K}_{r,a}(\mathcal{W})$. Then $w\in\operatorname{int}\mathcal{K}_{r,a}(\mathcal{W})$ and*

$$N^\top\big(w-g(a)\big) = \max_{\delta\in\Psi_a^0(w,r,\mathcal{W})} N^\top\delta\lambda(a). \tag{21}$$

Proof. Fix an action profile a and a pair $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma_a(\mathcal{W})$. Suppose that w is neither a stationary payoff, nor on the boundary of $\mathcal{K}_{r,a}(\mathcal{W})$. Since $(w, N) \in \Gamma_a(\mathcal{W})$ implies that $w \in \mathcal{K}_{r,a}(\mathcal{W})$, it follows that $w \in \operatorname{int} \mathcal{K}_{r,a}(\mathcal{W})$. By definition of $\Gamma_a(\mathcal{W})$, there exists at least one $\delta \in \Psi_a^0(w, r, \mathcal{W})$, for which $N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$, hence

$$\max_{\delta \in \Psi_a^0(w,r,\mathcal{W})} N^\top \delta\lambda(a) \geq N^\top\big(w - g(a)\big).$$

For the opposite inequality, consider an arbitrary $\delta \in \Psi_a^0(w, r, \mathcal{W})$. Since w is in the interior of $\mathcal{K}_{r,a}(\mathcal{W})$, either Condition (i) or (iii) of Lemma 6.8 must hold. Lemma 6.8 thus implies that $N^\top\big(g(a) + \delta\lambda(a) - w\big) \leq 0$. Since δ was arbitrary, (21) follows. □

Lemma D.3. *For any $w_* \in \mathcal{K}_{r,a}(\mathcal{W})$ and $\delta_* \in \Psi_a^0(w_*, r, \mathcal{W})$, there exists a continuous selector $\delta : \mathcal{K}_{r,a}(\mathcal{W}) \to \Psi_a^0$ with $\delta(w) \in \Psi_a^0(w, r, \mathcal{W})$ and $\delta(w_*) = \delta_*$.*

Proof. Since Ψ_a^0 is closed and convex, the intersection with the closed- and convex-valued affine map $\mathcal{L}_0$ is lower semi-continuous. Therefore, $\Psi_a^0(w, r, \mathcal{W}) = \Psi_a^0 \cap \mathcal{L}_0(w)$ admits the a continuous selector by Michael's theorem; see Theorem 21.1 in Arutyunov and Obukhovskii [6]. The continuous extension property ($\delta(w_*) = \delta_*$) follows from Theorem 21.2 in [6]. □

Lemma D.4. *Let $\mathcal{A}_w$ denote a set of action profiles that strictly and minimally decomposes $w \notin \mathcal{S}_r(\mathcal{W})$. Then for each $a \in \mathcal{A}_w$, there exists $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ and $\delta \in \Psi_a^0(w, r, \mathcal{W})$ with $N^\top\big(g(a) + \delta\lambda(a) - w\big) > 0$.*

Proof. Fix $w \notin \mathcal{S}_r(\mathcal{W})$ and $\mathcal{A}_w$ that strictly and minimally decomposes w. Suppose towards a contradiction that there exists $a \in \mathcal{A}_w$ such that $N^\top\big(g(a)+\delta\lambda(a)-w\big) \leq 0$ for every $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ and every $\delta \in \Psi_a^0(w, r, \mathcal{W})$. This is equivalent to saying that the intersection of the closed upper half space $H(w, N) := \big\{x \mid N^\top(x - w) \geq 0\big\}$ with the convex set $\mathcal{X}_a := g(a) + \Psi_a^0(w, r, \mathcal{W})\lambda(a)$ has non-empty interior for every $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Since w is not stationary, $\mathcal{X}_a$ is strictly separated from w, hence $H(w, N) \cap \mathcal{X}_a = \emptyset$ for every $N \in \operatorname{int} \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Since $\mathcal{A}_w$ strictly decomposes w, for every $N \in \partial\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$, there exists some $\hat{a} \in \mathcal{A}_w$ and $\hat{\delta} \in \Psi_{\hat{a}}(w, r, \mathcal{W})$ with $N^\top\big(g(a) + \delta\lambda(a) - w\big) > 0$. In particular, $\hat{a} \neq a$ and hence a does not contribute anything to the decomposition of w. This contradicts minimality of $\mathcal{A}_w$. □

Lemma D.5. *Let $w \in \partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$ be minimally and strictly decomposed by $\mathcal{A}_w$. Then $w \in \mathcal{B}_r(\mathcal{W})$, $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ for each $a \in \mathcal{A}_w$ and (7) is satisfied.*

Proof. Fix such a payoff pair $w \in \partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$ and a set of action profiles $\mathcal{A}_w$. By Lemma D.4, for each $a \in \mathcal{A}_w$, there exist $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ and $\delta \in \Psi_a^0(w, r, \mathcal{W})$ with $N^\top\big(g(a) + \delta\lambda(a) - w\big) > 0$. Thus, Lemma 6.8 implies that $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ for each $a \in \mathcal{A}_w$. To show that $w \in \mathcal{B}_r(\mathcal{W})$, we distinguish two cases.

Consider first the case, where there exists $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ such that $N^\top\big(g(a) + \delta\lambda(a) - w\big) \leq 0$ for every $a \in \mathcal{A}_w$ and every $\delta \in \Psi_a^0(w, r, \mathcal{W})$. Since $\mathcal{A}_w$ strictly decomposes w, this implies that $N \in \operatorname{int} \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Because $g(a) + \Psi_a^0(w, r, \mathcal{W})\lambda(a)$ are closed and convex for every action profile a, there must exist $a_1, a_2 \in \mathcal{A}_w$ and $\delta_j \in \Psi_{a_j}(w, r, \mathcal{W})$ for $j = 1, 2$ with $w = \nu\big(g(a_1) + \delta_1\lambda(a_1)\big) + (1 - \nu)\big(g(a_2) + \delta_2\lambda(a_2)\big)$ for some $\nu \in (0, 1)$. Therefore, Lemma D.1 applies and shows that $w \in \mathcal{B}_r(\mathcal{W})$.

Consider next the case, where for each $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ there exists some $a \in \mathcal{A}_w$ and $\delta \in \Psi_a^0(w, r, \mathcal{W})$ with $N^\top\big(g(a) + \delta\lambda(a) - w\big) > 0$. Then there must exist some weights $(\nu_a)_{a \in \mathcal{A}_w}$ and incentives $(\delta_{a,w})_{a \in \mathcal{A}_w}$ with $\delta_{a,w} \in \Psi_a^0(w, r, \mathcal{W})$ such that $N^\top(x - w) > 0$ for every outward normal $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$, where $x := \sum_{a \in \mathcal{A}_w} \nu_a\big(g(a) + \delta_a\lambda(a)\big)$. Therefore, the straight line through x and w has a non-empty intersection with $\operatorname{int} \mathcal{B}_r(\mathcal{W})$. Lemma D.1 thus implies that $w \in \mathcal{B}_r(\mathcal{W})$.

For the last statement, note that $\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W}) := \bigcap_{a \in \mathcal{A}_w} \mathcal{K}_{r,a}(\mathcal{W})$ is closed and convex since $\mathcal{K}_{r,a}(\mathcal{W})$ is closed and convex for each $\mathcal{K}_{r,a}(\mathcal{W})$. For each $a \in \mathcal{A}_w$, Lemma D.3 implies the existence of a continuous selector $\delta_a : \mathcal{K}_{r,a}(\mathcal{W}) \to \Psi_a^0$ such that $\delta_a(v) \in \Psi_a^0(v, r, \mathcal{W})$ for any $v \in \mathcal{K}_{r,a}(\mathcal{W})$ and $\delta_a(w) = \delta_{a,w}$. Let $B_\varepsilon(w)$ denote the closed ball around w with radius $\varepsilon > 0$. Suppose towards a contradiction that $B_\varepsilon(w) \cap \big(\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W}) \setminus \mathcal{B}_r(\mathcal{W})\big) \neq \emptyset$ for some $\varepsilon > 0$. Then there exists $v \in \mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W}) \setminus \mathcal{B}_r(\mathcal{W})$ sufficiently close to w such that the straight line segment L through $\sum_{a \in \mathcal{A}_w} \nu_a\big(g(a) + \delta_a(v)\lambda(a)\big)$ and v also enters the interior of $\mathcal{B}_r(\mathcal{W})$. This is impossible by Lemma D.1. Therefore, $\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})$ is locally contained in $\mathcal{B}_r(\mathcal{W})$ and hence $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \subseteq \mathcal{N}_w\big(\bigcap_{a \in \mathcal{A}_w} \mathcal{K}_{r,a}(\mathcal{W})\big)$. □

Proof of Proposition 6.9. Fix $(w, N) \in \mathcal{N}_{\mathcal{B}_r(\mathcal{W})} \cap \Gamma(r, \mathcal{W})$ such that w is not a stationary payoff. Lemma C.3 implies that $(w, N') \in \Gamma(r, \mathcal{W})$ for every $N' \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. In particular, the payoff pair w is decomposable. Suppose first that w is strictly decomposable. Let $\mathcal{A}_w$ denote a set of action profiles that strictly and minimally decomposes w. Lemma D.5 implies that $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ for each $a \in \mathcal{A}_w$ and that (7) holds. In particular, case (i) applies.

Suppose, therefore, that w is not strictly decomposable, i.e., there exists an extremal normal vector N, for which $N^\top\big(g(a) + \delta\lambda(a) - w\big) \leq 0$ for all $\delta \in \Psi_a^0(w, r, \mathcal{W})$ and every $a \in \mathcal{A}$. Because $(w, N) \in \Gamma(r, \mathcal{W})$, there must exist at least one action profile, for which the expression is non-negative and hence (6) holds. If $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) = \{N\}$, then the maximizer a_* of (6) minimally decomposes w. If w is a corner, then, because w is decomposable, there exists $a \in \mathcal{A}$ and $\delta \in \Psi_a^0(w, r, \mathcal{W})$ with $N'^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$ for arbitrarily slight rotations $N' \in \operatorname{int} \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ of the extremal normal vector N. Without loss of generality, suppose that N' are clockwise rotations of N. Since $\Psi_a^0(w, r, \mathcal{W})$ is closed, there must exist $\delta_* \in \Psi_a^0(w, r, \mathcal{W})$ with $N'^\top\big(g(a) + \delta_*\lambda(a) - w\big) \geq 0$ for $N' = N$ as well as for slight clockwise rotations N' of N. Therefore, a maximizes (6) and $g(a) + \delta_*\lambda(a)$ lies on the tangent to $\mathcal{B}_r(\mathcal{W})$ to the right of w as illustrated in Figure 19. It follows that a decomposes w.

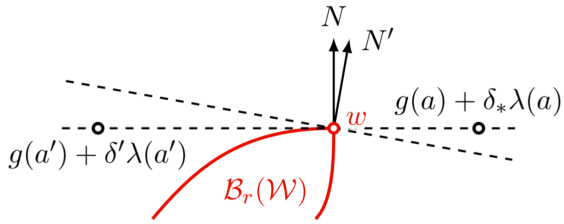

Figure 19: Since $N'^{\top}\big(g(a) + \delta_* \lambda(a) - w\big) \geq 0$ for slight clockwise rotations of N, $g(a) + \delta_* \lambda(a)$ must lie to the right of w and hence decompose w.

Clearly, case (ii.b) holds if $w \in \partial \mathcal{K}_{r,a_*}(\mathcal{W})$. Suppose, therefore, that w is in the interior of $\mathcal{K}_{r,a_*}(\mathcal{W})$. If the set of outward normals is not a singleton, then $(w, N_j) \in \Gamma_{a_*}(\mathcal{W})$ for three distinct N_1, N_2, N_3. Lemma D.2 thus implies that

$$N_i^{\top} w = \max_{\delta \in \Psi_{a_*}(w,r,\mathcal{W})} N_i^{\top}\big(g(a_*) + \delta \lambda(a_*)\big)$$

for $i = 1, 2, 3$, which is possible only if $w = g(a_*) + \delta \lambda(a_*)$ for some $\delta \in \Psi_{a_*}(w, r, \mathcal{W})$. This implies that w is stationary, which is a contradiction. We have thus shown that $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) = \{N\}$, hence case (ii.a) holds. Because the cases we have considered are mutually exclusive, this concludes the proof. □

E Proof of Theorem 6.10

The proof of Theorem 6.10 will proceed in three steps. First, we show that stationary payoffs are contained $\mathcal{B}_r(\mathcal{W})$, that is, we provide the formal proof of Lemma 6.2. Second, we show that $\operatorname{cl} \mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$ for any convex set $\mathcal{X} \subseteq \mathcal{V}^*$, whose boundary satisfies (i) and (ii) of Theorem 6.10. Finally, we show that $\partial \mathcal{B}_r(\mathcal{W})$ satisfies (i) and (ii) of Theorem 6.10 so that $\operatorname{cl} \mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$ by the previous claim, showing that $\mathcal{B}_r(\mathcal{W})$ is closed. We begin with the proof of Lemma 6.2.

Proof of Lemma 6.2. Because the union of two $\mathcal{W}$-relaxed self-generating sets is again $\mathcal{W}$-relaxed self-generating, $\mathcal{S}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$ will follow once we show that $\{w\}$ is $\mathcal{W}$-relaxed self-generating for an arbitrary stationary payoff pair w. By definition, there exist a, δ_0 such that $(0, \delta_0)$ enforces a, $w = g(a) + \delta_0 \lambda(a)$, and $w + r\delta_0(y) \in \mathcal{W}$ for every $y \in Y$. The constant strategy profile $A \equiv a$ is thus enforced by the continuation promise (β, δ) with $\beta \equiv 0$ and $\delta \equiv \delta_0$. A solution W to (2) starting in w with $M \equiv 0$ and A, β, δ as given thus has neither drift nor diffusion term. It follows that W remains in w until the arrival time σ_1 of the first event y, at which point W jumps to the payoff set $\mathcal{W}$ since $W_\sigma = W_{\sigma-} + r\delta_{\sigma-}(y) = w + r\delta_0(y) \in \mathcal{W}$. □

Proposition E.1. *Fix a closed, convex set $\mathcal{W}$ with $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$. Let $\mathcal{X} \subseteq \mathcal{V}^*$ be a convex set with $\partial \mathcal{X}$ satisfying (i) and (ii) of Theorem 6.10. Then $\operatorname{cl} \mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$.*

Lemma E.2. *Fix a closed, convex set $\mathcal{W}$ with $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$ and let $\mathcal{X}$ be a closed, convex subset of $\mathcal{V}^*$. Suppose that there exists a set $\mathcal{X}' \subseteq \partial\mathcal{X}$ such that:*

(i) For each $w \in \mathcal{X}'$, there exist $\varepsilon_w > 0$ and a $\mathcal{W}$-enforceable solution W^w to (2) such that W^w remains in $\mathcal{X}$ on $[\![0, \tau_w \wedge \sigma_1)\!)$ and $\min_{w \in \mathcal{X}'} \tau_w > 0$ a.s., where $\tau_w := \inf\{t \geq 0 \mid W_t^w \notin B_{\varepsilon_w}(w)\}$.

(ii) For every $w \notin \mathcal{X}'$, there exists a $\mathcal{W}$-enforceable solution W to (2) that remains on $\partial\mathcal{X}$ until time σ_1 or until W reaches $\mathcal{X}'$.

Then $\mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$.

Proof. Fix an arbitrary payoff pair $w \in \mathcal{X}$. We must show that w can be attained by a $\mathcal{W}$-enforceable solution to (2) that remains in $\mathcal{X}$ until time σ_1. Using public randomization at time 0, we can attain w with a continuation value W_0 on $\partial\mathcal{X}$. By condition (ii), we can attain W_0 with a $\mathcal{W}$-enforceable solution W to (2) that remains on $\partial\mathcal{X} \subseteq \mathcal{X}$ up to time $\tau_1 := \inf\{t \geq 0 \mid W_t \in \mathcal{X}'\}$. Due to condition (i) and Lemma A.2, on the event $\{\tau_1 < \sigma_1\}$ we can extend W to an enforceable solution W that stays in $\mathcal{X}$ up until time $\tau_2 := \tau_1 + \hat{\tau}$, where $\hat{\tau} = \min_{w \in \mathcal{X}'} \tau_w > 0$. On the event $\{\tau_2 < \sigma_1\}$, we can use public randomization to attain W_{τ_2} with continuation values on the boundary. An iteration of this procedure yields a $\mathcal{W}$-enforceable solution W that remains in $\mathcal{X}$ until time σ_1 since $\tau_{2n} \geq \sum_{k=1}^{n}(\tau_{2k} - \tau_{2k-1}) \to \infty$ a.s. as the sum of independent and identically distributed random variables. □

It remains to verify the conditions of Lemma E.2 for each boundary point of a set $\mathcal{X}$ as given in Proposition E.1. We begin with boundary points on C^1-solutions to (6) via the following lemma. It is the analogue to Corollary C.1 inside $\Gamma(r, \mathcal{W})$.

Lemma E.3. *Let $\mathcal{C}$ be a C^1-solution to (6). For every $w \in \mathcal{C}$, there exists a solution W to (2) with $W_0 = w$, $\beta = 0$, $M = 0$, and $\delta \in \Psi_a^0(W_-, r, \mathcal{W})$ such that W remains on $\mathcal{C}$ until an endpoint of $\mathcal{C}$ is reached or an infrequent event occurs.*

Proof. Fix such a curve $\mathcal{C}$ and let $a^*(w)$ and $\delta^*(w)$ denote the maximizers in (6) at $w \in \mathcal{C}$. In particular, $\delta^*(w) \in \Psi_{a^*(w)}(w, r, \mathcal{W})$ for any $w \in \mathcal{C}$. Consider first a subsegment $\mathcal{C}_{a_1}$ of $\mathcal{C}$, on which $a^* = a_1$. We refer to $w \in \mathcal{C}_{a_1}$ as the starting point or end point of $\mathcal{C}_{a_1}$, respectively, if $\mathcal{C}_{a_1} \cap H(w, T_+) = \{w\}$ and $\mathcal{C}_{a_1} \cap H(w, T_-) = \{w\}$, where $T_\pm := \pm\big(g(a_1) + \delta^*(w)\lambda(a_1)\big)$ and $H(w,T) := \big\{x \bigm| T^\top(x - w) \geq 0\big\}$ is the closed upper half-space in direction T. Fix $w \in \mathcal{C}_{a_1}$ that is not the end point of $\mathcal{C}_{a_1}$. It follows as in the proof of Lemma 6.4 that a solution to (2) starting at $W_0 = w$ with $A = a_1$, $\beta = 0$, $\delta = \delta^*(W_-)$, and $M \equiv 0$ remains on $\mathcal{C}_{a_1}$ until either an event y occurs or the end point w_0 of $\mathcal{C}_{a_1}$ is reached. If w_0 is an end point of $\mathcal{C}$, we have achieved our goal. Suppose, therefore, that w_0 is in the relative interior of $\mathcal{C}$.

We distinguish two cases. If there exist $a_2 \neq a_1$ such that w_0 is the end point of a segment $\mathcal{C}_{a_2}$ with $a^* = a_2$, then there must exist $\delta_{a_k} \in \Psi_{a_k}(w, r, \mathcal{W})$ such that $g(a_k) + \delta_{a_k}\lambda(a_k)$ for $k = 1, 2$ are on the tangent to $\mathcal{C}$ at w on opposite sides of w. Set

$$\nu_{a_k} := \frac{\|w_0 - g(a_k) - \delta_{a_k}\lambda(a_k)\|}{\|g(a_1) + \delta_{a_1}\lambda(a_1) - g(a_2) - \delta_{a_2}\lambda(a_2)\|}$$

Then x defined in (20) coincides with w_0. It follows as in the proof of Lemma D.5 that w_0 can be attained by an enforceable solution $\hat{W}$ to (2) that remains in w_0 until time σ_1 with $\hat{W}_{\sigma_1} \in \mathcal{W}$. A concatenation of W with $\hat{W}$ thus achieves our goal.

If there exists no such a_2, then w_0 must be the starting point of a segment $\mathcal{C}_{a_3} \subseteq \mathcal{C}$, on which $a^* = a_3$. Thus, w_0 can be attained by a solution to (2) until an infrequent event occurs or the end point of $\mathcal{C}_{a_3}$ is reached. An iteration of this procedure yields a solution that remains on $\mathcal{C}$ until time σ_1 or an end point of $\mathcal{C}$ is reached. □

Proof of Proposition E.1. Fix a closed, convex set $\mathcal{W}$ with $\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{W}$ and a set $\mathcal{X}$ whose boundary satisfies (i) and (ii) of Theorem 6.10. Let $\mathcal{X}_1 \subseteq \partial\mathcal{X}$ denote all boundary points of $\mathcal{X}$, where $\partial\mathcal{X}$ is locally given by a solution to (5) or (6). Let $\mathcal{X}_2 \subseteq \partial\mathcal{X}$ denote all payoff pairs that are strictly decomposable and denote by $\mathcal{X}'$ the set of all boundary payoffs that are not in $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{S}_r(\mathcal{W})$. We first show that $\mathcal{X}'$ is finite and that each $w \in \mathcal{X}'$ satisfies Condition (i) of Lemma E.2.

Since $\partial\mathcal{X}$ satisfies conditions (i) and (ii) of Theorem 6.10, it follows that for any $w \in \mathcal{X}'$, either (a) w is in $\partial\mathcal{K}_{r,a}(\mathcal{W})$ for some action profile a which maximizes (6) and decomposes w or (b) w is not in $\partial\mathcal{K}_{r,a}(\mathcal{W})$ for any a and the boundary smoothly transitions from a solution to (5) to a solution to (6) at w. In case (a), w is the starting point of a solution $\mathcal{C}_w$ to (6) of positive length that remains in $\operatorname{cl}\mathcal{X}$. Thus, Lemma E.2 implies that w can be attained by a $\mathcal{W}$-enforceable solution W^w to (2) that remains on $\mathcal{C}_w$ and hence in $\operatorname{cl}\mathcal{X}$ until time σ_1 or until time τ_w when the endpoint of $\mathcal{C}_w$ is reached. Consider now case (b). Let $\mathcal{C}_w$ be a C^1 solution to (6) with initial condition (w, N_w). If $\mathcal{C}_w$ stays within $\operatorname{cl}\mathcal{X}$ in a neighborhood of w, then the same argument as in case (a) applies. Suppose, therefore, that $\mathcal{C}_w$ escapes $\operatorname{cl}\mathcal{X}$ in a neighborhood U of w. Without loss of generality, suppose that this happens to the right of w, that is, $\partial\mathcal{X}$ is a strictly curved solution to (5) to the right of w. Let a_* denote the maximizer in (6) as $v \in \mathcal{C}_w$ approaches w from the right. Since $\mathcal{K}_{r,a_*}(\mathcal{W})$ is convex, it follows that $w' \in \mathcal{K}_{r,a_*}(\mathcal{W})$ for any $w' \in \partial\mathcal{X}$ and any $w' \in \mathcal{C}_w$ sufficiently close to w. Consider $w' \in \partial\mathcal{X}$ and $v' \in \mathcal{C}_w$ with identical outward normal vector N sufficiently close to w such that they both lie in $\mathcal{K}_{r,a_*}(\mathcal{W})$. This implies that $\Psi^0_{a_*}(v', r, \mathcal{W})$ and $\Psi^0_{a_*}(w', r, \mathcal{W})$ are both non-empty. Since $\mathcal{C}_w$ escapes $\operatorname{cl}\mathcal{X}$ and a_* maximizes (6) at (v', N), it follows that

$$N^\top w' < N^\top v' = \max_{x \in g(a_*) + \Psi^0_{a_*}(v', r, \mathcal{W})\lambda(a_*)} N^\top x \leq \max_{x \in g(a_*) + \Psi^0_{a_*}(w', r, \mathcal{W})\lambda(a_*)} N^\top x.$$

where we have used that $\Psi^0_{a_*}(w', r, \mathcal{W}) = \Psi^0_{a_*} \cap \mathcal{L}_0(w') = \Psi^0_{a_*} \cap (\mathcal{L}_0(v') + (v' - w')/r\mathbf{1})$ is non-empty in the last inequality. It follows that there exists $\delta \in \Psi^0_{a_*}(w', r, \mathcal{W})$

with strictly inward-pointing drift at w', implying that $(w', N) \in \Gamma(r, \mathcal{W})$. This is a contradiction to the fact that $\partial\mathcal{X}$ is a C^1 solution to (5) to the right of w. Therefore, $\mathcal{C}_w$ remains within $\operatorname{cl}\mathcal{X}$, hence w can be attained locally as in case (a).

Choose now $\varepsilon > 0$ such that for every $w \in \mathcal{X}'$, the curve $\mathcal{C}_w$ remains within $\operatorname{cl}\mathcal{X}$ for distance ε in the direction of the drift rate of W^w attaining w. Since the drift rate of W^w is uniformly bounded because $\mathcal{A}$ and $\mathcal{V}^*$ are bounded, the time τ it takes for W^w to reach the end point $\mathcal{C}_w$ is uniformly bounded from below. In particular, $\min_{w\in\mathcal{X}'} \tau_w > 0$ a.s., hence $\mathcal{X}'$ satisfies Condition (i) of Proposition E.1.

Observe that Condition (ii) of Proposition E.1 is trivially satisfied for stationary boundary points. Once we show that for any $j = 1, 2$, any $w \in \mathcal{X}_j$ can be attained by a $\mathcal{W}$-enforceable solution that remains in $\mathcal{X}_j$ until time σ_1 or until $\partial\mathcal{X} \setminus \mathcal{X}_j$ is reached, then condition (ii) of Proposition is satisfied for the entire boundary $\partial\mathcal{X} \setminus \mathcal{X}'$. Indeed, by iteratively concatenating the solutions on $\mathcal{X}_1$, $\mathcal{X}_2$, and $\mathcal{S}_r(\mathcal{W})$, such a solution either remains in $\partial\mathcal{X} \setminus \mathcal{X}'$ until time σ_1 or until $\mathcal{X}'$ is reached.

For $\mathcal{X}_1$, we have already shown the desired statement in Lemmas 6.4 and E.3, respectively. Consider, therefore, a payoff pair $w \in \mathcal{X}_2$ that is minimally and strictly decomposable by some set $\mathcal{A}_w$ with $\mathcal{N}_w(\mathcal{X}) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,\mathcal{A}_w}(\mathcal{W})\big)$. As in the proof of Lemma D.5, either $w \in \mathcal{B}_r(\mathcal{W})$ or there exist weights $(\nu_a)_{a\in\mathcal{A}_w}$ and continuation promises $(\delta_a)_{a\in\mathcal{A}_w}$ such that $\delta_a \in \Psi^0_a(w, r, \mathcal{W})$ for every $a \in \mathcal{A}_w$ and $N^\top(x_w - w) > 0$ for every outward normal $N \in \mathcal{N}_w(\mathcal{X})$, where $x_w := \sum_{a\in\mathcal{A}_w} \nu_a\big(g(a) + \delta_a\lambda(a)\big)$. For each $a \in \mathcal{A}_w$, let $\Xi(a)$ be defined as in the proof of Lemma D.1 and let $v_w \in \operatorname{int}\mathcal{X}$ be collinear with w and x_w such that $c\|v_w - w\| \geq \|x_w - w\|$ for a constant $c > 0$. Such a constant c exists because x_w is uniformly bounded in $w \in \mathcal{V}^*$. Let $(W, A, \beta, \delta, M, (J^y)_{y\in Y}, Z)$ be the solution to (2) with $W_0 = w$,

$$A = \sum_{a\in\mathcal{A}_w} a1_{\Xi(a)}, \qquad \delta = \sum_{a\in\mathcal{A}_w} \delta_a 1_{\Xi(a)},$$

$\beta \equiv 0$, and $\mathrm{d}M_t = (V_w - w)\,\mathrm{d}J'_t - \lambda'(v_w - w)\,\mathrm{d}t$, where J' is a Poisson process independent of $(J^y)_{y\in Y}$ with intensity $\lambda_w = \|x_w - w\|/\|v_w - w\| \leq c$ and V_w is a $\partial\mathcal{X}$-valued random variable with expected value v_w that is independent from J', $(J^y)_{y\in Y}$, and Z. It follows from an exact law of large numbers that

$$\mathrm{d}W_t = \sum_{y\in Y} \delta_w(y)\,\mathrm{d}J^y_t + (V_w - w)\,\mathrm{d}J'_t.$$

Therefore, W is a $\mathcal{W}$-enforceable solution that remains in w until time σ_1 or a jump in J' occurs. If $V_w \in \mathcal{X}_2$, we can repeat this procedure and concatenate the solutions. Since $\lambda_w \leq c$ for every $w \in \mathcal{X}_2$ an iteration will extend beyond $\tau := \sigma_1 \wedge \inf\{t \geq 0 \mid W_t \in \mathcal{X}' \cup \mathcal{S}_r(\mathcal{W}) \cup \mathcal{X}_1\}$ with certainty. Therefore, condition (ii) of Lemma E.2 is satisfied on $\partial\mathcal{X} \setminus \mathcal{X}'$, hence $\operatorname{cl}\mathcal{X} \subseteq \mathcal{B}_r(\mathcal{W})$ by Lemma E.2. □

Proof of Theorem 6.10. It follows from Lemma 6.2 that all stationary payoff pairs are contained in $\mathcal{B}_r(\mathcal{W})$, i.e., $\mathcal{S}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$. We proceed to show that $\partial\mathcal{B}_r(\mathcal{W})$ satisfies

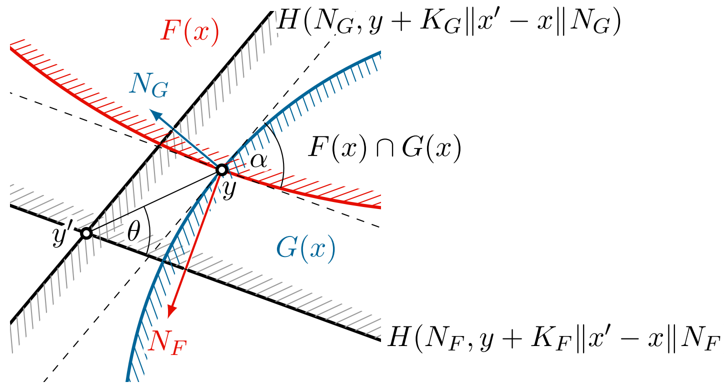

Figure 20: Since $\max(\theta, \alpha - \theta) \geq \alpha/2$, it follows that $\sin(\alpha/2)\|y' - y\| \leq \max(K_F, K_G)\|x' - x\|$

(i) and (ii). Lemma C.3 implies that the boundary $\partial\mathcal{B}_r(\mathcal{W})$ is continuously differentiable outside of $\Gamma(r, \mathcal{W})$. It then follows from statement (ii.a) of Proposition 6.9 that $\partial\mathcal{B}_r(\mathcal{W})$ is continuously differentiable outside of $\mathcal{S}_r(\mathcal{W}) \cup \mathcal{K}_r(\mathcal{W})$. Lemma C.4 and statement (ii.a) of Proposition 6.9 show that $\partial\mathcal{B}_r(\mathcal{W})$ is a solution to (5) and (6) within and outside of $\Gamma(r, \mathcal{W})$, respectively. This concludes the proof of the first statement. The second statement follows from statements (i) and (ii.b) of Proposition 6.9. It follows from Proposition E.1 that $\operatorname{cl}\mathcal{B}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$, hence $\mathcal{B}_r(\mathcal{W})$ is closed. Finally, Proposition E.1 also implies that $\mathcal{B}_r(\mathcal{W})$ is the largest of all such sets, thereby concluding the proof of Theorem 6.10. □

F Auxiliary results related to the optimality equation

Lemma F.1. *Let F and G be closed- and convex-valued maps that are locally Lipschitz continuous at $x_0 \in \operatorname{dom} F \cap G$. If $F(x_0) \cap \operatorname{int} G(x_0) \neq \emptyset$, then $F \cap G$ is locally Lipschitz continuous at x_0.*

Proof. By local Lipschitz continuity of the individual maps, there exists $\varepsilon > 0$ such that $F(x) \cap \operatorname{int} G(x) \neq \emptyset$ for any x in the closed ball $B_\varepsilon(x_0)$ and that F and G are Lipschitz on $B_\varepsilon(x_0)$. For any $x \in B_\varepsilon(x_0)$ and any $y \in \partial F(x) \cap \partial G(x)$, define

$$f(x, y) := \min_{N_F \in \mathcal{N}_y(F(x))} \min_{N_G \in \mathcal{N}_y(G(x))} {N_F}^\top N_G.$$

Observe that $f(x, y) > -1$ since the hyperplane with normal vector N_F through y would separate $F(x)$ and $G(x)$ otherwise, contradicting $F(x) \cap \operatorname{int} G(x) \neq \emptyset$. Since $y \mapsto \mathcal{N}_y(F(x))$ and $y \mapsto \mathcal{N}_y(G(x))$ are upper hemi-continuous, the infimum of $f(x, y)$ over $y \in \partial F(x) \cap \partial G(x)$ is attained, hence strictly larger than -1. Since $x \mapsto \partial F(x) \cap \partial G(x)$ is upper hemi-continuous as the intersection of two upper hemi-continuous maps, we deduce that

$$K := \inf_{x \in B_\varepsilon(x_0)} \inf_{y \in \partial F(x) \cap \partial G(x)} f(x, y) > -1.$$

Fix now any $x, x' \in B_\varepsilon(x_0)$ and any $y \in \partial F(x) \cap \partial G(x)$. Let K_F and K_G be the Lipschitz constants of the individual maps F and G, respectively, on $B_\varepsilon(x_0)$. Since F and G are convex-valued, it follows that $F(x') \subseteq H(N_F, y + K_F\|x' - x\|)$ and $G(x') \subseteq H(N_G, y + K_G\|x' - x\|)$ for any normal vectors $N_F \in \mathcal{N}_y(F(x))$ and $N_G \in \mathcal{N}_y(G(x))$, where $H(N, y) := \{y' \mid N^\top(y' - y) \leq 0\}$ denotes the lower half-space in direction N through y. Let y' denote the point in $\partial H(N_F, y + K_F\|x' - x\|) \cap \partial H(N_G, y + K_G\|x' - x\|)$ in the plane spanned by y, N_F, and N_G and let α denote the angle between the two half-spaces at y'; see Figure 20 for an illustration. Clearly

$$\|y' - y\| \leq \frac{\max(K_F, K_G)}{\sin(\alpha/2)}\|x' - x\| = \frac{\sqrt{2}\max(K_F, K_G)}{\sqrt{1 + {N_F}^\top N_G}}\|x' - x\|.$$

It follows that

$$\begin{aligned} F(x') \cap G(x') &\subseteq H(N_F, y + K_F\|x' - x\|N_F) \cap H(N_G, y + K_G\|x' - x\|N_G) \\ &\subseteq H(N_F, y) \cap H(N_G, y) + \frac{\sqrt{2}\max(K_F, K_G)}{\sqrt{1 + K}} B_{\|x'-x\|}(0) \end{aligned} \tag{22}$$

For $y \in \partial\big(F(x) \cap G(x)\big) \setminus \big(\partial F(x) \cap \partial G(x)\big)$ and $N \in \mathcal{N}_y\big(F(x) \cap G(x)\big)$, we have

$$F(x') \cap G(x') \subseteq H(N, y) + \frac{\sqrt{2}\max(K_F, K_G)}{\sqrt{1 + K}} B_{\|x'-x\|}(0). \tag{23}$$

Using the fact that the closed and convex set $F(x) \cap G(x)$ is the intersection over all bounding half-spaces, taking the intersection of (22) over all $y \in \partial F(x) \cap \partial G(x)$, $N_F \in \mathcal{N}_y(F(x))$, and $N_G \in \mathcal{N}_y(G(x))$, and of (23) over all $y \in \partial\big(F(x) \cap G(x)\big) \setminus \big(\partial F(x) \cap \partial G(x)\big)$ and any $N \in \mathcal{N}_y\big(F(x) \cap G(x)\big)$ yields

$$F(x') \cap G(x') \subseteq F(x) \cap G(x) + \frac{\sqrt{2}\max(K_F, K_G)}{\sqrt{1 + K}} B_{\|x'-x\|}(0).$$

This concludes the proof that $F \cap G$ is Lipschitz continuous on $B_\varepsilon(x_0)$. □

Lemma F.2. *Let $\mathcal{C} \subseteq \mathcal{W}$ be a C^1-solution to* (17) *for Lipschitz expansion $\mathcal{L}$ oriented by $w \mapsto N_w$ such that $\mathcal{N}_\mathcal{C} \cap \big(\Gamma(\mathcal{L}) \cup \mathcal{P}\big) = \emptyset$. For every $a \in \mathcal{A}$, there exists $K_a > 0$ such that for any $\gamma \geq 0$, any $(w, N_w) \in \mathcal{N}_\mathcal{C} \cap E_a(\mathcal{L})$, and any $(T_w\phi + N_w\chi, \delta)$ enforcing a with $\delta \in \mathcal{L}(w)$ and ${N_w}^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$, we have*

$$1 - \frac{\big(\|\phi\| - \gamma\|\chi\|\big)^2}{\big\|\phi(a, N_w, \delta)\big\|^2} \leq \frac{2K_a + 2\gamma}{\bar{\Psi}_a}\|\chi\|, \tag{24}$$

where $\bar{\Psi}_a := \inf_{w\in\mathcal{C}} \min_{\delta'\in\Psi_a(w,N_w,\mathcal{L})} \|\phi(a, N_w, \delta')\|^2 > 0$.

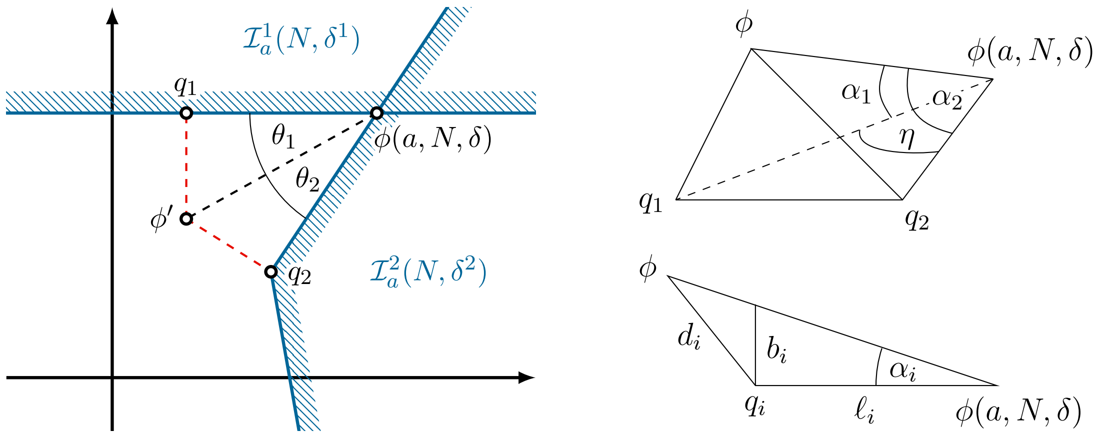

Figure 21: The left panel illustrates the positions of $\phi(a, N, \delta)$, ϕ', q_1, and q_2 relative to $\mathcal{I}_a^1(N, \delta^1)$ and $\mathcal{I}_a^2(N, \delta^2)$. The top right panel illustrates that $\alpha_j \geq \theta_j \geq \eta/2$. The bottom right panel illustrates the identity $\|\phi(a, N, \delta) - \phi\| \leq d_i(1 + 1/\tan(\alpha_i))$: because q_i is the projection of ϕ onto $\mathcal{I}_a^i(N, \delta^i)$, the angle at q_i is at least 90°. Therefore, $d_i \geq b_i = \ell_i \tan(\alpha_i)$ and the desired identity follows from the triangle inequality.

Proof. Fix an action profile $a \in \mathcal{A}$. Since $\Gamma(\mathcal{L})$ is closed by Lemma B.5, $\mathcal{N}_{\mathcal{C}}$ is bounded away from $\Gamma(\mathcal{L})$ and hence $\bar{\Psi}_a$ is indeed strictly positive. Observe that it is sufficient to show the existence of a constant $K_a > 0$ such that

$$\big\|\phi(a, N_w, \delta)\big\| - \|\phi\| \leq K_a \|\chi\| \tag{25}$$

holds for any $(w, N_w) \in \mathcal{N}_{\mathcal{C}} \cap E_a(\mathcal{L})$ and any $(T_w\phi + N_w\chi, \delta)$ enforcing a with $\delta \in \mathcal{L}(w)$ and $N_w{}^{\top}\big(g(a) + \delta\lambda(a) - w\big) \geq 0$. Indeed, then

$$1 - \frac{\|\phi\| - \gamma\|\chi\|}{\big\|\phi(a, N_w, \delta)\big\|} \leq \frac{\big\|\phi(a, N_w, \delta)\big\| - \|\phi\| + \gamma\|\chi\|}{\bar{\Psi}_a} \leq \frac{K_{a,2}(K_a + 1) + \gamma}{\bar{\Psi}_a}\|\chi\| \tag{26}$$

and hence (24) follows from (26) in conjunction with $1 - x \geq \frac{1}{2}(1 - x^2)$.

It remains to show (25). Clearly, the inequality is satisfied if $\|\phi\| \geq \|\phi(a, N, \delta)\|$, hence suppose that $\|\phi\| < \|\phi(a, N, \delta)\|$. For $i = 1, 2$, define

$$\mathcal{I}_a^i(N, \delta^i) := \big\{\phi \in \mathbb{R}^d \;\big|\; (T^i\phi, \delta^i) \text{ satisfies (3) for player } i\big\}$$

for any direction N and any $\delta \in \Psi_a$. Because $\mathcal{I}_a^i(N, \delta^i)$ is the intersection of closed half-spaces, it is a closed convex polyhedron. Observe that the normal vectors to its hyperfaces are given by the column vectors of $M^i(a)$ and (N, δ) determines only the location of those hyperfaces. Observe further that $\phi(a, N, \delta)$ is the shortest vector in the intersection $\Phi_a(N, \delta) := \mathcal{I}_a^1(N, \delta^1) \cap \mathcal{I}_a^2(N, \delta^2)$.

For $i = 1, 2$, let q_i denote the point in $\mathcal{I}_a^i(N, \delta^i)$ closest to ϕ and let $d_i = \|\phi - q_i\|$. Note that $\phi + N^i/T^i\chi \in \mathcal{I}_a^i(\delta^i)$ for $i = 1, 2$ and hence $|N^i/T^i|\|\chi\| \geq d_i$. Let ϕ' be the projection of ϕ onto the plane through $\phi(a, N, \delta)$, q_1, and q_2. Let $j \in \{1, 2\}$ be the index i for which the angle θ_i between $\phi(a, N, \delta) - \phi'$ and $\phi(a, N, \delta) - q_i$ is maximal. Then $\theta_j \geq \eta/2$, where η is the angle between $\phi(a, N, \delta) - q_1$ and $\phi(a, N, \delta) - q_2$. Let α_i denote the angle between $\phi(a, N, \delta) - \phi$ and $\phi(a, N, \delta) - q_i$; See Figure 21 for an illustration. Observe that $\alpha_i \geq \theta_i$ and hence

$$\|\phi(a, N, \delta) - \phi\| = d_j \left(1 + \frac{1}{\tan(\alpha_j)}\right) \leq \left(1 + \frac{1}{\tan(\eta/2)}\right) \left|\frac{N^j}{T^j}\right| \|\chi\|.$$

Note that η cannot be 0 since then $\phi(a, N, \delta)$ would not be the shortest vector in $\Phi_a(N, \delta)$. Since changes in N and δ do not change the direction of the hyperplanes bounding $\mathcal{I}_a^i(N, \delta)$, a uniform lower bound $\underline{\eta}$ for η is given by taking the minimum over all strictly positive angles between the finitely many hyperfaces of $\mathcal{I}_a^1(N, \delta)$ and $\mathcal{I}_a^2(N, \delta)$. Therefore, $\|\phi(a, N, \delta) - \phi\| \leq K_a \|\chi\|$ holds for

$$K_a = \left(1 + \frac{1}{\tan(\underline{\eta}/2)}\right) \sup_{w \in \mathcal{C}} \left(\left|\frac{N_w^1}{T_w^1}\right| + \left|\frac{N_w^2}{T_w^2}\right| \right).$$

Equation (25) now follows from the reverse triangle inequality. □