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 continuously, governed by a Brownian motion, 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 whose solutions approximate the boundary. Equilibrium strategies that attain payoff pairs on the efficient frontier are elicitable from the limiting solution. The analysis reveals the drastic influence of abrupt information on the equilibrium payoff set: because the presence of abrupt information enables equilibrium incentives through value burning, the equilibrium payoff set may contain corners and straight line segments outside the set of static Nash payoffs—two features that are precluded in games with a continuous stream of information.

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

1 Introduction

In continuous-time games with imperfect monitoring, information may arrive both continuously through a noisy signal and intermittently as occurrences of infrequent but informative events. In many applications, there is a clear distinction between the two types of information. Consider a climate agreement that obligates each signatory to reduce its greenhouse gas emissions. Countries cannot measure each other's emissions directly, hence they must rely on imperfect information to enforce the agreement. When one country violates the terms of the agreement and emits more than its agreed-upon share of greenhouse gases, other countries may observe

*I would like to thank Christoph Frei, Tilman Klumpp, participants of the 11^{th} World Congress of the Econometric Society, the 5^{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 Bern for valuable comments and interesting discussions. Financial support by the Natural Sciences and Engineering Research Council of Canada through grant RGPIN/402585-2011 as well as the Swiss National Science Foundation through grant P2SKP1_171737 is gratefully acknowledged.

†Department of Economics, UCLA, Los Angeles, CA 90095; benbernard@ucla.edu.

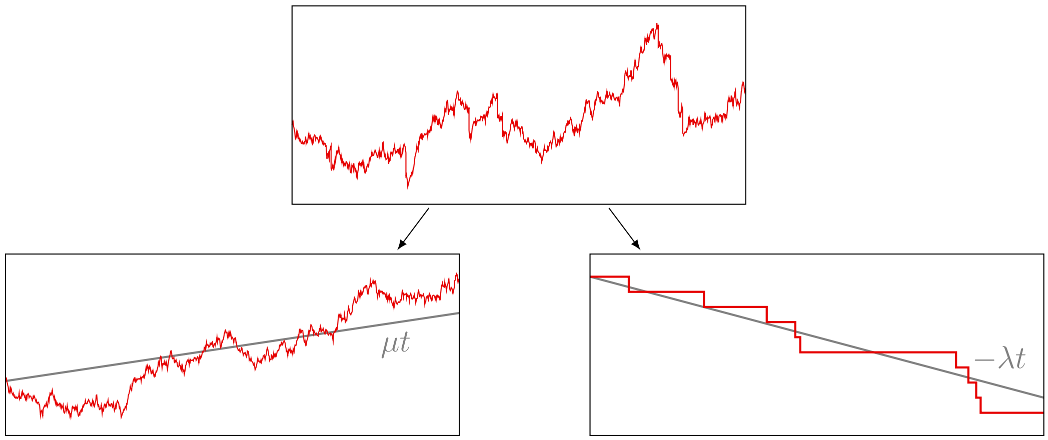

Figure 1: The total revenue of the partnership (top panel) consists of normal market fluctuations (bottom left panel) and shocks due to bad press (bottom right panel). Both components carry important information about past play: the expected rise in continuous revenue μ and the intensity of the scandals λ depend on the chosen effort levels.

an increase in industrial production or an increase in the atmospheric greenhouse gas concentration—information that is suggestive but not conclusive proof that the country has violated the agreement. These observations or measurements are possible at any point in time and are best modeled using a continuous but noisy signal. In addition to these continuous processes, countries also observe infrequent but informative political and economic events such as the passing of an environmental bill or the commissioning of a coal power plant. These events are inherently discrete and are better modeled using counting processes that jump when one of these events is observed. In another example, consider a partnership between two firms, where each firm chooses hidden effort levels and observes only the total revenue of the partnership. The total revenue moves continuously due to day-to-day fluctuations in supply and demand conditions and it is subject to demand shocks when one of the firms receives bad press. As illustrated in Figure 1, a decomposition of the information leads to two separate signals that are both indicative of the partner's effort level: the continuous increase in total revenue without the impact of demand shocks and the frequency at which scandals occur. The goal of this paper is to characterize the set of all equilibrium payoffs in two-player games, where information may arrive both continuously and abruptly. Compared to the existing literature, the more general information structure allows for a wider class of incentives that can be provided to players, thereby fundamentally changing the shape of the equilibrium payoff set.

Repeated games provide a very tractable framework to model sustained interaction between strategic decision makers. To a large extent, the tractability stems from the time homogeneity of the decision problem: because players face the same decision tomorrow as they do today, subgame perfect equilibria can be constructed in a recursive fashion as demonstrated in the works of Abreu, Pearce and Stacchetti [2, 3].

While these recursive techniques have been extremely fruitful in understanding incentives in repeated interactions, the continuous-time methods may give rise to explicit results that are not available in discrete time. In his seminal paper, Sannikov [20] 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. Such an explicit characterization of the equilibrium payoff set for any level of discounting is a result without analogue in discrete time, where results on equilibrium payoffs are often limited to folk theorems and payoff bounds; see Fudenberg, Levine and Maskin [10] and Fudenberg and Levine [8], respectively. Crucial for Sannikov's result is the assumption that information is Brownian, that is, information arrives continuously in small but noisy increments. This assumption, however, does not come without loss of generality.[1] Aside from the modeling perspective that certain events are inherently discrete, the assumption on Brownian information also restricts the types of incentives that can be provided to players in equilibrium. As argued in Sannikov and Skrzypacz [22], it is too costly to attach mutual punishments to undesirable outcomes of Brownian information: because Brownian information is so noisy, the chance of wrongfully entering a punishment phase is too large and too much value is destroyed in expectation. Therefore, incentives at extremal equilibrium payoff pairs in Sannikov [20] are restricted to tangential transfers of value between players. While these types of incentives are sufficient to prove a folk theorem, we know from Green and Porter [11] that the destruction of value upon the arrival of an undesirable signal can be an efficient way of providing equilibrium incentives to impatient players. This paper characterizes the equilibrium payoff set for any level of discounting in an information structure that allows for both types of incentives in equilibrium. This is achieved by complementing the Brownian information structure with the observation of infrequent events, whose arrival times are governed by Poisson processes.

When information arrives also abruptly, the equilibrium payoff set may have corners outside the set of static Nash payoffs. These corners correspond to *stationary payoffs* that may arise when incentives are provided through the abrupt information only. When incentives depend on abrupt information exclusively, only two factors impact the players' continuation values: the current extraction of instantaneous flow payoffs and the expected punishments/rewards that players receive upon the arrival of a discrete event. At extremal equilibrium payoff pairs, these two forces offset

[1]In the theory of discrete-time repeated games with public monitoring, barely any distributional assumption have to be made on the public signal. The only restriction is that the signals in different periods are independent and conditionally identically distributed, given the played action profile, to ensure homogeneity over time periods. Public signals that satisfy a similar time homogeneity property in a continuous-time setting are Lévy processes—a class of signals that includes not only Brownian motion, but also discontinuous processes with independent increments.

each other precisely so that the players' continuation values remain locally constant. At stationary payoffs, any equilibrium profile prescribes a locally constant play of actions until the next infrequent event is observed. Indeed, since the continuous information is not used to provide incentives, relevant information arrives only with the arrival of infrequent events. Until such an event occurs, players receive no new information that warrants an adjustment of their chosen actions.

New in this paper is also the the characterization of the equilibrium payoff set when the observed information is Brownian but does not satisfy the pairwise identifiability condition assumed in Sannikov [20]. In particular, the characterization is new for any game with a one-dimensional signal, which contains the important applications of a Cournot duopoly in a single homogeneous good and partnership games, where only the total revenue is observed. When action profiles fail to be pairwise identifiable, deviations by the two players cannot necessarily be distinguished statistically by observing the public signal. Fudenberg, Levine, and Maskin [10] have shown that a pairwise identifiable action profile can be enforced by transferring value between players at any rate. Without pairwise identifiability, this is no longer the case and instead, there may be a minimal/maximal rate at which players are willing to transfer value. While this may result in a collapse of the equilibrium payoff set to the set of static Nash payoff in some games, in other games the players locally keep transferring value at these limiting rates. In these games, the boundary of the equilibrium payoff set may have straight line segments, where the slope corresponds to the minimal/maximal rate at which players transfer value in equilibrium.

At corners or straight line segments of the equilibrium payoff set, the unique equilibrium incentives may necessitate that the continuation value enters the interior of the equilibrium payoff set with certainty. This stands in contrast to the setting of Sannikov [20], where the continuation value is absorbed on the boundary once the boundary is reached. This shows that a bang-bang property does not hold for continuous-time games if players observe abrupt information or if the continuous information does not satisfy the pairwise identifiability condition.

We find that neither type of information dominates the other in terms of impact on equilibrium payoffs. The two types of information merely serve different purposes. Abrupt information may enlarge the set of enforceable action profiles by attaching large punishments/rewards to the observation of infrequent events. Information through these infrequent events, however, arrives too sparsely for players to react to each other dynamically. Therefore, incentives through the continuous signal are necessary outside the set of stationary payoffs. Moreover, because abrupt information is tied to the burning of value, it is used on the efficient frontier only if insufficient incentives can be provided through transfers, that is, abrupt information is used to provide the residual incentives that cannot be provided by the continuous information.

From a methodological point of view, the treatment of abrupt information is quite challenging. Sannikov [20] shows that when information is Brownian, only local in-

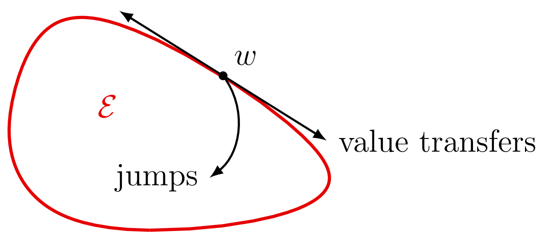

Figure 2: Because Brownian information is used to transfer value between players along tangents, only local information about the geometry of the equilibrium payoff set $\mathcal{E}$ is needed to structure incentives at a payoff pair w on the boundary. In contrast, abrupt information may induce jumps in the continuation value that are required to remain in $\mathcal{E}$—requiring global information about $\mathcal{E}$.

formation 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 about a function to encode its global properties. Crucial in this respect is that Brownian information is used only to transfer value between players tangentially. Equilibrium incentives at the boundary thus depend on the equilibrium payoff set only through the direction of the tangent, i.e., only through local information about the geometry of the equilibrium payoff set. In contrast, when information arrives according to Poisson processes, such signals can be used also to provide incentives via value burning. 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 the equilibrium incentives, global information about the set of equilibrium payoffs is required. The differential equation describing the boundary of the equilibrium payoff set 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. Therefore, this paper also contributes to the literature on computing equilibrium payoffs, providing an alternative to Judd, Yeltekin and Conklin [14] for two-player games and an extension of Abreu and Sannikov [4] to imperfect monitoring. The concept of relaxed self-generating payoff sets may be of interest in itself for future works in continuous-time games where certain information arrives according to Poisson processes. This is the case, for example, in continuous-time stochastic games with finitely many states.

A variety of papers have analyzed discrete-time repeated games with frequent actions, where the length of the time period shrinks to 0. Abreu, Milgrom, and Pearce [1] have shown that in games with Poisson signals, the limit behavior as the one-period discount factor converges to 1 crucially depends on whether the discount rate is decreased or the frequency of actions is increased. Increasing the frequency of actions may destroy incentives if the signals are not sufficiently informative and hence the likelihood of erroneous punishments is high. Sannikov and Skrzypacz [21] show that incentives are always destroyed in a Cournot duopoly with Brownian information as players act more frequently, leaving only static Nash equilibria in the limit. Fudenberg and Levine [9] analyze asymptotic efficiency in games, where players observe cumulative signals of either a Poisson or a diffusion process. They find that efficient limit behavior is obtained when players' actions affect the volatility of the diffusion and extreme realizations are indications of cheating. The first paper to model the joint arrival of information through continuous and discontinuous processes is Sannikov and Skrzypacz [22]. Combining the recursive techniques in [3] with the restrictions on the use of continuous and abrupt information, they establish a payoff bound for discrete-time games when the length of the time period approaches zero. The current paper deviates from this stream of literature by directly studying the continuous-time game. This allows a nuanced comparison of the different effects that continuous and discontinuous information have on equilibrium payoffs: the trade-off between the two types of information is captured in a sequence of ODEs that approximates the boundary of the equilibrium payoff set. Moreover, our techniques allow us to elicit equilibrium strategies that attain extremal equilibrium payoff pairs. A numerical solution of the characterizing sequence of ODEs computes the unique enforceable action profile and its enforcing equilibrium incentives in each state of the limiting solution. This determines equilibrium play uniquely on the boundary. For example, in the partnership game with demand shocks presented in Section 3, any payoff pair on the efficient frontier can be attained by play of one-sided effort until some random event occurs, followed by a forgiving grim trigger strategy profile.

In terms of provided incentives, this paper differs from [22] in the fact that only bounded amounts of value can be transferred or destroyed upon the arrival of an infrequent event. This is not a difference in the underlying model, but rather in the result that we prove: Cooperation between players requires larger punishments or rewards when players are impatient. To enforce equilibrium behavior, the continuation payoffs after the punishments/rewards have to remain in the bounded equilibrium payoff set. As players get more and more impatient, fewer incentives can be provided through the observation of infrequent events. When players get too impatient, the provided incentives could be insufficient to support equilibrium bahavior outside the set of static Nash profiles and the equilibrium payoff set could collapse to the set of static Nash payoffs. This is a feature not observed in [22] as their payoff bound relies on incentives that can be provided when players are arbitrarily patient.

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 in Section 5 and show how it is applied to approximate the equilibrium payoff set. In Section 6, we present our main results: the characterization of any relaxed-self generating payoff set and the iterative construction of equilibria. We discuss important implications of our main result in Section 7 as well as how our result relates in more detail to the existing literature. 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–E.

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 $\mathcal{A} = \mathcal{A}^1 \times \mathcal{A}^2$. Rather than directly observing each other's actions, players see only the impact of the chosen actions on the distribution of a random signal. The public signal contains continuous, but noisy information modeled by a d-dimensional Brownian motion Z and informative, but infrequent observations of events of type $y \in Y$. We assume that there are finitely many (possibly zero) different types of events in $Y = \{y_1, \ldots, y_m\}$. Events arrive according to Poisson processes $(J^y)_{y \in Y}$ that are independent from 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 as a public randomization device. Events of different types are thus observable but their underlying intensities are not. Because we study perfect public equilibria, a player's choice of action at time t must be based solely on information in $\mathcal{F}_t$, which is formalized in the following definition.

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 $\lambda(y \mid \cdot\,) : \mathcal{A} \to (0, \infty)$ 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. Let $\lambda(a) := \big(\lambda(y_1 \mid a), \ldots, \lambda(y_m \mid a)\big)^\top$ denote the vector of intensities of all events. We assume that events of any type y are possible after any history, that is, it is a game of full support public monitoring.

Assumption 1 (Full support)**.** $\lambda(y \mid a) > 0$ 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 $Q^A = \big(Q^A_t\big)_{t\geq 0}$, under which players observe the public signal.[2] On $[0,T]$ for any $T > 0$, the public signal signal takes the form

$$X_t = \int_0^t \mu(A_s)\,\mathrm{d}s + Z^A_t + \sum_{y\in Y} h(y)\,J^y_t,$$

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

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

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

Indeed, the information carried by $\tilde{X}$ is identical to the information in $X = \sigma\tilde{X}$.[3]

Anderson [5] and Simon and Stinchcombe [23] demonstrate that seemingly simple strategies need not necessarily lead to unique outcomes in continuous-time games of perfect monitoring. This is not a problem in our model because actions taken by agents do not immediately generate information: Assumption 1 in conjunction with the unbounded support of the normal distribution implies that any outcome is possible after play of any strategy profile. In public monitoring games, one can identify the probability space with the path space of all publicly observable processes, and hence, a realized path $\omega \in \Omega$ leads to the unique outcome $A(\omega)$. This is analogous to discrete-time repeated games with full-support public monitoring; see Mailath and Samuelson [17] for a thorough exposition of discrete-time games.

[2] 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 $Q^A = \big(Q^A_t\big)_{t\geq 0}$ is defined via its density process relative to P, given by

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

[3] Contrary to discrete-time models with frequent actions (c.f. Fudenberg and Levine [9]), 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.

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

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

$$W^i(A) \geq W^i\big(\tilde{A}^i, A^{-i}\big) \text{ a.e.,}^4$$

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

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

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

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 $r\mathrm{e}^{-r(s-t)}$ 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 $\mathcal{V} := \operatorname{conv}\{g(a) \mid a \in \mathcal{A}\}$. 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 $\mathcal{V}^N := \operatorname{conv}\big\{g(a) \,\big|\, a \in \mathcal{A}^N\big\}$ 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.

	0	1
0	0, 0	4, -1
1	-1, 4	2, 2

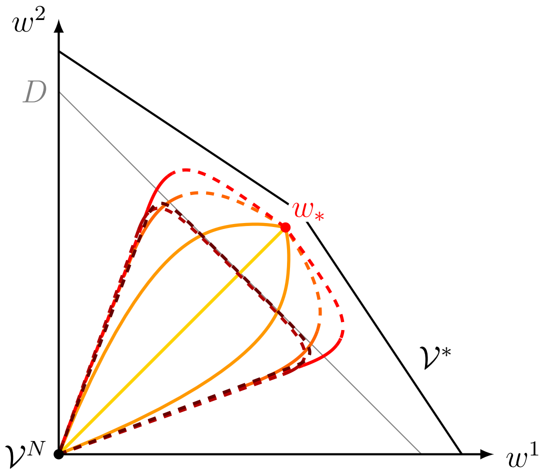

Figure 3: The left panel shows the matrix of stage game payoffs. The right panel shows the equilibrium payoff set $\mathcal{E}_\gamma(5)$ as a subset of $\mathcal{V}^*$ for γ ranging from $\gamma = 0$ (dark red) to $\gamma = 1$ (yellow) in increments of size 0.2. The boundary is dashed where the unique equilibrium incentives do not make use of the abrupt information. If the continuous signal is relatively uninformative (γ high), abrupt information is required to enforce non-trivial equilibria. As the informativeness of the continuous signal increases, players can transfer larger amounts of value in equilibrium, which widens the equilibrium payoff set. If the abrupt information is relatively uninformative (γ low), the action profile of mutual effort is not enforceable, hence $\mathcal{E}_\gamma(5)$ collapses below the negative diagonal D.

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 cannot observe each other's effort levels and instead see only the 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. The continuous stream of revenue satisfies $\mathrm{d}X^\gamma_t = \mu_\gamma(A_t)\,\mathrm{d}t + \mathrm{d}Z^{A,\gamma}_t$ for a Brownian motion $Z^{A,\gamma}_t$, where

$$\mu_\gamma(a) = (1 - \gamma)\big(4(a^1 + a^2) - a^1a^2\big).$$

On average, the continuous stream of revenue increases with effort but the increase diminishes with additional effort. The arrival of demand shocks is governed by a Poisson process $J^{A,\gamma}$ with instantaneous intensity $\lambda_\gamma(A_t)$, where

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

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

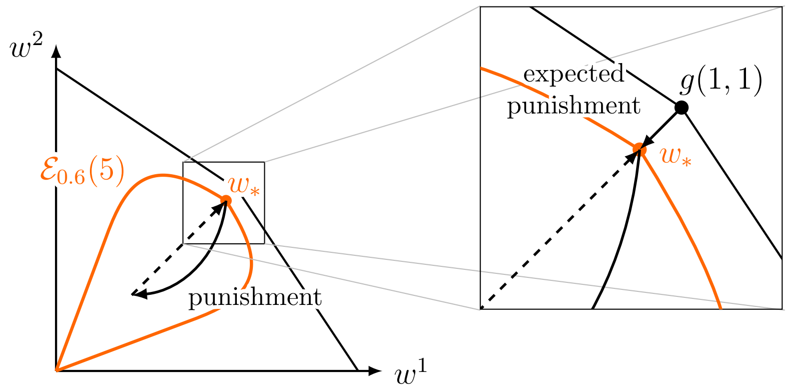

Figure 4: The equilibrium payoff set $\mathcal{E}_{0.6}(5)$ is shown together with the path of the continuation value of a forgiving grim-trigger strategy profile attaining w_*. The solid arrow indicates the minimum amount of value that has to be burnt upon the arrival of a demand shock to enforce the action profile of mutual effort. During the Nash punishment phase, 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.

Demand shocks arrive rarely when both players are exerting effort but they arrive not that much more frequently if neither player is exerting effort than when only one player is exerting effort. 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. Note that the characterization is new even for the continuous-monitoring game ($\gamma = 0$) because the signal is one-dimensional.

As the figure shows, neither type of information dominates the other in terms of impact on equilibrium payoffs. The two types of information merely serve different purposes. Abrupt information may enlarge the set of enforceable action profiles because of the added possibility for value burning: by burning at least $5/(8\gamma)$ payoff units for each player upon the arrival of a demand shock, the action profile of mutual effort is enforceable. If the demand shocks are sufficiently informative ($\gamma \geq 1/3$), the action profile of mutual effort can be enforced in equilibrium and the most efficient symmetric equilibrium payoff pair is $w_* = (1.875, 1.875)$. The payoff pair w_* is stationary, that is, the players' equilibrium continuation values remain constant while both players exert effort. Only when a demand shock occurs, players burn value by, say, entering a Nash punishment phase for a certain amount of time; see Figure 4. If the demand shocks are uninformative ($\gamma < 1/3$), the equilibrium payoff set collapses below the diagonal D, which connects the stage-game payoffs of one-sided effort $g(1,0)$ and $g(0,1)$. Information through demand shocks, however, arrives too sparsely for players to react to each other dynamically. Because the time interval between two demand shocks is unbounded, players cannot rely exclusively on abrupt information outside the set of stationary payoffs. The continuous information serves for players to

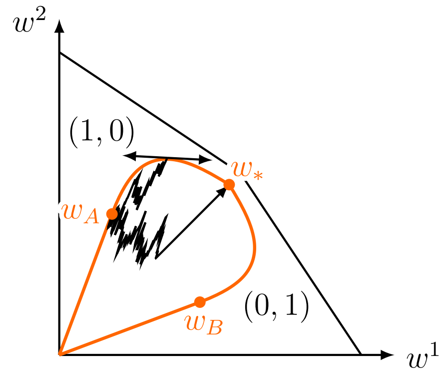

Figure 5: On curved parts of the boundary, only one player exerts effort. The action profile of one-sided effort is enforced using tangential transfers. At the payoff pairs w_A and w_B, where the boundary becomes straight, the continuation value of any equilibrium enters the interior of $\mathcal{E}_{0.6}(5)$. Two such sample paths starting at w_A are shown.

react to each other dynamically. As the informativeness of the continuous information increases, Figure 3 shows that the equilibrium payoff set gets wider because players are able to transfer larger amounts of value between each other in equilibrium.

Because the use of abrupt information is necessarily tied to the burning of value where the boundary is curved (as jumps into the set lie below the tangent), abrupt information is used sparingly on the efficient frontier. Figure 3 shows that abrupt information is used on the boundary only where the curvature is large: on the efficient frontier, incentives are provided through value transfers based on the continuous information for as long as these incentives are sufficient. Because smaller tangential transfers are possible in equilibrium where the curvature is large, these incentives are insufficient and abrupt information is used to provide the residual 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 behavior 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 strategy profile as illustrated in Figure 4, but other forms of punishment are possible in equilibrium as we will elaborate in Section 7. For lower values of γ, the demand shocks arrive less frequently, hence more severe punishments are necessary to deter deviations. For $\gamma = 1/3$, the only sufficiently strong punishment 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 conceivable 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 5. The action profiles of one-sided effort are enforced by transferring value between players tangentially to the set. The minimal/maximal rates, at which players are willing to transfer value, are reached at the payoff pairs w_A and w_B, respectively, as indicated in Figure 5. Where the boundary is strictly curved, the continuation value locally remains on the boundary until a demand shock occurs. At the payoff pairs w_A and w_B, however, the continuation value enters the interior of the equilibrium payoff set with certainty; see Figure 5. Players keep transferring value at the extremal rate until the continuation value either reaches the boundary again or it reaches the positive diagonal of symmetric payoff pairs. In the former case, the same player keeps exerting effort and incentives are provided through tangential transfers. Because of the volatility in the players' continuation value, the diagonal is reached eventually either in the interior of the set or at w_*. On the diagonal, players adopt a forgiving grim-trigger strategy profile (starting with the punishment phase if the continuation value is in the interior of the set). Therefore, any extremal equilibrium payoff pair can be attained by the equilibrium profile of one-sided effort followed by a forgiving grim-trigger strategy profile.

4 Enforceability and self-generation

As in any game with imperfect monitoring, players' incentives are tied to the public signal. We thus start this section by stating the dependence of the continuation value on the public signal. The following stochastic differential representation is the extension of Proposition 1 in Sannikov [20] to games with abrupt information.

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

(i) W is the discounted expected payoff under A.

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

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

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 $\mathbb{E}_{Q_T^A}\left[\int_0^T |\beta_t^i|^2\,\mathrm{d}t\right] < \infty$ and $\mathbb{E}_{Q_T^A}\left[\int_0^T |\delta_t^i(y)|^2 \lambda(y\,|A_t)\,\mathrm{d}t\right] < \infty$ for any $T \geq 0$.

The process $r\beta^i$ is the sensitivity of player i's continuation value to the continuous component of the public signal and the processes $r\delta^i(y)$ are the impacts on player i's continuation value when an event of type $y \in Y$ occurs. Note that in expectation, the impact of the public signal on the continuation value averages out as $Z - \int_0^\cdot \mu(A_t)\,\mathrm{d}t$ and $J^y - \int_0^\cdot \lambda(y|A_t)\,\mathrm{d}t$ for every $y \in Y$ are martingales under Q^A. Nevertheless, the exposure to the public signal is relevant to provide incentives as we shall see below. In expectation, player i's continuation value moves away from the expected flow payoff rate $g^i(A)$ towards $W^i(A)$: if player i currently extracts a higher/lower payoff rate than he receives in expectation, this has to be balanced out in the future by decreasing/increasing the player's continuation value. To keep the notation succinct, we use δ^i to refer to the row vector $(\delta^i(y_1), \ldots, \delta^i(y_m))$ containing the impacts on i's continuation payoffs for all types of events.

In discrete-time games, incentives are provided by a *continuation promise* that maps the public signal to a promised continuation payoff for every player; see, for example, Abreu, Pearce and Stacchetti [3]. The representation in (2) shows that in continuous-time games, the continuation value is linear in the public signal and hence, so is the continuation promise. Similarly to Sannikov [20] and Sannikov and Skrzypacz [22], the incentive compatibility conditions take the following form.

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 $\tilde{a}^i \notin \mathcal{A}^i \setminus \{a^i\}$,

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

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.

The expression in (3) highlights that a deviation of a player i impacts his discounted expected future payoff in two ways: by the change of expected flow payoff $r\big(g^i(\tilde{A}^i_t, A^{-i}_t) - g^i(A_t)\big)\,\mathrm{d}t$ and by the change in the distribution of the public signal. For sensitivities (β, δ) in (2), this change of distribution has an expected impact on player i's continuation value of $r\beta_t\big(\mu(\tilde{A}^i_t, A^{-i}_t) - \mu(A_t)\big)\,\mathrm{d}t + r\delta_t\big(\lambda(\tilde{A}^i_t, A^{-i}_t) - \lambda(A_t)\big)\,\mathrm{d}t$. Thus, if the sensitivities (β, δ) of the continuation value to the public signal are the continuation promises made to enforce A—that is, the promises are kept—then no player has an incentive to deviate from A. This is formalized in the following lemma, which is a generalization of Proposition 2 in Sannikov [20] to our setting.

Lemma 4.3. *A strategy profile is a PPE if and only if (β^1, β^2) and $\big(\delta^1(y), \delta^2(y)\big)_{y\in Y}$ 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). Because there is no notion of a terminal value, the theory of backward stochastic differential equations cannot be applied to infinitely repeated games. Instead, we use the concept of self-generating payoff sets to construct forward solutions similarly to discrete time.

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

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. The proof is easily derived from the proof of Lemma 2 in Bernard and Frei [7] as that proof works more generally for signals given by any Lévy process.

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

5 Iteration over arrival of abrupt information

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 6. In this section, we illustrate that in the presence of abrupt information, these restrictions depend on the equilibrium payoff set itself. Any attempt at describing the boundary of the equilibrium payoff set via equilibrium incentives thus leads to a self-referential description. This motivates the introduction of an iterative procedure over the arrival times of infrequent events that will lead to an iterative construction of equilibrium profiles as well as an approximation of the equilibrium payoff set.

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 $\mathcal{N}_w(\mathcal{W}) := \{N \in S^1 \mid N^\top(w - v) \geq 0 \text{ for all } v \in \mathcal{W}\}$ 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:

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

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

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

Sannikov [20] shows that 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 (i) and (ii) 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

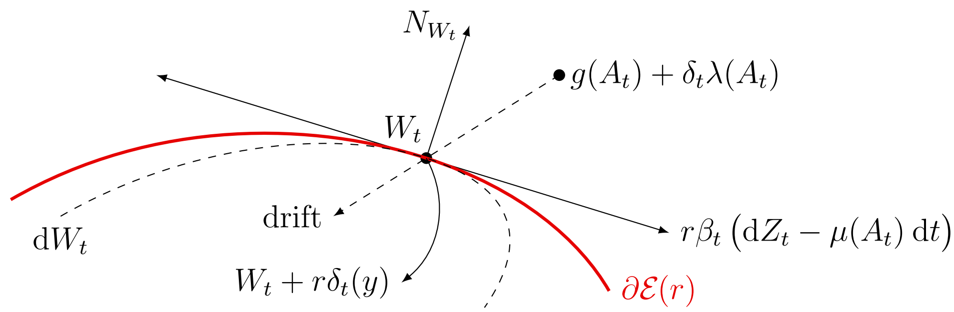

Figure 6: Because $\mathcal{E}(r)$ is self-generating, the continuation value W of a PPE can never escape the set $\mathcal{E}(r)$. At the boundary $\partial\mathcal{E}(r)$, the drift rate $r\big(W_t - g(A_t) - \delta_t\lambda(A_t)\big)\,\mathrm{d}t$ thus has to point towards the interior of the set. Moreover, the diffusion $r\beta_t\,\big(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t\big)$ has to be tangential to $\partial\mathcal{E}(r)$ as the continuation value would escape $\mathcal{E}(r)$ otherwise due to the unbounded variation of Brownian motion. Finally, an event of type $y \in Y$ incurs a jump in the continuation value of size $r\delta_t(y)$. Since W cannot jump outside of $\mathcal{E}(r)$, it is necessary that $W + r\delta(y) \in \mathcal{E}(r)$ for every $y \in Y$.[5]

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 the third informational restriction is a global restriction involving the precise shape of the equilibrium payoff set. A local description of $\mathcal{E}(r)$ using restrictions (i)–(iii) 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 restriction (iii) is relaxed to the jumps landing in a fixed payoff set $\mathcal{W}$.

Definition 5.1. Let σ_n denote the occurrence of the n^{th} infrequent event.

(i) We say that a payoff set $\mathcal{X}$ is $\mathcal{W}$-*relaxed self-generating* for a payoff set $\mathcal{W}$ if for every $w \in \mathcal{X}$, there exists a solution W to (2) for processes A, β, δ, and M such that (β, δ) enforces A, $W_0 = w$ a.s., $W_\tau \in \mathcal{X}$ a.s. for every stopping time $\tau < \sigma_1$, and $W_{\sigma_1} \in \mathcal{W}$ a.s.

(ii) For a convex and compact set $\mathcal{W} \subseteq \mathbb{R}^2$, let $\mathcal{B}_r(\mathcal{W})$ denote the largest $\mathcal{W}$-relaxed self-generating set. Observe that this is well defined since the convex hull of two $\mathcal{W}$-relaxed self-generating sets is again $\mathcal{W}$-relaxed self-generating.

Any payoff pair in a $\mathcal{W}$-relaxed self-generating payoff set $\mathcal{X}$ can be attained by an enforceable strategy profile whose continuation value remains in $\mathcal{X}$ until the arrival of the first infrequent event, at which point the continuation value jumps to $\mathcal{W}$. At the boundary of $\mathcal{X}$, players must play an action profile such that the continuation value

[4] Jump times of Poisson processes are totally inaccessible, which means that players have absolutely no way of anticipating 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 punishments/rewards upon the arrival of an event are consistent with equilibrium behavior.

has inward-pointing drift, tangential volatility, and jumps that land in $\mathcal{W}$. Because the incentives provided through the abrupt information do not depend on $\mathcal{X}$, the set of admissible incentives at the boundary depends on $\mathcal{X}$ only through local information, which makes a description of the boundary via incentives possible.

The operator $\mathcal{B}_r$ is a continuous-time analogue to 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 at time σ_1 that lies in $\mathcal{W}$. If $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$, then the payoff pair W_{σ_1} can be attained by an enforceable strategy profile until the arrival of the next event and so on. The following lemma characterizes the relation between the operator $\mathcal{B}_r$ and 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.*

An n-fold application of $\mathcal{B}_r$ to a set $\mathcal{W}$ thus ensures that the continuation value after the first n events is in $\mathcal{W}$. Because Poisson processes have only countably many jumps, taking the limit as n goes to infinity covers all events. We thus obtain the following algorithm to compute $\mathcal{E}(r)$ iteratively.

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

This algorithm is similar to the algorithm in Abreu, Pearce and Stacchetti [3]. However, unlike its discrete-time counterpart, we show in the next section that the boundary of the resulting set at each step of the iteration admits a characterization by a differential equation. This is possible because in each step, the restriction on the use of abrupt information is fixed. The algorithm in Proposition 5.3 thus provides an alternative to the discrete-time implementation by Judd, Yeltekin, and Conklin [14].

6 Characterization of equilibrium payoffs

In this section we show how the iterative procedure from the previous section is used to construct equilibrium profiles. The basic idea is to construct enforceable solutions to (2) subject to the informational restrictions (i)–(iii). Due to Proposition 5.3 it is sufficient to do so up until the arrival of the first infrequent event, i.e., we will construct equilibrium profiles attaining payoff pairs in $\mathcal{B}_r(\mathcal{W})$ for some fixed convex payoff set $\mathcal{W}$. Because the boundary is a priori unknown, we abstract away from informational restrictions (i)–(iii) with the following definition.

Definition 6.1. 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 $\Xi_a(w, N, r, \mathcal{W})$ all such continuation promises is non-empty.

In the following two subsections, we will construct restricted-enforceable strategy profiles whose continuation values locally remain either constant or on a continuously differentiable curve. These two techniques will serve as building blocks to construct enforceable strategy profiles attaining payoff pairs on the boundary of $\mathcal{B}_r(\mathcal{W})$.

6.1 Stationary payoffs

The simplest $\mathcal{W}$-relaxed self-generating payoff sets consist of only a single payoff pair w. These payoff pairs are attainable by a strategy profile whose continuation value remains in w until an infrequent event occurs. We call such a payoff pair $\mathcal{W}$-stationary. It is convenient to work with the following more elementary but equivalent definition of stationary payoffs.

Definition 6.2. 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 $y \in Y$.

We denote by $\mathcal{S}_r(\mathcal{W})$ the set of all $\mathcal{W}$-stationary payoffs. Observe that any Nash payoff pair in $\mathcal{W}$ is $\mathcal{W}$-stationary for $\delta_0 = 0$, i.e., $\mathcal{V}^N \cap \mathcal{W} \subseteq \mathcal{S}_r(\mathcal{W})$. The set of stationary payoffs is contained in $\mathcal{B}_r(\mathcal{W})$ as formalized by the following lemma.

Lemma 6.3. *$\mathcal{S}_r(\mathcal{W}) \subseteq \mathcal{B}_r(\mathcal{W})$.*

Proof. Because the union of two $\mathcal{W}$-relaxed self-generating sets is again $\mathcal{W}$-relaxed self-generating, it is sufficient to 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 $W_\sigma = W_{\sigma-} + r\delta_{\sigma-}(y) = w + r\delta_0(y) \in \mathcal{W}$. $\square$

6.2 Restricted-enforceable strategy profiles on curves

The construction of restricted-enforceable strategy profiles, whose continuation value remains on a curve, requires players to continuously adjust the tangential transfers provided by the continuous information. For players to be able to do so, the set of directions, in which an action profile is restricted-enforceable, has to be closed. We show in the appendix that if an action profile a is restricted-enforceable along a convergent sequence of directions $(N_n)_{n \geq 0}$, then a is also restricted-enforceable along the limiting direction $\lim_{n\to\infty} N_n$ unless the limiting direction is a *coordinate* direction in $\{\pm e_1, \pm e_2\}$. To ensure that this property holds true also for coordinate limiting directions, we impose Assumption 2 below.

For any action profile $a \in \mathcal{A}$, let Ψ_a^i denote the set of $\delta^i \in \mathbb{R}^m$, for which continuation promise $(0, \delta^i)$ provides sufficient incentives to player i to support play of a, i.e., $(0, \delta^i)$ is a solution to (3) for player i. We make the following non-empty interiority assumption, stating that if the discontinuous component of the public signal is sufficient to incentivize player i to play a^i, given a^{-i}, then player i can be strictly incentivized to play a^i, given a^{-i}, using only the discontinuous information. Observe that the assumption is satisfied for generic choices of g and λ.

Assumption 2. Suppose Ψ_a^i for $i = 1, 2$ has non-empty interior for any $a \in \mathcal{A}$.

Remark 6.1. Assumption 2 is a generalization of Assumption 2.(i) in Sannikov [20] as it reduces to a unique best response assumption when $Y = \emptyset$. Note however, that we do not require action profiles to be pairwise identifiable. We are able to relax this condition despite our general framework by analyzing continuity properties of incentives that are "optimal" on the boundary; see Appendices B and E.2 for details.

Without pairwise identifiability, action profiles may be restricted-enforceable in some directions but not in others. For payoff pairs and directions, at which they are restricted-enforceable with non-zero tangential transfers, enforceable strategy profiles on curves can be constructed with the following lemma.

Lemma 6.4. *Suppose that Assumptions 1 and 2 hold. Let $\mathcal{C}$ be a continuously differentiable curve oriented by the Gauss map $w \mapsto N_w$ such that:*

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

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

(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 $A = a^*(W)$, $\delta = \delta^*(W_-)$, $\beta = \beta^*(W)$, and $M \equiv 0$ remains on $\mathcal{C}$ until an endpoint of $\mathcal{C}$ is reached or an event occurs.

The above lemma enables us to construct enforceable strategy profiles that remain on a curve with curvature (4). While the presence of abrupt information creates many technical challenges that we address in Appendices B and C, the intuition behind Lemma 6.4 is similar to Sannikov [20] because the diffusion term is the main driver behind Lemma 6.4. The strategy profile constructed in Lemma 6.4 and its continuation promises are Markovian in the continuation value. At any point w on

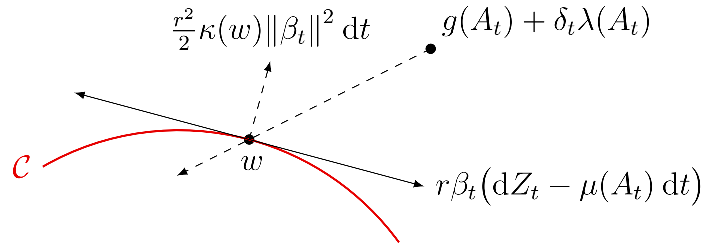

Figure 7: Incentives related to the continuous component of the public signal lead to tangential volatility of the continuation value. Because Brownian motion has unbounded variation, the tangential volatility leads to a second-order outward drift. At points where the curvature is larger, the tangential volatility leads to a stronger outward drift. For continuation promise (β, δ) in Lemma 6.4, this outward drift is precisely counteracted by the inward drift $rN_w{}^\top\big(w - g(A_t) - \delta_t\lambda(A_t)\big)$ so that the continuation value remains on $\mathcal{C}$ until an end point is reached or an infrequent event occurs.

the curve, incentives provided through the continuous component of the public signal are parallel to the curve. Because of the infinitesimally fast oscillation of Brownian motion, the tangential volatility leads to an outward-pointing drift; see Figure 7. It follows form Itō's formula that the outward drift is proportional to the curvature and the square of the tangential incentives. For the constructed strategy profile, this outward drift is precisely counteracted by the inward-pointing drift so that at w,

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

and the continuation value remains on the curve $\mathcal{C}$. Because δ is chosen such that $W + r\delta(y) \in \mathcal{W}$ almost everywhere, it is guaranteed that the continuation value jumps to $\mathcal{W}$ after the occurrence of any event $y \in Y$. This construction of enforceable strategy profiles will be used several times in the characterization of $\partial\mathcal{B}_r(\mathcal{W})$.

6.3 Decomposition of extremal payoff pairs

Consider a payoff set, whose boundary consists of only stationary payoff pairs and solutions to (4) for some fixed set $\mathcal{W}$. Then any payoff pair on the boundary is attainable by an enforceable strategy profile, whose continuation value remains on the boundary until an infrequent event occurs and jumps to $\mathcal{W}$ at the time of the first event. Such a payoff set is thus contained in $\mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. One may thus wonder whether the boundary of $\mathcal{B}_r(\mathcal{W})$ consists only of stationary payoff pairs and solutions to (4). Unfortunately the answer is negative in general. There may be payoff pairs on the boundary that are neither stationary nor require the continuous information to provide incentives.

Definition 6.5. A payoff pair $w \in \partial\mathcal{B}_r(\mathcal{W})$ is called $\mathcal{W}$-*decomposable* if there exist a and δ such that $(0, \delta) \in \Xi_a(w, N, r, \mathcal{W})$ for all outward normal vectors $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. We say that such a pair (a, δ) *decomposes* w.

Remark 6.2. This is similar to the one-period decomposition of payoffs in Fudenberg, Levine, and Maskin [10], where a payoff vector on the boundary of the equilibrium payoff set is decomposed into a current-period payoff vector outside the set and a continuation value within the set. This is not surprising because a continuous-time game with abrupt information only can be mapped to a discrete-time game with time periods of random length. Time period n lasts from the occurrence of the $n-1^{\text{st}}$ event to the occurrence of the n^{th} event. At the end of period n, the signal Y_n is equal to y if and only if the n^{th} event is of type y. However, while the techniques related to abrupt information are often similar to those used in discrete time, they are not identical because of the crucial difference that the interarrival times between two abrupt events are unbounded. Because players do not know when they will receive new information, incentives cannot be adjusted dynamically.

We denote by $\mathcal{D}_r(\mathcal{W})$ the set of all $\mathcal{W}$-decomposable payoff pairs. Referring back to the definition of Ξ_a in Definition 6.1, we note that the defining characteristics of a $\mathcal{W}$-decomposable payoff pair are the following: incentives are provided through the abrupt information only such that the drift rate points towards the interior of $\mathcal{B}_r(\mathcal{W})$ and the continuation payoff after the occurrence of an event is in $\mathcal{W}$. The following lemma establishes that at most one of these defining conditions holds strictly.

Lemma 6.6. *Suppose that Assumptions 1 and 2 are satisfied and that $\mathcal{W}$ has non-empty interior. For any $\mathcal{W}$-decomposable payoff pair w, it is impossible that there exists (a,δ) such that two of the following conditions hold simultaneously:*

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

(ii) $N^\top\big(g(a)+\delta\lambda(a)-w\big)>0$ *for some outward normal* $N\in\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$,

(iii) $w+r\delta(y)\in\operatorname{int}\mathcal{W}$ *for every* $y\in Y$.

Proof. Suppose that there exists such a pair (a,δ). Since $\mathcal{W}$ and Ψ_a have non-empty interior, there exists δ' sufficiently close to δ such that all three conditions (i)–(iii) hold simultaneously. Since the conditions are strict, they all hold for v and N' sufficiently close to w and N. Because of condition (i), there exists $\phi\in\mathbb{R}^d$ sufficiently small such that $(T\phi,\delta')$ enforces a for any direction $T\in S^1$. Let $\mathcal{C}_{w_0,\phi}$ be a solution to (4) with initial value (w_0,N) and selectors $a^*(w)=a$, $\delta^*(w)=\delta'$, and $\beta^*=T_v\phi$, where T_v is the tangent vector to $\mathcal{C}_{w_0,\phi}$ at a point v. Because of conditions (ii) and (iii), choosing $w_0\notin\operatorname{cl}\mathcal{B}_r(\mathcal{W})$ sufficiently close to w and choosing $\|\phi\|$ sufficiently small guarantees that $\mathcal{C}$ enters the interior of $\mathcal{B}_r(\mathcal{W})$ on both sides of w_0 as illustrated in the left panel of Figure 8. Lemma 6.4 thus shows that $w_0\in\mathcal{B}_r(\mathcal{W})$, a contradiction. □

A consequence to Lemma 6.6 is that corners are either stationary or that incentives are binding for at least one player i such that $w+r\delta(y)\in\partial\mathcal{W}$ for some event $y\in Y$. Indeed, if condition (ii) of Lemma 6.6 is violated and $N^\top\big(g(a)+\delta\lambda(a)-w\big)=0$

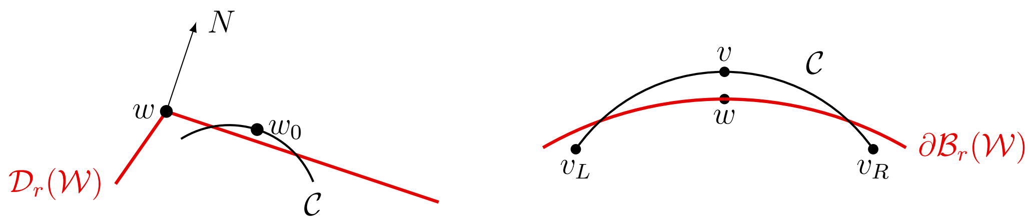

Figure 8: The left panel illustrates that it is not possible that two constraints of Lemma 6.6 are satsfied at $w \in \mathcal{D}_r(\mathcal{W})$ simultaneously. Otherwise w_0 outside of $\mathcal{B}_r(\mathcal{W})$ could be attained by an enforceable strategy profile that remains on a curve $\mathcal{C}$ that reaches the interior of $\mathcal{B}_r(\mathcal{W})$, showing that $w_0 \in \mathcal{B}_r(\mathcal{W})$ by Lemma 6.4, a contradiction. Similarly, a curve with curvature in (5) cannot intersect $\partial\mathcal{B}_r(\mathcal{W})$ outside of $\mathcal{D}_r(\mathcal{W})$ as depicted in the right panel.

for all outward normal vectors N to $\partial\mathcal{B}_r(\mathcal{W})$ at a corner w, then $w = g(a) + \delta\lambda(a)$ and hence w is stationary. If condition (ii) is satisfied, incentives have to be binding and $w + r\delta(y) \in \partial\mathcal{W}$ for at least one $y \in Y$. Therefore, there must exist an action profile a such that the corner lies on the boundary of the set

$$\mathcal{K}_{r,a}(\mathcal{W}) := \{w \mid \exists\delta \in \Psi_a \text{ with } w + r\delta(y) \in \mathcal{W} \text{ for every } y \in Y\}.$$

Corners of $\mathcal{D}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$ are thus contained in the set $\mathcal{K}_r(\mathcal{W}) = \partial \bigcup_{a\in\mathcal{A}} \mathcal{K}_{r,a}(\mathcal{W})$.[6] A perturbation argument in Appendix C shows that this is, in fact, true for all extremal payoff pairs in $\mathcal{D}_r(\mathcal{W})$.

Proposition 6.7. *Suppose that Assumptions 1 and 2 are satisfied. Any payoff pair $w \in \operatorname{ext}\mathcal{D}_r(\mathcal{W})\setminus\mathcal{S}_r(\mathcal{W})$ is decomposable by an action profile a such that $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ and $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \subseteq \mathcal{N}_w\big(\mathcal{K}_{r,a}(\mathcal{W})\big)$.*

The key restriction of the characterization of decomposable payoff pairs in Proposition 6.7 is the restriction on outward normal vectors. It implies that every corner in $\mathcal{D}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$ is a corner of $\mathcal{K}_{r,a}(\mathcal{W})$ for some action profile a and any line segment in $\mathcal{D}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$ with a strictly positive curvature also lies on $\partial\mathcal{K}_{r,a}(\mathcal{W})$ for some action profile a. The fact that corners are either stationary payoffs or that incentives have to be binding at corners is similar in spirit to Abreu and Sannikov [4], who find an improvement on the algorithm in Judd, Yeltekin, and Conklin [14] for two-player games with perfect monitoring. They find that extremal points in the discrete-time analogue of $\mathcal{B}_r(\mathcal{W})$ are attainable either by a stationary strategy profile, or by a current-period action profile with binding incentives.

[6] Observe that $\partial\mathcal{K}_{r,a}(\mathcal{W})$ is reasonably nice since $\mathcal{K}_{r,a}$ is convex: for any $w_0, w_1 \in \mathcal{K}_{r,a}(\mathcal{W})$, there exist (δ_0, δ_1) with $w_k + r\delta_k(y) \in \mathcal{W}$ for every $y \in Y$ and $k = 1, 2$. Define $w_\gamma := \gamma w_1 + (1-\gamma)w_0$ and $\delta_\gamma := \gamma\delta_1 + (1-\gamma)\delta_0$ and observe that for every $\gamma \in [0, 1]$, $(0, \delta_\gamma)$ enforces a and $w_\gamma + r\delta_\gamma(y) \in \mathcal{W}$ for every $y \in Y$ by convexity of Ψ_a and $\mathcal{W}$, respectively.

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

The following result is a complete characterization of $\mathcal{B}_r(\mathcal{W})$, building on the previous results in this section. The proof is contained in Appendices B–E.

Theorem 6.8. *Suppose that Assumptions 1 and 2 hold and that $\mathcal{W} \subseteq \mathcal{V}^*$ is compact and convex with non-empty interior. Then $\mathcal{B}_r(\mathcal{W})$ is the largest closed convex subset of $\mathcal{V}^*$ such that* $\operatorname{ext} \mathcal{D}_r(\mathcal{W}) \subseteq \mathcal{S}_r(\mathcal{W}) \cup \mathcal{K}_r(\mathcal{W})$ *and $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is continuously differentiable with curvature at almost every point w given by*

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

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

The characterization of decomposable payoff pairs follows straight from Proposition 6.7. Outside of $\mathcal{D}_r(\mathcal{W})$, the curvature of the boundary is of the same form as in Lemma 6.4, where the selectors correspond to the maximizers of the expression on the right-hand side. The intuition behind this is the same as in Sannikov [20]: if the curvature of $\partial\mathcal{B}_r(\mathcal{W})$ at an extremal payoff pair w was smaller than the expression in (5), then a solution to (5) starting at a payoff pair v slightly outside of $\mathcal{B}_r(\mathcal{W})$ would intersect $\partial\mathcal{B}_r(\mathcal{W})$ and reach payoff pairs in the interior of $\mathcal{B}_r(\mathcal{W})$; see the right panel in Figure 8. This implies by Lemma 6.4 that $v \in \mathcal{B}_r(\mathcal{W})$, which is a contradiction. Note that this argument requires continuity of (4) in initial conditions, which we establish in Appendix B. Similarly, the curvature of $\partial\mathcal{B}_r(\mathcal{W})$ cannot be larger than the maximum curvature in (5) because otherwise, there would exist no restricted enforceable strategy profile that remains in $\mathcal{B}_r(\mathcal{W})$, contradicting the fact that $\mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}$-relaxed self-generating. There are many details needed to make this argument rigorous, which can be found in Appendices C and E.2.

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 6.8, it follows from Proposition 5.3 that $\mathcal{E}(r)$ is compact. An application of Theorem 6.8 for $\mathcal{W} = \mathcal{E}(r)$ thus provides a fixed-point characterization of $\mathcal{E}(r)$ since $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$.

7 Discussion

In this section, we discuss important insights that can be gained by our main results. We also discuss in greater detail how our result relates to the literature. While some of these points have already been highlighted in the example of Section 3, the tools developed in Sections 4–6 help gaining some additional insights.

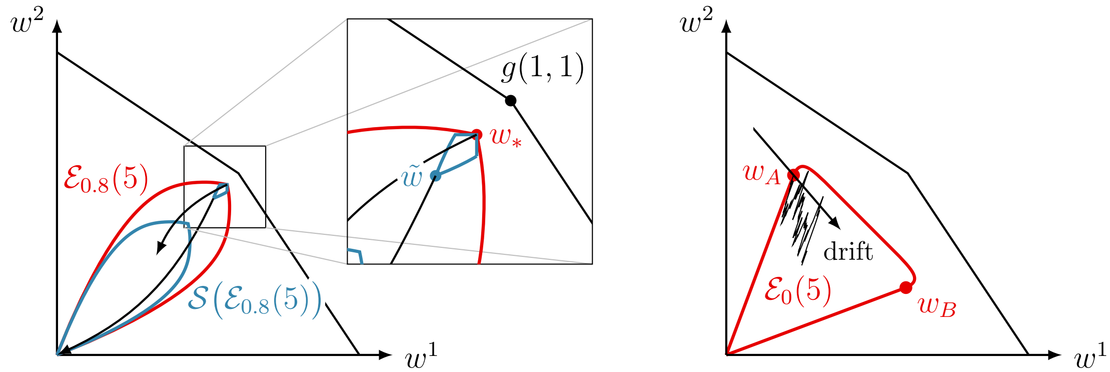

Figure 9: The left panel shows the set of stationary equilibrium payoffs as a subset of the equlibrium payoff set in the partnership example of Section 3. The two disconnected components correspond to the stationary payoffs decomposed by the action profiles of mutual effort and mutual shirking, respectively. The payoff pair $\tilde{w}$ can be attained by an unforgiving grim-trigger strategy profile. Because the value burnt upon the arrival of a demand shock is about twice as large as in the minimal equilibrium punishment, the payoff pair $\tilde{w}$ is about twice as far below the expected flow payoff pair $g(1,1)$ as w_*. The right panel illustrates that the continuation value of any equilibrium attaining w_A enters the interior of the equilibrium payoff set with certainty due to the drift term.

7.1 Bang-bang property and pairwise identifiability

In the model of Sannikov [20], the continuation value of a PPE is absorbed on the boundary of the equilibrium payoff set once the boundary is reached. This can be interpreted as a continuous-time analogue to the bang-bang proprty of Abreu, Pearce, and Stacchetti [3]. We have already highlighted in the example of Section 3 that the bang-bang property does not hold when players observe finitely many discrete types of events: the extremal payoff pair w_* in the partnership game of Section 3 is decomposed by the minimal value burning necessary to enforce mutual effort. For $\gamma > 1/3$ this minimal punishment has a continuation value in the interior of the equilibrium payoff set. Note that players cannot agree to use a larger punishment and, say, revert to the static Nash equilibrium forever as this would destroy too much value in expectation to attain w_*; see Figure 9. The fact that a permanent Nash punishment is not consistent with equilibirium behavior at w_* can also be observed using our construction of equilibria: a solution to (2) with a Nash punishment would have drift rate in the direction of $w_* - \tilde{w}$, violating the inward-drift condition.

A possible explanation for the failure of the bang-bang property in the presence of abrupt information is the fact that players observe only finitely many different types of events in our model. The embedding of the abrupt information into a discrete-time game as in Remark 6.2 thus generates a discrete-time game with a finite signal space, for which a bang-bang property does not hold in discrete time either.[7] More

[7]We emphasize that this is merely one possible explanation. The discrete-time embedding in Remark 6.2 differs from the model in Abreu, Pearce, and Stacchetti [3] crucially in the fact that

surprising is thus the fact that the bang-bang property may fail to hold even when information is entirely Brownian but action profiles are not pairwise identifiable.

An action profile is said to be pairwise identifiable if deviations of any two players can be statistically distinguished by observing the public signal. For a pairwise identifiable and enforceable action profile, sufficient incentives can be provided to players by transferring value between each other along any non-coordinate tangent (cf. Fudenberg, Levine, and Maskin [10]). If action profiles are not pairwise identifiable as in our setting, players cannot transfer value on any tangent but there may be limiting directions along which the action profile can be enforced. In the continuous-monitoring partnership game of Section 3 with $\gamma = 0$, this occurs at the payoff pairs w_A and w_B; see the right panel of Figure 9. Lemma 6.4 establishes that the continuation value of an equilibrium profile remains on the boundary of the equilibrium payoff set where the boundary is curved. As highlighted in Figure 7, the unbounded variation of the tangential transfers leads to an outward-pointing drift proportional to the curvature that is precisely counteracted by the inward-pointing drift from the extraction of expected flow payoffs. At w_A and w_B, where the curvature becomes 0, there is no outward-pointing drift anymore and hence the continuation value of any equilibrium enters the interior of $\mathcal{E}(r)$ with probability 1.

7.2 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, it is sufficient to verify that the "jumps within the set" condition of equilibrium incentives holds at the boundary. Indeed, two events of independent Poisson processes happen at the same time with probability 0. Thus, if any event requires punishments/rewards with a continuation value in the interior of the set, players can use public randomization before another event occurs such that the continuation value after the randomization lies on the boundary again.

Despite the two important uses of public randomization in the derivation of our result, there are many instances in which the equilibrium payoff set $\mathcal{E}(r)$ with public randomization coincides with the equilibrium payoff set without public randomization. This is the case, for example, if information arrives continuously. Clearly, $\mathcal{E}(r)$ contains the equilibrium payoff set without public randomization. To see that the two sets coincide, it is thus sufficient to show that any payoff pair in $\mathcal{E}(r)$ can be attained by an equilibrium profile without public randomization. Indeed, in a game without abrupt information, the continuation value in (2) is continuous, hence any

the lengths of time periods are random, hence the same result need not apply. Studying our model for a continuum of possible types of events to verify whether a bang-bang property would hold is a rather involved exercise: with a continuum of Poisson processes, an iterative construction over arrival times does not work as there would be uncountably many jumps in the continuation value.

payoff pair in the interior can be attained locally by playing any enforceable action profile until the boundary is reached. Stationary payoff pairs on the boundary and payoff pairs on the boundary, where the curvature is strictly positive, can be attained locally by strategy profiles constructed in the proof of Lemma 6.3 and in Lemma 6.4, respectively. Payoff pairs on the boundary, where the curvature is 0, can be attained locally by using action profile and incentives form either end point of the straight line segment: the transfers are tangential and the drift points towards the interior at any point on the straight line segment as illustrated in the right panel of Figure 9. The continuation value will thus enter the interior of the equilibrium payoff set with certainty, hence a concatenation of these three local constructions will yield an enforceable strategy profile on $[0, \infty)$, which is an equilibrium profile by Lemma 4.3.

If players observe also abrupt information, the two payoff sets do not coincide in general. It is, however, fairly simple to verify whether they do coincide for a given equilibrium payoff set $\mathcal{E}(r)$ with public randomization. A sufficient condition for the two sets to coincide are:

(i) There exists at least one action profile a_0 that is enforceable without the use of abrupt information, and

(ii) the boundary of $\mathcal{E}(r)$ does not have any straight line segments outside the set of stationary payoffs.

If the two conditions are satisfied, then any payoff pair in the interior of $\mathcal{E}(r)$ can be attained locally by playing a_0 until the boundary is reached: by (i) no incentives through the abrupt information are required, hence the continuation value is continuous and does not jump past the boundary. By (ii), any extremal payoff pairs can be attained locally with stationary strategy profiles or by strategy profiles constructed in Lemma 6.4, which do not require public randomization. If the boundary of $\mathcal{E}(r)$ has straight line segments, however, the above argument may break down: it is possible that the equilibrium rewards/punishments at an end point of a straight line segment do not remain in $\mathcal{E}(r)$ if translated along the line segment. Nevertheless, in the partnership game of Section 3 the two sets do coincide for any value of γ: any payoff pair on the straight line segments can be attained by local play of the static Nash profile until the extremal payoff pairs w_A and w_B are reached; see, for example, Figure 5.

7.3 Relation to Sannikov and Skrzypacz [22]

Because of the similarity of our model with Sannikov and Skrzypacz [22], 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.[8] 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 and whether the induced sequence is decreasing.

This paper moves beyond the payoff bound and gives a way to compute $\mathcal{E}(r)$ for any value of the discount rate r. The precise shape is characterized by the tradeoff among the two types of information. It is also worth noting that the payoff bound M may be significantly larger than the equilibrium payoff set for discount rates that are bounded away from 0: in the partnership example of Section 3, our methods indicate that the equilibrium payoff set collapses below the negative diagonal D (as illustrated in Figure 3) if and only if $r \leq 15\gamma$, i.e., when the abrupt information is not sufficiently informative relative to the players' patience. The payoff bound M_γ of [22], however, extends beyond the negative diagonal D for any $\gamma > 0$. Finally, as we have already noted in Section 3, the techniques in this paper allow us to elicit the equilibirum 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 6.8 numerically. We illustrate the convergence of an iterated application of $\mathcal{B}_r$ to the equilibrium payoff set with the partnership game of Section 3. Finally, Section 8.3 provides an improvement on the algorithm in Proposition 5.3 for numerical implementation.

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

The general procedure for computing $\mathcal{B}_r(\mathcal{W})$ is the following: compute the set $\mathcal{S}_r(\mathcal{W})$ of stationary payoffs and then find the largest solution to the ODE (5) that contains $\mathcal{S}_r(\mathcal{W})$, where we check for non-stationary payoff pairs in $\mathcal{D}_r(\mathcal{W})$ "on the fly". Indeed, non-stationary payoff pairs in $\mathcal{D}_r(\mathcal{W})$ are precisely the payoff pairs where the expression in (5) is unbounded for some action profile a. Outside the set of stationary payoffs, the boundary of $\mathcal{B}_r(\mathcal{W})$ can thus be viewed as an extended solution to (5), which follows $\partial\mathcal{K}_{r,a}(\mathcal{W})$ if the expression in (5) is unbounded for some action profile a.

We begin by illustrating how to compute the set of stationary payoff pairs. A stationary payoff pair w can be written as $w = g(a) + \delta\lambda(a)$ for some $a \in \mathcal{A}$ and $\delta \in \Psi_a$, where we denote $\Psi_a = \Psi_a^1 \times \Psi_a^2$. The condition that the continuation value after an event y comes from the set $\mathcal{W}$ can thus be expressed as $g(a) + f_y(\delta) \in \mathcal{W}$, where

[8] This follows from the fact that $\mathcal{E}(r)$ is closed by Corollary 8.2. 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.

$f_y(\delta) = \delta\big(\lambda(a) + re_y\big)$. This eliminates the variable w and allows us to parametrize the set of stationary payoffs via incentives δ. Such incentives have to come from Ψ_a and they have to satisfy the jump condition for every $y \in Y$, i.e., the set of all such incentives is given by $\mathcal{X}_a(\mathcal{W}) := \Psi_a \cap \bigcap_{y\in Y} f_y^{-1}\big(\mathcal{W} - g(a)\big)$, 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

$$\mathcal{S}_r(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \big(g(a) + \mathcal{X}_a(\mathcal{W})\lambda(a)\big), \tag{6}$$

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 onto $\mathbb{R}^2$ in the direction $\lambda(a)$. Note that Ψ_a is a convex polytope characterized by the affine inequalities in (3). For a discretization $\mathcal{W}'$ of $\mathcal{W}$ with extremal points $x_1, \dots, x_n$ and corresponding normal vectors $N_1, \dots, N_n$, the set $\mathcal{X}_a(\mathcal{W}')$ is a convex polytope again, characterized by the affine inequalities in (3) and

$$N_j^\top\big(\delta(\lambda(a) + re_y)\big) \le N_j^\top x_j, \quad j = 1, \dots, n,\ y \in Y.$$

One can thus compute extremal points $z_1, \dots, z_n$ of the set $\mathcal{X}_a(\mathcal{W}')\lambda(a)$ by maximizing $N_j^\top\delta\lambda(a)$ for $j = 1, \dots, n$ over $\delta \in \mathcal{X}_a(\mathcal{W}')$. This is an efficient numerical procedure as it maximizes a linear function under a set of affine constrains. The computation is particularly simple if there is only one type y of events as then

$$\mathcal{X}_a(\mathcal{W})\lambda(a) = \left(\Psi_a \cap \frac{\mathcal{W} - g(a)}{\lambda(a) + r}\right)\lambda(a).$$

The set $g(a) + \mathcal{X}_a(\mathcal{W})\lambda(a)$ is the set of all stationary payoff pairs that are attainable locally by play of a. Equation (6) thus shows that $\mathcal{S}_r(\mathcal{W})$ consists of up to $|\mathcal{A}|$ disjoint components. Figures 9 and 10 show the two components of $\mathcal{S}_5\big(\mathcal{E}_{0.8}(5)\big)$ and $\mathcal{S}_5(\mathcal{V}^*)$ in the partnership game of Section 3 for $\gamma = 0.8$ and $\gamma = 0.4$, respectively.

It remains to find the largest solution to (5) that contains the set of stationary payoffs. In principle, this is achieved similarly as in Sannikov [20], but we additionally need to account for straight line segments and non-stationary payoff pairs in $\mathcal{D}_r(\mathcal{W})$. Since $\mathcal{B}_r(\mathcal{W})$ is convex, we parametrize the boundary via tangential angle θ. Let $w(\theta)$ denote the set of payoff pairs in $\mathcal{B}_r(\mathcal{W})$ with normal vector $N(\theta) = \big(\cos(\theta), \sin(\theta)\big)^\top$. Note that $w(\theta)$ is unique where the curvature of $\partial\mathcal{B}_r(\mathcal{W})$ is strictly positive, hence one can solve

$$\frac{\mathrm{d}w(\theta)}{\mathrm{d}\theta} = \frac{T(\theta)}{\kappa(\theta)} \tag{7}$$

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 maximization in (5) is unbounded for some action profile a, we check whether $w(\theta) \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ and $N(\theta) \in \mathcal{N}_{w(\theta)}\big(\mathcal{K}_{r,a}(\mathcal{W})\big)$. If this is

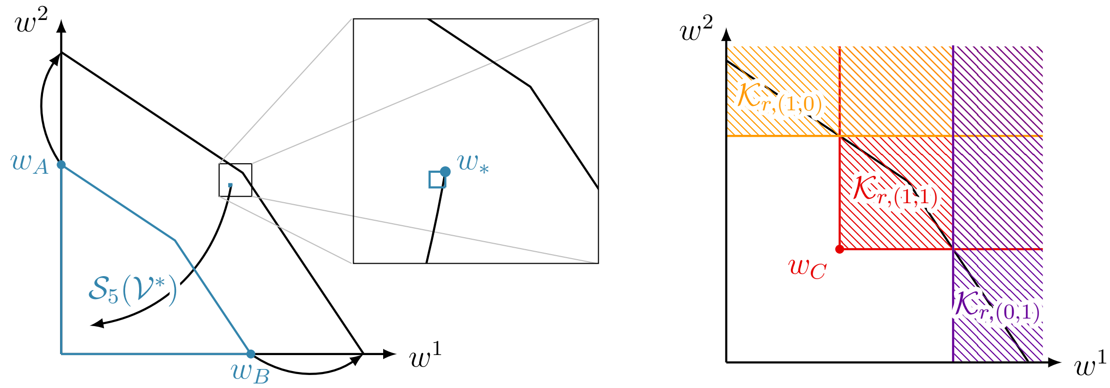

Figure 10: The left panel shows $\mathcal{S}_5(\mathcal{V}^*)$ in the partnership game of Section 3 for $\gamma = 0.4$. The smaller component is decomposed using the action profile of mutual effort by attaching mutual punishments to demand shocks. The larger component is decomposed using the static Nash profile by rewarding players when a demand shock occurs. Even though these incentives are not necessary to support the static Nash profile, they are allowed for the decomposition of payoffs. The right panel shows $\mathcal{K}_{5,a}(\mathcal{V}^*)$ for action profiles $a = (0, 1), (1, 0)$ and $(1, 1)$. A non-stationary corner $w \in \mathcal{B}_5(\mathcal{V}^*) \setminus \mathcal{S}_5(\mathcal{V}^*)$ is a corner of $\mathcal{K}_{5,a}(\mathcal{V}^*)$ by Proposition 6.7, where w is decomposable by a. The only such candidate w_C, however, is in the interior of $\operatorname{conv} \mathcal{S}_5(\mathcal{V}^*)$. Therefore w_C, cannot be a corner of $\mathcal{B}_5(\mathcal{V}^*)$.

the case, it is possible that the solution has reached a corner or a segment of positive length in $\mathcal{D}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$. Because $\mathcal{B}_r(\mathcal{W})$ is the largest solution to (5), we search for the maximal angle at corners and the longest segments in $\partial\mathcal{K}_{r,a}(\mathcal{W})$, respectively, for which a closed solution to (5) exists. If the maximization in (5) is taken over empty sets at some point $w(\theta)$, it is possible that $\partial\mathcal{B}_r(\mathcal{W})$ has a straight line segment outside of $\mathcal{D}_r(\mathcal{W})$. Again, we search for the longest straight line segments orthogonal to $N(\theta)$, for which a closed solution to (5) exists.

We illustrate this procedure by computing $\mathcal{B}_r(\mathcal{V}^*)$ in the partnership example of Section 3 for $\gamma = 0.4$. We first compute the set of stationary payoffs as illustrated in the left panel of Figure 10. 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 $\mathcal{B}_r(\mathcal{V}^*)$. Therefore, payoff pairs w_A and w_B in Figure 10 thus serve as starting points for solving (5). We thus search for the maximal angle at w_A, for which a solution connects to either w_B or w_*. The largest such solution yields $\partial\mathcal{B}_r(\mathcal{V}^*)$, which is depicted as the outermost curve in the left panel of Figure 11. If the intersection $\mathcal{S}_r(\mathcal{W}) \cap \partial\mathcal{V}^*$ was empty instead, starting points of symmetric games may be found by searching over the positive diagonal with initial angles $\pi/4$ or $5\pi/4$. If the game is asymmetric, one may use an iterative procedure as in Section 8 of Sannikov [20].

In this game, there are no corners outside the set of stationary payoffs as the solution to (2) is bounded everywhere. This is illustrated in the right panel of Figure 10, which shows $\mathcal{K}_{r,a}(\mathcal{V}^*)$ for the different action profiles $a \in \mathcal{A}$. By Proposition 6.7, the boundary of $\mathcal{B}_r(\mathcal{V}^*)$ may have corners outside the set of stationary payoffs only

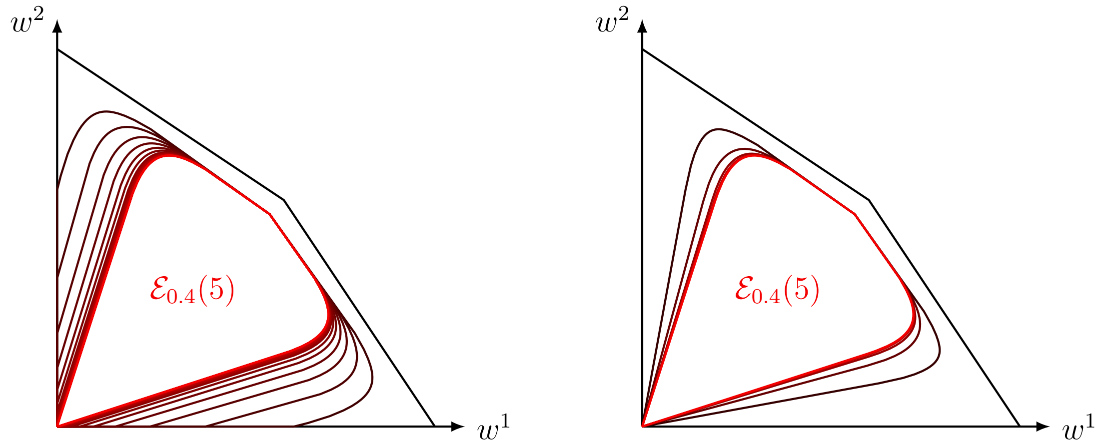

Figure 11: The left panel shows the convergence of the algorithm in Proposition 5.3. The right panel shows the convergence of the algorithm in Proposition 8.4, where outward jumps are excluded.

at corners of $\mathcal{K}_{r,a}(\mathcal{V}^*)$. The point w_C is the only such candidate, but it lies in the interior of the convex hull of stationary payoffs, hence also in the interior of $\mathcal{B}_r(\mathcal{V}^*)$. We conclude that all corners of $\mathcal{B}_r(\mathcal{V}^*)$ are stationary.

8.2 Computing the equilibrium payoff set

The equilibrium payoff set can be computed with the algorithm in Proposition 5.3. The sequence $(\mathcal{W}_n)_{n\geq 0}$ defined by $\mathcal{W}_n = \mathcal{B}_r(\mathcal{W}_{n-1})$ for any $n > 0$ starting at $\mathcal{W}_0 = \mathcal{V}^*$ is computed iteratively as described in Section 8.1. We stop the approxmation if the difference between two consecutive sets $\mathcal{W}_n$ and $\mathcal{W}_{n-1}$ is sufficiently small. The left panel of Figure 11 illustrates the convergence to $\mathcal{E}_{0.4}(5)$ in the partnership game. 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. Consider the payoff pair $w_A \in \mathcal{S}_5(\mathcal{V}^*)$ in Figure 10 as an example, which is decomposed using the static Nash profile and incentives that give player 2 his highest feasible and individually rational payoff upon the arrival of a demand shock. In the limit as $\mathcal{B}_r$ is applied infinitely often, incentives using outward jumps cannot be efficient to 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.

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

We begin by defining the set of payoff pairs that can be decomposed using inward jumps. Similarly to the definition of restricted-enforceability, such a definition should not depend on the a priori unknown set $\mathcal{B}_r(\mathcal{W})$. For any action profile $a \in \mathcal{A}$ and any direction $N \in S^1$, let

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

be the half-space of payoff pairs that can be decomposed with inward-jumps with respect to the direction N. Observe that $H_a(N)$ cannot extend past $g(a)$, i.e., $g(a)$ is either on the boundary or outside 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 immediately from the definition of $\mathcal{Q}_a$, stating that any payoff pair $w \in \mathcal{D}_r(\mathcal{W})$ that is decomposed by a is either contained in $\mathcal{Q}_a$ or requires outward jumps to be enforced.

Lemma 8.1. *Let $w \in \mathcal{D}_r(\mathcal{W})$ be decomposable by $a \in \mathcal{A}$. Then 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 $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$.*

While outward jumps may be necessary to decompose payoff pairs in $\mathcal{D}_r(\mathcal{W})$, for an arbitrary set $\mathcal{W}$, outward jumps cannot support equilibrium behavior because $\mathcal{E}(r)$ is convex. Therefore, we obtain the following fixed-point characterization of $\mathcal{E}(r)$ as a consequence to Proposition 5.3, Theorem 6.8, and Lemma 8.1.

Corollary 8.2. *Under Assumptions 1 and 2, $\mathcal{E}(r)$ is the largest closed subset of $\mathcal{V}^*$ such that $\mathcal{D}_r\big(\mathcal{E}(r)\big) \subseteq \bigcup_{a \in \mathcal{A}} \big(\mathcal{S}_r(\mathcal{E}(r)) \cup \partial\mathcal{K}_{r,a}\big(\mathcal{E}(r)\big)\big) \cap \mathcal{Q}_a$ and $\partial\mathcal{E}(r) \setminus \mathcal{D}_r\big(\mathcal{E}(r)\big)$ is continuously differentiable with curvature at almost every point w given by*

$$\kappa(w) = \max_{a \in \mathcal{A}} \max_{(\beta,\delta) \in \Xi_a(w, r, N_w, \mathcal{E}(r))} \frac{2{N_w}^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\beta\|^2}, \tag{8}$$

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

At a first glance, it may seem that the characterization has become more difficult with the exclusion of outward jumps. This is only a notational difficulty. The refinement allows us to exclude many points that would have to be considered as potential corners and extremal points in $\mathcal{D}_r\big(\mathcal{E}(r)\big)$. The following algorithm clarifies that we can excludes all of these points straight from the beginning, leading to a faster computation of $\mathcal{E}(r)$ through a refinement of the operator $\mathcal{B}_r$. Figure 11 illustrates that the sequence of approximating payoff sets converges much faster.

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

(i) its boundary is a solution to (5) at all points that are not decomposable, and

(ii) all points on the boundary that are decomposable by some $a \in \mathcal{A}$ are contained in the set $\big(\mathcal{S}_r(\mathcal{W}) \cup \partial\mathcal{K}_{r,a}(\mathcal{W})\big) \cap \mathcal{Q}_a$.

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 $\bigcap_{n\geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*

9 Conclusion

This paper characterizes the equilibrium payoff set for a class of continuous-time two-player games with imperfect public observation, where information may arrive both continuously through the observation of a noisy signal and discontinuously (abruptly) as the occurrences of infrequent but informative events. The presence of abrupt information allows the use of equilibrium incentives through value burning. This additional way of providing incentives has a drastic impact on the set of equilibrium strategies and the attainable payoff pairs. Stationary payoff pairs can be attained locally through strategies that disregard the continuous stream of information completely. At these payoff pairs, the equilibrium payoff set may have corners and straight line segments outside the set of static Nash payoffs, which is precluded in models without abrupt information. The characterization of the equilibrium payoff set is new even when the signal is continuous but fails the widely assumed pairwise identifiability condition. Our methods thus allow the computation of the equilibrium payoff set in games where the signal is one-dimensional such as a partnership game or a duopoly in a single homogeneous good.

The hallmark of continuous-time repeated games with imperfect monitoring is the ability to describe the equilibrium payoff set via an ODE of its boundary that involves the unique equilibrium incentives at each point. In models with abrupt information, any such description is self-referential because the amount of value that can be burnt at each point depends on the equilibrium payoff set itself. We show that the equilibrium payoff set can be computed with an iterative procedure over the arrival times of abrupt information. In each step of the algorithm, one computes the largest payoff set that is relaxed self-generating with respect to the set from the previous step. Doing so eliminates any self-referentiality and hence each set of the algorithm is described by an explicit ODE. The notion of largest relaxed self-generating payoff set is the continuous-time analogue to the standard set-operator in Abreu, Pearce and Stacchetti [3]. Our iterative procedure thus resembles the algorithm in [3]. However, unlike its discrete-time analogue, the set in each step is computed efficiently as the numerical solution to an ODE.

The concept of relaxed self-generating payoff sets is an important methodological contribution as it will be useful in subsequent research on continuous-time games that involve discontinuous information, such as stochastic games with a finite state space. Indeed, finite-state Markov processes in continuous time are precisely described by a set of Poisson processes, whose intensities govern the rate at which the underlying state changes. The study of continuous-time stochastic games may provide new insights that are not available in discrete time. In their discrete-time treatment of stochastic games, Hörner et al. [13] assume that the underlying Markov process is irreducible so that the distribution over states converges to the stationary distribution. As a result, the payoff bound as players get arbitrarily patient becomes independent of the initial state, making the analysis tractable. In contrast, the techniques in this paper do not require that players become arbitrarily patient. It seems, therefore, plausible that one could describe the correspondence $x \mapsto \mathcal{E}_x(r)$ for initial states x as a set of coupled differential equations. This would allow us to investigate richer sets of questions that involve the equilibrium payoff impact of the underlying state such as, for example, contestable democracies, where players can pay to influence the current state. Moreover, if convergence to a stationary distribution is not required to obtain the result, the analysis would not be limited to games in which the underlying Markov process is irreducible, thereby extending the analysis to a wider class of stochastic games than is currently possible in a discrete-time setting.

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.

A Proofs of auxiliary results in the main text

A.1 Dynamics of the continuation value and continuation promises

For the proofs of Lemmas 4.1 and 4.3, we draw from the arguments in Bernard and Frei [7]. In the interest of space we present in this subsection only the additional arguments required due to abrupt information.

Proof of Lemma 4.1. This proof mirrors the proof of Lemma 1 in Bernard and Frei [7], with the following additional arguments for the jump processes. 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 (Theorem IV.36 in Protter [19]). Therefore, we obtain the unique martingale representation property for a square-integrable martingale by Corollary 1 to Theorem IV.37 in [19]. That is, for a bounded $\mathcal{F}_T$-measurable random variable w^i_T, there exists an $\mathcal{F}_0$-measurable c^i_T, predictable processes $\big(\beta^i_{t,T}\big)_{0\leq t\leq T}$, $\big(\delta^i_{t,T}(y)\big)_{0\leq t\leq T}$ for all $y\in Y$ with $\mathbb{E}_{Q^A_T}\Big[\int_0^T \beta^i_{t,T}\,\mathrm{d}t\Big]<\infty$ 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

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

The remainder of the equivalence between (i) and (ii) follows along the same lines as the proof in Bernard and Frei [7], requiring only one additional argument to ensure that $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\delta^i_s(y)\,\big(\mathrm{d}J^y_s-\lambda(y|\tilde{A}_s)\,\mathrm{d}s\big)$ is a martingale under $Q^{\tilde{A}}$. Indeed, since $\int r\delta^i(y)(\mathrm{d}J^y_t-\mathrm{d}t)$ has bounded jumps by construction for any $y\in Y$, it follows that $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\delta^i_s(y)\,\big(\mathrm{d}J^y_s-\lambda(y|A_s)\,\mathrm{d}s\big)$ is a martingale with bounded mean oscillation (BMO) under Q^A_u up to any time $u\in(t,\infty)$. Assumption 1 implies that the jumps of $\big(\lambda(y|A_s)-1\big)\Delta J^y_s$ in Footnote 2 are bounded from below by $-1+\varepsilon$ for some $\varepsilon>0$ and any $y\in Y$. Therefore, Remark 3.3 and Theorem 3.6 in Kazamaki [16] imply that $\int_t^\cdot r\mathrm{e}^{-r(s-t)}\delta^i_s(y)\,\big(\mathrm{d}J^y_s-\lambda(y|\tilde{A}_s)\,\mathrm{d}s\big)$ is a *BMO*-martingale under $Q^{\tilde{A}}_u$. □

Proof of Lemma 4.3. This proof is analogous to the proof of the second statement of Lemma 1 in Bernard and Frei [7]. □

A.2 Convergence of the algorithm

In this appendix we prove the convergence of the algorithm in Proposition 5.3 to $\mathcal{E}(r)$. We first show that $\mathcal{B}_r(\mathcal{W})$ is monotone in $\mathcal{W}$.

Lemma A.1. *Let $\mathcal{W}\subseteq\mathcal{W}'$. Then $\mathcal{B}_r(\mathcal{W})\subseteq\mathcal{B}_r(\mathcal{W}')$.*

Because $\mathcal{B}_r(\mathcal{W})$ is defined through payoff pairs being attainable by a solution to the stochastic differential equation (2), it is necessary to discuss the solution concept used. Note that the SDE (2) does not admit strong solutions in general, that is, there might not exist processes (W,A,β,δ,M) that solve (2) for a fixed Brownian motion Z and fixed Poisson processes $(J^y)_{y\in Y}$. However, the SDE (2) does admit weak solutions, where Z and $(J^y)_{y\in Y}$ and in fact the entire filtered probability space $(\Omega,\mathcal{F},\mathbb{F},P)$ on which the processes are defined are part of the solution.

Definition A.2. We call $(\Omega,\mathcal{F},\mathbb{F},P,Z,(J^y)_{y\in Y})$ a *stochastic framework* if $(\Omega,\mathcal{F},P)$ is a probability space containing the filtration $\mathbb{F}$, with respect to which Z is a Brownian motion and $(J^y)_{y\in Y}$ are Poisson processes that are pairwise independent and independent of Z.

Proof of Lemma A.1. For any $w \in \mathcal{B}_r(\mathcal{W})$, there exists a solution (W, A, β, δ, M) to (2) in a stochastic framework such that (β, δ) enforces A, $W_0 = w$ a.s., $W_\tau \in \mathcal{X}$ a.s. for every stopping time $\tau < \sigma_1$, and $W_{\sigma_1} \in \mathcal{W}$ a.s. It follows that $W_{\sigma_1} \in \mathcal{W}'$ a.s., hence $\mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}'$-relaxed self-generating, hence contained in $\mathcal{B}_r(\mathcal{W}')$ by maximality. □

In the proof of Lemma A.1, it was sufficient to attain a fixed payoff pair with an enforceable solution to (2). For the piecewise construction of equilibrium strategy profiles it will be necessary to concatenate solutions at stopping times, at which point the continuation value is random. Because the filtration, in which the continuation value and the stopping times live, is part of the solution to (2), some care is needed at these concatenation points. We will employ the following lemma for concatenations.

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

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

(ii) There exist a strategy profile A, square-integrable predictable processes β and δ, a bounded semimartingale W, and a martingale M strongly orthogonal to Z and $(J^y)_{y \in Y}$ such that $W, A, \beta, \delta, M, Z$ and $(J^y)_{y \in Y}$ satisfy (2), (β, δ) enforces A, $W_0 = X$ P-a.s., $W_\tau \in \mathcal{B}_r(\mathcal{W})$ P-a.s. for every $\mathbb{F}$-stopping time $\tau < \sigma_1$, and $W_{\sigma_1} \in \mathcal{W}$ P-a.s. for every $y \in Y$.

Proof. The proof works similarly to the proof of Lemma 8 in Bernard and Frei [7] with the additional observation that the path space of cadlag functions admits a metric that makes it complete and separable by Theorem A.2.2 in Kallenberg [15]. Therefore, it is possible to choose the path space of the public information in a complete and separable way even in the presence of abrupt information. This guarantees the existence of regular conditional probabilities, which are required to establish the equivalence as in the proof of Lemma 8 in Bernard and Frei [7]. □

We are now ready to prove Lemma 5.2 and Propositions 5.3 and 8.4. The proof of Lemma 5.2 will show how Lemma A.3 is applied to concatenate strategy profiles.

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. To that end, fix a payoff pair $w \in \mathcal{B}_r(\mathcal{W})$ arbitrarily and let it be attained by 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 (β, δ) enforcing A, $W_0 = w$ P-a.s., $W \in \mathcal{B}_r(\mathcal{W})$ on $[\![0, \sigma_1)\!)$, and $W_{\sigma_1} \in \mathcal{W}$ P-a.s.[9] Define the processes

$$\tilde{Z} = Z_{\cdot + \sigma_1} - Z_{\sigma_1}, \qquad \tilde{J}^y = J^y_{\cdot + \sigma_1} - J^y_{\sigma_1}, \text{ for every } y \in Y$$

[9]For two stopping times σ and τ, the set $[\![\sigma, \tau)\!) := \{(\omega, t) \in \Omega \times [0, \infty) \mid \sigma(\omega) \leq t < \tau(\omega)\}$ 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.

and the filtration $\tilde{\mathbb{F}} = \left(\tilde{F}_t\right)_{t\geq 0}$ defined by $\tilde{\mathcal{F}}_t := \mathcal{F}_{t+\sigma_1}$. Because Brownian motion and Poisson processes have independent and identically distributed increments, $\tilde{Z}$ is an $\tilde{\mathbb{F}}$-Brownian motion and $\tilde{J}^y$ is an $\tilde{\mathbb{F}}$-Poisson process for any $Y \in Y$. Therefore, $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde{Z}, (\tilde{J}^y)_{y\in Y})$ is a stochastic framework with $W_{\sigma_1} \in \tilde{\mathcal{F}}_0$. Since $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$ by assumption, it follows that $W_{\sigma_1} \in \mathcal{B}_r(\mathcal{W})$ and hence the equivalence in Lemma A.3 implies the existence of processes $\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}$, and $\tilde{M}$ solving (2) on the stochastic framework $(\Omega, \mathcal{F}, \tilde{\mathbb{F}}, P, \tilde{Z}, (\tilde{J}^y)_{y\in Y})$ such that $(\tilde{\beta}, \tilde{\delta})$ enforces $\tilde{A}$, $\tilde{W}_0 = W_{\sigma_1}$ P-a.s., $\tilde{W} \in \mathcal{B}_r(\mathcal{W})$ on $[\![0, \tilde{\sigma}_1)\!)$, and $\tilde{W}_{\tilde{\sigma}_1} \in \mathcal{W}$ P-a.s., where $\tilde{\sigma}_n$ refers to the n^{th} time any of the processes $\left(\tilde{J}^y\right)_{y\in Y}$ jump. These two solutions are concatenated by setting

$$\hat{W} = W 1_{[\![0,\sigma_1)\!)} + \tilde{W}_{\cdot - \sigma_1} 1_{[\![\sigma_1,\infty)\!)}$$

and similarly for $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, and $\hat{M}$. Observe that $(\hat{\beta}, \hat{\delta})$ enforces $\hat{A}$ and $\hat{W}_0 = w$ P-a.s. The concatenations are solutions to (2) in the "concatenated" stochastic framework defined by

$$\hat{Z} = Z 1_{[\![0,\sigma_1)\!)} + (\tilde{Z} + Z_{\sigma_1}) 1_{[\![\sigma_1,\infty)\!)}, \qquad \hat{J}^y = J^y 1_{[\![0,\sigma_1)\!)} + (\tilde{J} + J^y_{\sigma_1}) 1_{[\![\sigma_1,\infty)\!)}.$$

Note, however, that $\hat{Z} = Z$ and $\hat{J}^y - J^y$ for all events $y \in Y$, that is, the concatenated stochastic framework is identical to the original framework. Therefore, $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ is a solution to (2) in $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$. Moreover, $\tilde{\sigma}_n = \sigma_{n+1}$ and hence $\hat{W}$ is contained in $\mathcal{B}_r(\mathcal{W})$ on $[\![0, \sigma_2)\!)$ and $\hat{W}_{\sigma_2} = \tilde{W}_{\tilde{\sigma}_1} \in \mathcal{W}$ a.s. We have thus constructed an enforceable solution to (2) that remains in $\mathcal{B}_r(\mathcal{W})$ up until the arrival of the second event, such that at the time of the second event, the continuation payoff comes from $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$. 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, observe that self-generation implies that for any $w \in \mathcal{W}$, there exists an enforceable strategy profile with continuations that remain in $\mathcal{W}$. In particular, $W_\sigma \in \mathcal{W}$ a.s. It follows that $\mathcal{W} \subseteq \mathcal{B}_r(\mathcal{W})$ by maximality of $\mathcal{B}_r(\mathcal{W})$. □

Proof of Proposition 5.3. Lemma 5.2 implies that $\mathcal{E}(r) \subseteq \mathcal{B}\big(\mathcal{E}(r)\big)$ and that $\mathcal{B}\big(\mathcal{E}(r)\big)$ is self-generating. Since $\mathcal{E}(r)$ is the largest bounded self-generating set, it follows that $\mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$. Let now $\mathcal{W}_0 = \mathcal{V}^*$. Since any payoff in $\mathcal{B}_r(\mathcal{V}^*)$ can be attained by a locally enforceable strategy profile with a continuation payoff in $\mathcal{V}^*$, it follows that every player i's continuation value is at least as large as his minmax payoff at any point in time. This implies that $\mathcal{W}_1 \subseteq \mathcal{W}_0$. Monotonicity of $\mathcal{B}_r$ implies that $\mathcal{E}(r) \subseteq \mathcal{W}_1 \subseteq \mathcal{W}_0$. An iterated application of Lemma 5.2 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 5.2. □

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 $\mathcal{W}_\infty$.

Corollary 8.2 asserts that $\tilde{\mathcal{B}}_r\big(\mathcal{E}(r)\big) = \mathcal{B}_r\big(\mathcal{E}(r)\big) = \mathcal{E}(r)$. 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 is locally Lipschitz continuous at almost every point, so that locally, it admits a unique solution. We show in Lemma C.4 that $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is C^1, hence $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is the unique C^1 solution to the optimality equation. For any fixed $r > 0$, $a \in \mathcal{A}$, and closed and convex $\mathcal{W} \subseteq \mathcal{V}$, consider the optimality equation in the following form:

$$\kappa_a(w, N) = \max_{(\phi,\delta)\in\Xi_a(w,N,r,\mathcal{W})} \frac{2N^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\phi\|^2}. \tag{9}$$

We start by reducing the two-variable optimization problem to a one-variable optimization by expressing the control ϕ in terms of δ. 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. For players $i = 1, 2$, define

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

for any direction N and any $\delta^i \in \mathbb{R}^{|Y|}$. Because $\mathcal{I}^i_a(N, \delta^i)$ is the intersection of closed half-spaces, it is a (possibly unbounded or empty) closed convex polytope. Therefore, so is $\Phi_a(N, \delta) := \mathcal{I}^1_a(N, \delta^1) \cap \mathcal{I}^2_a(N, \delta^2)$, the set of all vectors $\phi \in \mathbb{R}^d$ such that $(T\phi, \delta)$ enforces a. Let $\phi(a, N, \delta)$ denote the vector of smallest length in $\Phi_a(N, \delta)$.

Lemma B.1. *Fix $a \in \mathcal{A}$. Then $(N, \delta) \mapsto \phi(a, N, \delta)$ is locally Lipschitz continuous where $\Phi_a(N, \delta) \neq \emptyset$ and N is not a coordinate direction, i.e., $N \notin \{\pm e_1, \pm e_2\}$.*

In an intermediate step, we will show that the set-valued map $(N, \delta) \mapsto \Phi_a(N, \delta)$ is locally Lipschitz continuous for N different from coordinate directions. We refer to Aubin and Frankowska [6] for a detailed overview of set-valued maps and their properties and state here only the most central property.

Definition B.2. 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.

Proof of Lemma B.1. Let $\mathcal{I}^i_a(\delta^i) := \{\beta \in \mathbb{R}^d \mid (\beta, \delta^i) \text{ satisfies (3) for player } i\}$ be the solution set to (3) for player i and observe that it is a closed convex polytope. Its hyperfaces have normal vectors $\Delta\mu^i_{j_i} := \mu(a) - \mu(a^i_{j_i}, a^{-i})$, where $a^i_1, \ldots, a^i_{m_i}$ is an enumeration of $\mathcal{A}^i \setminus \{a^i\}$. The parameter δ^i determines the location of these hyperfaces. Observe that a change from δ^i to $\tilde{\delta}^i$ shifts face j_i by $(\tilde{\delta}^i - \delta^i)\Delta\lambda^i_{j_i}$, where $\Delta\lambda^i_{j_i} := \lambda(a) - \lambda(a^i_{j_i}, a^{-i})$. Therefore, the triangle inequality implies that

$$\mathcal{I}^i_a(\delta^i) \subseteq \mathcal{I}^i_a(\tilde{\delta}^i) + B_1(0) \sum_{j_i=1,\ldots,m_i} \|\Delta\lambda^i_{j_i}\| \, \|\tilde{\delta}^i - \delta^i\|,$$

i.e., $\mathcal{I}^i_a(\delta^i)$ is Lipschitz continuous in δ^i. It is clear that $\mathcal{I}^i_a(N, \delta^i) = \frac{1}{T^i}\mathcal{I}^i_a(\delta^i)$ for $i = 1, 2$ is locally Lipschitz continuous in (N, δ^i) for N different from coordinate directions. To conclude that also the intersection $\Phi_a(N, \delta) = \mathcal{I}^1_a(N, \delta^1) \cap \mathcal{I}^2_a(N, \delta^2)$ is locally Lipschitz continuous, we verify the conditions of the technical Lemma E.2. Because the scaling of $\mathcal{I}^i_a(\delta^i)$ does not affect the direction of its hyperfaces, the normal vectors of the hyperfaces of $\mathcal{I}^i_a(N, \delta^i)$ are constant in (N, δ^i). Thus, Lemma E.2 applies, which establishes that $(N, \delta) \mapsto \Phi_a(N, \delta)$ is locally Lipschitz continuous for non-coordinate directions N. The statement now follows from the following lemma. □

Lemma B.3. *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

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

Lemma B.1 significantly simplifies the constraints in the maximization in (9) because we are left with 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. Instead of requiring $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and some fixed set $\mathcal{W}$, we will require that $\delta \in D(w)$ for an affine, compact- and convex-valued correspondence $w \mapsto D(w) \subseteq \mathbb{R}^{2\times m}$. We study the optimality equation of the form

$$\kappa_a(w, N) = \max_{\delta \in \Psi_a(w,N,r,D)} \frac{2N^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\phi(a, N, \delta)\|^2}, \tag{10}$$

where $\Psi_a(w, N, r, D) := \big\{\delta \in D(w) \;\big|\; \Phi_a(N, \delta) \neq \emptyset \text{ and } N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0\big\}$. Observe that the ODE (10) reduces to (9) for $D(w) = (\mathcal{W} - w)^m/r$, where we denote by $\mathcal{X}^m$ the m-fold product of a set $\mathcal{X} \subseteq \mathbb{R}^2$, i.e.,

$$\mathcal{X}^m := \big\{x \in \mathbb{R}^{2\times m} \;\big|\; (x_k^1, x_k^2)^\top \in \mathcal{X} \text{ for every } k = 1, \ldots, m\big\}.$$

Lemma B.4. *Let $w \mapsto D(w)$ be affine, compact- and convex-valued. Then for any $a \in \mathcal{A}$, the map $(w, N) \mapsto \Psi_a(w, N, r, D)$ is compact- and convex-valued. Moreover, it is locally Lipschitz continuous for N different from coordinate directions.*

Proof. Identify $\mathbb{R}^{2\times|Y|}$ with $\mathbb{R}^{2|Y|}$ by setting $\delta \approx (\delta^1, \delta^2)$. Let $\Psi_a(w, N)$ and $\mathcal{J}_a(N)$ denote the sets of all δ, for which $N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$ and $\Phi_a(N, \delta) \neq \emptyset$, respectively, are satisfied. We begin by showing that $\mathcal{J}_a(N)$ is closed and convex, hence so is $\Psi_a(w, N, r, D) = \mathcal{J}_a(N) \cap \Psi_a(w, N) \cap D(w)$ as intersection of such sets. Indeed, let $\delta_1, \delta_2 \in \mathcal{J}_a(N)$. Then there exist ϕ_1, ϕ_2 such that $(\delta_j, T\phi_j)$ for $j = 1, 2$ satisfy (3) for every $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$ and $i = 1, 2$. By linearity of (3), so does $(\delta_\nu, T\phi_\nu)$ for $\nu \in [0, 1]$, where we set $\delta_\nu := \nu\delta_1 + (1 - \nu)\delta_2$ and $\phi_\nu := \nu\phi_1 + (1 - \nu)\phi_2$. This shows that $\mathcal{J}_a(N)$ is convex. Let $(\delta_n)_{n\geq 0}$ be a sequence in $\mathcal{J}_a(N)$. Then there exists $(\phi_n)_{n\geq 0}$ such that $(\delta_n, T\phi_n)$ satisfies (3). Since the inequalities in (3) are not strict, $(\lim_{n\to\infty} \delta_n, T \lim_{n\to\infty} \phi_n)$ satsfies (3), hence $\mathcal{J}_a(N)$ is closed. Compactness of $\Psi_a(w, N, r, D)$ now follows because $D(w)$ is compact.

It remains to show (local) Lipschitz continuity. Observe that $w \mapsto \Psi_a(w, N)$ and $w \mapsto \mathcal{J}_a(N)$ are affine functions. Lipschitz continuity of $w \mapsto \Psi_a(w, N, r, D)$ thus follows from Lemma E.1. For local Lipschitz continuity in N, we are going to use the fact that arbitrary unions of Lipschitz-continuous functions with uniformly bounded Lipschitz constant are again Lipschitz continuous. To that end, introduce the auxiliary sets $\mathcal{J}_a(N, \phi) := \big\{\delta \in \mathbb{R}^{2|Y|} \;\big|\; (T\phi, \delta) \text{ enforces } a\big\}$ for $\phi \in \mathbb{R}^d$. For $i = 1, 2$, let $a_1^i, \ldots, a_{m_i}^i$ be an enumeration of $\mathcal{A}^i \setminus \{a^i\}$ and abbreviate $\Delta\mu_{j_i}^i := \mu(a) - \mu(a_{j_i}^i, a^{-i})$ and $\Delta\lambda_{j_i}^i := \lambda(a) - \lambda(a_{j_i}^i, a^{-i})$ as in the proof of Lemma B.1. Then $\mathcal{J}_a(N, \phi)$ is a closed convex polytope, whose hyperfaces have normal vectors

$$\begin{pmatrix} \Delta\lambda_{j_1}^1 \\ 0 \end{pmatrix}, \quad j_1 = 1, \ldots, m_1, \qquad \begin{pmatrix} 0 \\ \Delta\lambda_{j_2}^2 \end{pmatrix}, \quad j_2 = 1, \ldots, m_2. \tag{11}$$

Because N only determines the location of these hyperfaces $N \mapsto \mathcal{J}_a(N, \phi)$ is Lipschitz continuous similarly as in the proof of Lemma B.1, with a Lipschitz constant that depends only on the vectors $\Delta\mu^i_{j_i}$. In particular, the Lipschitz constant of $N \mapsto \mathcal{J}_a(N,\phi)$ is uniformly bounded in ϕ.

Let $\bar{\mathcal{W}}$ be a bounded polytope containing $D(w)$ such that none of its normal vectors are arbitrarily close to being linearly dependent to any $2|Y|-1$ normal vectors of $\Psi_a(w,N)$ or any of the vectors in (11). Then $\mathcal{J}_a(N,\phi)\cap\bar{W}$ and $\Psi_a(w,N)\cap\bar{\mathcal{W}}$ are convex bounded polytopes. Let $\mathcal{X}(N)$ be any subset of normal vectors to $\Psi_a(w,N)$ and $\mathcal{J}_a(N,\phi)$. If there exists a linear combination amongst the vectors in $\mathcal{X}(N)$, then there exists a linear combination also in $\mathcal{X}\big(\tilde{N}\big)$ for $\tilde{N}$ arbitrarily close to N by multiplying the coefficients by $\tilde{N}^i/N^i$, respectively. Therefore, Lemma E.2 applies and shows that $N \mapsto \Psi_a(w,N)\cap\mathcal{J}_a(N,\phi)\cap\bar{\mathcal{W}}$ is locally Lipschitz continuous in N for non-coordinate directions N. Since $D(w)\subseteq\bar{\mathcal{W}}$ is constant in N and the intersection of a Lipschitz continuous map with a convex and compact set is Lipschitz continuous, it follows that for any $\phi\in\mathbb{R}^d$, $N\mapsto\Psi_a(w,N)\cap\mathcal{J}_a(N,\phi)\cap D(w)$ is Lipschitz continuous. Local Lipschitz continuity of $N\mapsto\Psi_a(w,N,r,D)$ now follows from the fact that the arbitrary union of Lipschitz continuous maps with uniformly bounded Lipschitz constants is Lipschitz again. □

So far we have shown that (10) is locally Lipschitz continuous for almost every direction N, where $\phi(a,N,\delta)$ is well defined and bounded away from 0. Define

$$E_a(r,D) := \big\{(w,N)\in\mathbb{R}^2\times S^1 \;\big|\; \Psi_a(w,N,r,D)\neq\emptyset\big\}$$
$$\Gamma_a(r,D) := \big\{(w,N)\in\mathbb{R}^2\times S^1 \;\big|\; \exists\delta\in\Psi_a(w,N,r,D) \text{ with } \phi(a,N,\delta)=0\big\}$$

and $\Gamma(r,D) := \bigcup_{a\in\mathcal{A}}\Gamma_a(r,D)$. Denote by $\mathcal{P} := \mathbb{R}^2\times\{\pm e_1,\pm e_2\}$ the set of points $(w,N)\in\mathbb{R}^2\times S^1$ with a coordinate normal vector N.

Lemma B.5. *Let D be an affine, compact- and convex-valued correspondence. If a sequence $(w_n,N_n)_{n\geq0}$ converges to $(w,N)\notin\mathcal{P}$ such that $\Psi_a(w_n,N_n,r,D)\neq\emptyset$ for all $n\geq0$, then $\Psi_a(w,N,r,D)\neq\emptyset$.*

Proof. Let $\delta_n\in\Psi_a(w_n,N_n,r,D)$. Because $D(w_n)$ is uniformly bounded by $D(\mathcal{V})$, the sequence $(\delta_n)_{n\geq0}$ is uniformly bounded as well. Therefore, $(\delta_n)_{n\geq0}$ converges along a subsequence $(n_k)_{k\geq0}$ to some finite limit δ with $N^\top\big(g(a)+\delta\lambda(a)-w\big)\geq0$. Since D is closed-valued and Lipschitz continuous, $\delta(y)\in D(w)$ for every $y\in Y$. It remains to show that $\Phi_a(N,\delta)\neq\emptyset$. Suppose towards a contradiction that the converse is true. Then closedness of $\mathcal{I}^i_a(N,\delta^i)$ for $i=1,2$ implies that $\mathcal{I}^1_a(N,\delta^1)$ and $\mathcal{I}^2_a(N,\delta^2)$ are strictly separated. By continuity, $\mathcal{I}^1_a(N_{\ell_k},\delta^1_{\ell_k})$ and $\mathcal{I}^2_a(N_{\ell_k},\delta^2_{\ell_k})$ are separated as well for k sufficiently large, a contradiction. □

Corollary B.6. *For any $a\in\mathcal{A}$ and $\varepsilon\geq0$, $E_a(r,D)\cup\mathcal{P}$ and $\Gamma_a(r,D)$ are closed. Therefore, so is $\Gamma(r,D)$.*

Proof. Closedness of $E_a(r,D) \cup \mathcal{P}$ is established in Lemma B.5. The proof that $\Gamma_a(r,D)$ is closed follows along the same lines with the additional observation that $0 \in \Phi_a(N,\delta)$ for some $N \in S^1$ if and only if $0 \in \Phi_a(N,\delta)$ for all $N \in S^1$. Finally, $\Gamma(r,D)$ is closed as finite union of closed sets. $\square$

Proposition B.7. *Suppose that $\mathcal{W}$ has non-empty interior and that Assumption 2 is satisfied. For any affine, compact- and convex-valued correspondence D,*

$$\kappa(w,N) = \max_{a\in\mathcal{A}} \max_{\delta\in\Psi_a(w,N,r,D)} \frac{2N^\top\big(g(a)+\delta\lambda(a)-w\big)}{r\|\phi(a,N,\delta)\|^2} \tag{12}$$

is locally Lipschitz continuous outside of $\Gamma(r,D)$, except where (w,N) leaves or enters $E_a(r,D)$ of the maximizing action profile a. Here, we interpret $\kappa(w,N) = 0$ on $\bigcap_{a\in\mathcal{A}} E_a(r,D)^c$, i.e., where the maxima are taken over empty sets.

When we refer to a solution to (12), we will always mention explicitly with respect to which map D (12) is being solved.

Proof. Suppose first that N is not a coordinate direction, that is, $(w,N) \in E_a(r,D) \setminus \big(\Gamma_a(r,D) \cup \mathcal{P}\big)$. We first show local Lipschitz continuity of κ_a in (10) for fixed $a \in \mathcal{A}$. Since $\Gamma_a(r,D)$ is closed by Corollary B.6, there exists an open neighbourhood U of (w,N) bounded away from $\Gamma_a(r,D)\cup\mathcal{P}$. Therefore, $\inf_{N,\delta} \|\phi(a,N,\delta)\| \geq c$ and hence the function that is maximized in the right hand side of (10) is Lipschitz continuous on U by Lemma B.1. It follows that κ_a is Lipschitz continuous by Lemmas B.3 and B.4. Because (12) is the maximum over finitely many functions κ_a, it is Lipschitz continuous except where (w,N) leaves the domain of the maximal function κ_a.

Suppose now that N is a coordinate direction and without loss of generality, let $N \in \{\pm e_1\}$. Let a denote the maximizing action profile at (w,N). Because we show Lipschitz continuity only where maximizers in (12) do not change, we may assume that a maximizes (12) in a neighborhood of (w,N). If κ_a is identically zero in a neighborhood, the statement is trivial, hence suppose that a solution is strictly curved. Let $\mathcal{I}^2_a(N',\delta^2)$ be defined as in the proof of Lemma B.1, and let $\phi^2(a,N',\delta^2)$ denote the shortest vector in $\mathcal{I}^2_a(N',\delta^2)$. Note that $\mathcal{I}^2_a(N',\delta^2)$ and $\phi^2(a,N',\delta^2)$ are Lipschitz continuous in a neighborhood of N. It follows with an identical argument as in the proof of Lemma B.4 that $\Psi^2_a(w',N',r,D) := \Psi^2_a(w',N',r,D) \cap \Psi^1_a \times \mathbb{R}^d$ is locally Lipschitz continuous in (w',N') in a neighborhood of (w,N). Moreover, $\Psi^2_a(w,N,r,D)$ has non-empty interior by assumption, hence $\Psi^2_a(w',N',r,D)$ is non-empty for (w',N') in a neighborhood of (w,N) by Lipschitz continuity. Therefore,

$$\tilde{\kappa}_a(w',N') := \max_{\delta\in\Psi^2_a(w',N',r,D)} \frac{2N'^\top\big(g(a)+\delta\lambda(a)-w'\big)}{r\big(\phi^2(a,N',\delta^2)\big)^2}$$

is Lipschitz continuous in a neighborhood of (w, N) as well. Note that $\tilde{\kappa}_a \le \kappa_a$ and $\kappa_a(w, N) = \tilde{\kappa}_a(w, N)$, hence it follows that

$$\kappa_a(w, N) - \kappa_a(w', N') \le \tilde{\kappa}_a(w, N) - \tilde{\kappa}_a(w', N') \le K\|w - w'\| + K\|N - N'\|,$$

i.e., κ_a is locally Lipschitz continuous as well. □

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

We start by showing that $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is given by the optimality equation. Because the continuous part of the signal is what creates the curvature, these steps are similar in ideas to Sannikov [20]. Some technical bounds on the provision of incentives and proximity of solutions to (12) for different choices of D are deferred to Appendix E.2. We begin with the proof of Lemma 6.4.

Proof of Lemma 6.4. We prove the more general result, where we require $\delta \in D(w)$ for an affine, compact- and convex-valued correspondence $w \mapsto D(w) \subseteq \mathbb{R}^{2\times m}$ instead of $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$. 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 $B_\eta(w)$, $\partial\mathcal{W}$ 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, where T_w is the vector obtained by rotating N_w by 90° in clockwise direction. Denote by $\pi(v) = \big(\hat{v}, f(\hat{v})\big)$ the projection of $v \in B_\eta(w)$ onto $\partial\mathcal{W}$ 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) with initial condition $W_0 = w$ such that $M \equiv 0$ and for all $t \ge 0$, $A_t = a^*\big(\pi(W_t)\big)$, $\delta_t = \delta^*\big(\pi(W_{t-})\big)$, and $\beta_t = T_t\phi^*\big(\pi(W_t)\big)$ on $[\![0, \tau)\!)$, where we abbreviated $N_t = N_{\pi(W_t)}$ and $T_t = T_{\pi(W_t)}$, and define $\tau := \sigma_1 \wedge \inf\{t \ge 0 \mid W_t \notin B_\eta(w)\}$, where σ_1 indicates the first time any infrequent event occurs. Since $\delta \in \Psi_A(\pi(W), N, r, D)$ a.e. by construction, it follows that the solution satisfies (b) in Lemma 4.1 up to time τ. Since the maximizer of a measurable function is measurable and π is measurable, A, β and δ are all predictable. Moreover, because δ^* is bounded and ϕ is a Lipschitz-continuous function of δ^*, they are both square-integrable.

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 [15]) that

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

The volatility term is zero because $N^\top T = 0$. 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 $[\![0,\tau)\!)$,

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

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 $D_t = D_0\mathrm{e}^{rt}$, which is identically zero because $D_0 = 0$. On $\{\tau < \sigma_1\}$ we can repeat this procedure and concatenate the solutions 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 $\mathcal{C}$. Then $D_0 = 0$ on $[\![0,\rho\wedge\sigma_1)\!)$ implies that $\pi(W) = W$ and hence $\delta \in \Psi_A(W, N_W, r, D)$. □

Corollary C.1. *For any affine, compact- and convex-valued D, let $\mathcal{C}$ be a C^1 solution to* (12) *with positive curvature throughout. Then any payoff in the relative interior is attainable by a strategy profile A, enforced by (β,δ) with $\delta \in D\big((W(A)\big)$ such that $W(A)$ remains on $\mathcal{C}$ until either an endpoint of $\mathcal{C}$ is reached or an event occurs.*

Proof. For any $w \in \mathcal{C}$, let $a^*(w)$ and $\delta^*(w)$ denote the maximizers in (12). Since $\mathcal{C}$ is assumed to have positive curvature throughout, the maximization in (12) is not taken over empty sets. By Corollary B.6, the maximizers are attained. □

The following two lemmas establish that locally, $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ coincides with a solution to (5). Lemma C.2 states that it is impossible for a solution to (5) to cut through $\mathcal{B}_r(\mathcal{W})$. For a curve $\mathcal{C}$ with positive curvature throughout, we denote by $\mathcal{N}_\mathcal{C} := \big\{(w,N) \in \mathcal{C}\times S^1 \;\big|\; N^\top(w-v) \geq 0 \;\;\forall\, v \in \mathcal{C}\big\}$ its *outward normal bundle.*

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

$$D(w) := \big\{\delta \in \mathbb{R}^2 \;\big|\; \exists\gamma \in [0,1] \textit{ such that } \gamma w + (1-\gamma)\pi(w) + r\delta \in \mathcal{W}\big\}. \tag{13}$$

It is impossible for a C^1 solution $\mathcal{C}$ to (12) *for D oriented by $v \mapsto N_v$ with end points $v_L, v_R \in U$ to simultaneously satisfy*

(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) $\mathcal{N}_\mathcal{C} \cap \big(\Gamma(r,D) \cup \mathcal{P}\big) = \emptyset$,

(v) *for any* $a \in \mathcal{A}$, $\mathcal{N}_\mathcal{C} \cap \partial E_a(r,D) = \emptyset$.

Proof. Suppose towards a contradiction that there exists such a curve $\mathcal{C}$. Since D is affine, compact- and convex-valued, it follows from Conditions (iv) and (v) as well as Proposition B.7 that $\mathcal{C}$ is C^2 at almost every point. By Condition (iii), there exists a local parametrization f of $\mathcal{C}$ in the direction N'. Define the orthogonal projection $\hat{v} = T'^{\top} v$ onto the tangent for any $v \in U$, where T' is the counterclockwise rotation of N' by 90°. Denote by $\hat{\pi}(v) = \big(\hat{v}, f(\hat{v})\big)$ the projection of $v \in U$ onto $\mathcal{C}$ in the direction N'. By definition of $\mathcal{B}_r(\mathcal{W})$, there exists a solution $\big(W, A, \beta, \delta, M, Z, (J^y)_{y\in Y}\big)$ to (2) with $W_0 = v_0 + \eta N'$ such that on $[\![0, \sigma_1)\!)$, (β, δ) enforces A with $\delta \in D(W)$. Define the stopping time $\tau_1 := \inf\{t \geq 0 \mid W_t \notin U\}$.

Suppose first that $\mathcal{N}_{\mathcal{C}} \subseteq E_a(r, D)$ for some $a \in \mathcal{A}$, i.e., $\mathcal{C}$ is a non-trivial solution to (12). Let $N_t := N_{\hat{\pi}(W_t)}$ and $T_t := T_{\hat{\pi}(W_t)}$ and observe that these projections are well defined on $[\![0, \tau_1)\!)$. We measure the distance of W to $\mathcal{C}$ by $D_t = N'^{\top} W_t - f\big(\hat{W}_t\big)$. Denote $\ell_t := 1/(T_t{}^{\top} T')$ and $\gamma_t := \ell_t N_t{}^{\top} T'$ 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 [12], 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

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

$\xi_t = r\ell_t N_t{}^{\top}\beta_t$, $\rho_t(y) = r\ell_{t-}N_{t-}{}^{\top}\delta_t(y)$ and $\tilde{M}_t = \int_0^t r\ell_{t-}N_{t-}{}^{\top}\mathrm{d}M_t$. Define the stopping time $\tau_2 := \inf\{t \geq 0 \mid D_t \leq 0\}$ and observe that $\tau_2 \leq \tau_1$ a.s. by Condition (i). We will show that there exists an equivalent probability measure R under which the drift rate of D_t is bounded from below by rD_t. Then, D_t becomes arbitrarily large with positive R-probability, and hence positive Q^A-probability. Because it may take arbitrarily long until an accident arrives, this leads to a contradiction because $\mathcal{V}$ is bounded.

Let Ξ_1 denote the set where $N^{\top}\big(\hat{\pi}(W) - g(A) - \delta\lambda(A)\big) \geq 0$. On Ξ_1, $\zeta_t \geq rD_t$, hence there is no need to change the probability measure. It follows from Condition (iv) that $\beta \neq 0$ on Ξ_1^c. Let $\Xi_2 \subseteq \Xi_1^c$ be the set where $\mathcal{N}_{\hat{\mathcal{C}}} \subseteq E_A(r, D)$, i.e., $\Psi_A(\hat{\pi}(W), N, r, D) \neq \emptyset$. Set

$$\hat{\delta} \in \operatorname*{arg\,min}_{x\in\Psi_A(\hat{\pi}(W),N,r,D)} \big\|x - \delta^1\big\| + \big\|x - \delta^2\big\|,$$

then (12) implies that

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

Denote $\Lambda := \max_{a\in\mathcal{A}} \sum_{y\in Y} \lambda(y|a)$ and observe that $N^\top\big(g(A) + \hat{\delta}\lambda(A) - \hat{\pi}(W)\big)$ is uniformly bounded above by the constant $K_1 := \operatorname{diam}\mathcal{V} + \sup(\mathcal{W} - \mathcal{V})\Lambda < \infty$. The condition that $W + r\delta(y) \in \mathcal{W}$ implies that $\delta(y) \in D(W)$ on $[\![0, \tau_2)\!)$ for every $y \in Y$. Due to Lemma E.3, there exist constants K_2, $\bar{\Psi}$ such that

$$\zeta \geq rD - r\ell\Lambda K_2\big\|N^\top\beta\big\| - r\ell K_1\frac{2K_2 + 2\bar{\gamma}}{\bar{\Psi}}\big\|N^\top\beta\big\| =: rD_t - K_3\|\xi\|.$$

On the set $\Xi_1^c \cap \Xi_2^c$, condition (v) implies that $\mathcal{N}_{\mathcal{C}}$ is bounded away from $E_A(r, D) \cup \mathcal{P}$ by virtue of Corollary B.6. Lemma E.4 thus implies that $\big\|N^\top\beta\big\| \geq K_4$ for some constant K_4 and hence

$$\zeta_t \geq rD_t - r\ell_t K_1 \geq rD_t - \frac{K_1}{K_4}\|\xi_t\|.$$

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

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

where

$$\psi_t = K_3\frac{\xi_t}{\|\xi_t\|}1_{\Xi_2} + \frac{K_1}{K_4}\frac{\xi_t}{\|\xi_t\|^2}1_{\Xi_1^c\cap\Xi_2^c}.$$

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

$$D_t \geq D_0 + \int_0^t rD_s \,\mathrm{d}s + \int_0^t \xi_s{}^\top \,\mathrm{d}Z_s' + \tilde{M}_t + \sum_{y\in Y}\int_0^t \rho_s(y)\,\mathrm{d}J_s^y. \tag{14}$$

Since W is bounded, $\int_0^t \xi_s{}^\top \,\mathrm{d}Z_s$ is a $BMO(Q^A)$-martingale. Therefore, $\int_0^t \xi_s{}^\top \,\mathrm{d}Z_s'$ is a $BMO(R)$-martingale by Theorem 3.6 in Kazamaki [16]. 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 (14) 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.,

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

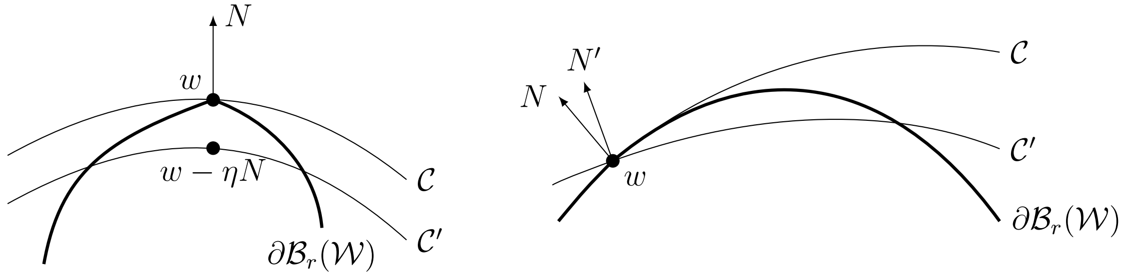

Figure 12: Construction of a curve $\mathcal{C}'$ that cuts through $\mathcal{B}_r(\mathcal{W})$.

where the last inequality follows from the optional stopping theorem and because R is equivalent to Q_T^A. This is a contradiction.

Suppose now that $\mathcal{N}_\mathcal{C} \subseteq \bigcap_{a\in\mathcal{A}} E_a(r,D)^c$, i.e., $\mathcal{C}$ is a straight line segment. Let D denote the distance of W to $\mathcal{C}$ in the direction of the normal vector N' of $\mathcal{C}$. Condition (iv) makes it possible to apply Lemma E.4, hence any (β,δ) enforcing A it follows that $\|N^\top\beta\| \geq K$ for some constant K. Similarly as before, the drift of D_t is thus bounded from below by $rD_t - K_1/\big(Kr\ell_t\|N^\top\beta_t\|\big)$. Therefore, there exists an equivalent probability measure under which D grows arbitrarily large with positive probability, a contradiction. □

Lemma C.3. *Fix $w \in \mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ with outward normal N, where* (5) *is locally Lipschitz continuous. Then $\partial\mathcal{B}_r(\mathcal{W})$ coincides with a solution to* (5) *in a neighbourhood of* (w,N).

Proof. We first show that a solution to (12) with D given in (13) coincides with $\partial\mathcal{B}_r(\mathcal{W})$, which implies that also a solution to (5) stays on $\partial\mathcal{B}_r(\mathcal{W})$. In a sufficiently small neighbourhood of (w,N), (12) admits a unique C^2 solution that is continuous in initial values. Let $\mathcal{C}$ be solution with initial value (w,N) and suppose towards a contradiction that $\mathcal{C}$ escapes $\operatorname{cl}\mathcal{B}_r(\mathcal{W})$ in a neighbourhood of w. Then we can change initial conditions slightly to obtain a curve $\mathcal{C}'$ that cuts through $\mathcal{B}_r(\mathcal{W})$. Specifically:

- If $\partial\mathcal{B}_r(\mathcal{W})$ is not C^1 at w, we obtain $\mathcal{C}'$ as a solution to (5) with initial conditions $(w-\eta N, N)$ for $\eta > 0$ sufficiently small.
- If $\partial\mathcal{B}_r(\mathcal{W})$ is C^1 at w, we obtain $\mathcal{C}'$ for initial conditions (w,N'), where N' is a slight rotation of N as illustrated in the left panel of Figure 12.

Because the set where (12) fails to be locally Lipschitz continuous is closed by Corollary B.6 and Proposition B.7, a small enough perturbation satisfies $\mathcal{N}_{\mathcal{C}'} \cap \Gamma_a(r,D) = \emptyset$ for every $a \in \mathcal{A}$ and either $\mathcal{N}_{\mathcal{C}'} \subseteq E_a(r,D)$ or $\mathcal{N}_{\mathcal{C}'} \cap \big(E_a(r,D) \cup \mathcal{P}\big) = \emptyset$ for any $a \in \mathcal{A}$, that is, $\mathcal{C}'$ satisfies conditions (iv) and (v) of Lemma C.2. By choosing η or N' suitably, we can get conditions (i)–(iii) to hold as well, hence $\mathcal{C}'$ is impossible due to Lemma C.2. We conclude that $\partial\mathcal{B}_r(\mathcal{W})$ is C^1 where (12) is locally Lipschitz continuous and that a solution to (12) cannot escape $\operatorname{cl}\mathcal{B}_r(\mathcal{W})$.

Suppose towards a contradiction that $\mathcal{C}$ falls into the interior of $\mathcal{B}_r(\mathcal{W})$ in a neighbourhood 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 (12), 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 Figure 13. Let v be close enough to w such that (12) with D is Lipschitz continuous on an open neighbourhood of $\mathcal{N}_{\mathcal{C}} := \{(\tilde{w}, N_{\tilde{w}}) \mid \tilde{w} \in \mathcal{C} \text{ between } w \text{ and } v\}$. Let $\delta > 0$ such that the closed ball $B_\delta(v)$ is contained in the interior of $\mathcal{B}_r(\mathcal{W})$. For $\zeta > 0$ to be chosen later, let $\mathcal{W}_\zeta := \{w \in \mathcal{V} \mid d(w, \mathcal{W}) \leq \zeta\}$, where $d(w, \mathcal{W})$ denotes the minimal distance of d from $\mathcal{W}$. Set

$$D_\zeta(w) := \big\{\delta \in \mathbb{R}^2 \;\big|\; \exists \kappa \in [0,1] \text{ such that } \kappa w + (1-\kappa)\pi(w) + r\delta \in \mathcal{W}_\zeta\big\},$$

where π is the projection onto $\partial\mathcal{B}_r(\mathcal{W})$ in the direction N. Observe that for ζ sufficiently small, (12) with D_ζ is Lipschitz continuous in a neighbourhood of $\mathcal{N}_{\mathcal{C}}$, hence it admits a unique solution $\mathcal{C}_\zeta$. Choose now ζ small enough such that Lemma E.5 asserts the existence of $v' \in \mathcal{C}_\zeta \cap B_\delta(v)$.

Because $\mathcal{C}_\zeta$ is continuous in initial conditions, a solution $\mathcal{C}'_\zeta$ to (12) with D_ζ for a slight rotation N' of N reaches a neighbourhood of v' in $\mathcal{B}_r(\mathcal{W})$. As illustrated in Figure 13, $\mathcal{C}'_\zeta$ will escape $\operatorname{cl} \mathcal{B}_r(\mathcal{W})$ to the right of w and enter $\mathcal{B}_r(\mathcal{W})$ to the left of w. Thus, for N' close enough to N, there exist $v_L, v_R \in \mathcal{C}'_\zeta \cap \mathcal{B}_r(\mathcal{W})$, such that $\|\tilde{w} - \pi(\tilde{w})\| \leq \zeta$ for all $\tilde{w} \in \mathcal{C}'_\zeta$. By Corollary C.1, for any $w' \in \mathcal{C}'_\zeta$ there exists a solution to (2) with $W_0 = w'$ such that $\delta \in \Psi_A(W, N, r, D_\zeta)$ on $[\![0, \sigma_1)\!)$ and $W \in \mathcal{C}'_\zeta$ until it reaches an end point of $\mathcal{C}'_\zeta$ or an event occurs. Let $\tau := \inf\big\{t \geq 0 \;\big|\; W_t \in \{v_L, v_R\}\big\}$ and observe that $W_\tau \in \mathcal{B}_r(\mathcal{W})$ on $\{\tau < \sigma_1\}$. The condition that $\delta \in D_\zeta(W)$ a.e. for every $y \in Y$ implies that $x + r\delta_t(y) \in \mathcal{W}_\zeta$ for some x between W_t and $\pi(W_t)$. On $[\![0, \tau \wedge \sigma_1)\!)$ it holds that $\|W_t - x\| \leq \zeta$, and hence $\delta \in \Psi_A(W, N_W, r, D)$. Because $W_\tau \in \mathcal{B}_r(\mathcal{W})$ on $\{\tau < \sigma_1\}$, by definition of $\mathcal{B}_r(\mathcal{W})$ there exists a solution $\big(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}\big)$ with $\tilde{W}_0 = W_\tau$ such that on $[\![\sigma_1, \sigma_2)\!)$, $\big(\tilde{\beta}, \tilde{\delta}\big)$ enforces $\tilde{A}$ and $\tilde{W} + r\tilde{\delta}(y) \in D\big(\tilde{W}\big)$ a.e. Therefore, a concatenation of (W, A, β, δ) with $(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta})$ satisfies the same properties, which shows that $\mathcal{C}'_\zeta \subseteq \mathcal{B}_r(\mathcal{W})$, which is a contradiction to $w' \notin \operatorname{cl} \mathcal{B}_r(\mathcal{W})$. □

Finally, because (5) is Lipschitz continuous almost everywhere, we need to show that $\partial\mathcal{B}_r(\mathcal{W})$ is C^1 to grant uniqueness of the solution. By convexity, $\mathcal{B}_r(\mathcal{W})$ cannot have inward corners, and it will follow with another escaping argument that it cannot have outward corners outside of $\mathcal{D}_r(\mathcal{W})$ either.

Lemma C.4. *$\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is C^1 where (5) fails to be Lipschitz continuous. Moreover, outside of $\mathcal{P}$, the set of all points in $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$, where (5) fails to be Lipschitz continuous, has relative measure 0.*

Proof. We already know from Lemma C.3 that $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is C^1 where it is locally Lipschitz continuous. Suppose, therefore, that $\partial\mathcal{B}_r(\mathcal{W})$ has a corner at $w \in$

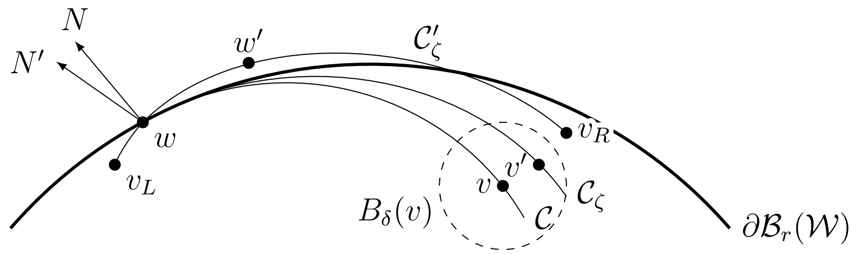

Figure 13: If $\mathcal{C}$ falls into the interior of $\mathcal{B}_r(\mathcal{W})$, there exists a solution $\mathcal{C}'_\zeta$ to (9) with initial conditions (w, N') and a slight reduction over controls $\delta \in D_\zeta(w)$ such that $\mathcal{C}'_\zeta$ escapes $\mathcal{B}_r(\mathcal{W})$. For small ζ and N' close to N, there exists an enforceable strategy profile attaining $w' \notin \mathcal{B}_r(\mathcal{W})$ which reaches $\mathcal{B}_r(\mathcal{W})$ with certainty. This leads to a contradiction.

$\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ where (5) fails to be locally Lipschitz continuous. Proposition B.7 implies that (w, N_L) and (w, N_R) are cointained in $\partial E_a(r, \mathcal{W})$ and $\partial E_{a'}(r, \mathcal{W})$, respectively, where N_L and N_R are the extremal normal vectors to $\partial\mathcal{B}_r(\mathcal{W})$ at w. Since $E_a(r, \mathcal{W}) \cup \mathcal{P}$ is closed, there exists an open neighborhood U of w and a set of normal vectors $\mathcal{T} \subseteq \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ such that $U \cap \operatorname{int}\mathcal{B}_r(\mathcal{W}) \times \mathcal{T}$ is either contained in or has empty intersection with E_a. Since $w \notin \mathcal{G}$, it follows that $\{w\} \times \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \cap \Gamma(r, \mathcal{W}) = \emptyset$. Since $\Gamma(r, \mathcal{W})$ is closed, $U \cap \operatorname{int}\mathcal{B}_r(\mathcal{W}) \times \mathcal{T} \cap \Gamma(r, \mathcal{W}) = \emptyset$ for U sufficiently small. We can thus construct a solution to $\mathcal{C}$ to (12) for D given in (13) with initial conditions $(v, N) \in U \cap \operatorname{int}\mathcal{B}_r(\mathcal{W}) \times \mathcal{T}$ that cuts through $\mathcal{B}_r(\mathcal{W})$ with $\mathcal{N}_{\mathcal{C}} \subseteq U \cap \operatorname{int}\mathcal{B}_r(\mathcal{W}) \times \mathcal{T}$. By choice of U and $\mathcal{T}$, the conditions in Lemma C.2 are satisfied, which is a contradiction.

For the second statement, suppose that there exists $\mathcal{C} \subseteq \partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ of positive length. By shortening the line segment we may assume that $\mathcal{N}_{\mathcal{C}} \subseteq \mathcal{P}$ or $\mathcal{N}_{\operatorname{int}\mathcal{C}} \cap \mathcal{P} = \emptyset$. Suppose towards a contradiction that $\mathcal{N}_{\operatorname{int}\mathcal{C}} \cap \mathcal{P} = \emptyset$. Then Proposition B.7 shows that (w, N) enters and leaves $E_a(r, D)$ of the maximizing action profile a at almost every $(w, N) \in \mathcal{N}_{\mathcal{C}}$. Because $\mathcal{A}$ is finite we may assume that this is the same action profile. This implies that $\mathcal{N}_{\mathcal{C}} \subseteq \partial E_a(r, D)$ and hence $\mathcal{N}_{\operatorname{int}\mathcal{C}} \subseteq E_a(r, D)$ by Corollary B.6, a contradiction. □

Next, we prove Proposition 6.7, showing that extremal points of any C^1 segment in $\mathcal{D}_r(\mathcal{W})$ and any corners of $\mathcal{D}_r(\mathcal{W})$ must lie in $\mathcal{K}_r(\mathcal{W})$.

Proof of Proposition 6.7. Suppose first that w is a corner of $\mathcal{B}_r(\mathcal{W})$ in $\mathcal{D}_r(\mathcal{W}) \setminus \mathcal{S}_r(\mathcal{W})$. By definition of $\mathcal{D}_r(\mathcal{W})$, there exists (a, δ_0) such that $\delta_0 \in \Psi_a$, $w + r\delta_0(y) \in \mathcal{W}$ for every $y \in Y$, and $N^\top\big(g(a) + \delta_0\lambda(a) - w\big) \geq 0$ for all $N \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Since w is not a stationary point, $w \neq g(a) + \delta_0\lambda(a)$ and hence $N^\top\big(g(a) + \delta_0\lambda(a) - w\big) > 0$ for almost all normal vectors in $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Lemma 6.6 thus readily implies that $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$. It remains to show that the normal vectors to the sets $\mathcal{B}_r(\mathcal{W})$ and $\mathcal{K}_{r,a}(\mathcal{W})$ satisfy the desired inclusion property. Suppose towards a contradiction that the converse holds, that is, $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big) \not\subseteq \mathcal{N}_w\big(\mathcal{K}_{r,a}(\mathcal{W})\big)$. Since both $\mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ and $\mathcal{N}_w\big(\mathcal{K}_{r,a}(\mathcal{W})\big)$ are closed, there exists $N \in \mathcal{N}_w\big(\mathcal{K}_{r,a}(\mathcal{W})\big) \setminus \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$ with $N^\top\big(g(a) + \delta_0\lambda(a) - w\big) > 0$.

By definition of the normal vector, $w_\varepsilon := w + \varepsilon T \in \operatorname{int}\mathcal{K}_{r,a} \setminus \operatorname{cl}\mathcal{B}_r(\mathcal{W})$ for $\varepsilon > 0$ sufficiently small, where T is orthogonal to N with $T^\top N' < 0$ for all $N' \in \mathcal{N}_w\big(\mathcal{B}_r(\mathcal{W})\big)$. Fix such an ε sufficiently small and define $w_{\varepsilon,\gamma} := \gamma w_+(1-\gamma)w\varepsilon$ for $\gamma \in [0,1]$. Since w_ε is in the interior of $\mathcal{K}_{r,a}(\mathcal{W})$, there exists $\delta' \in \operatorname{int}\Psi_a$ with $w_\varepsilon + r\delta'(y) \in \operatorname{int}\mathcal{W}$ for every $y \in Y$. Let $\delta_\gamma := \gamma\delta_0 + (1-\gamma)\delta'\ \delta_\varepsilon$. Since $\mathcal{W}$ and Ψ_a are convex, it follows that for any $\gamma > 0$, $\delta_\gamma \in \operatorname{int}\Psi_a$ and $w_{\varepsilon,\gamma} + r\delta_\gamma(y) \in \operatorname{int}\mathcal{W}$ for every $y \in Y$. Moreover for γ sufficiently small, $N^\top\big(g(a) + \delta_\gamma\lambda(a) - w_{\varepsilon,\gamma}\big) > 0$. A contradiction can thus be obtained in the same way as in the proof of Lemma 6.6.

Suppose next that $\mathcal{D}_r(\mathcal{W})$ contains a C^1 line segment $\mathcal{C}$ of positive length and positive curvature. Let w be in the relative interior of $\mathcal{C}$ and denote by N_w the unique normal vector to $\partial\mathcal{B}_r(\mathcal{W})$ at w. If there exists (a, δ_w) with $\delta_w \in \Psi_a \cap (\mathcal{W} - w)^m/r$ with ${N_w}^\top\big(g(a) + \delta_w\lambda(a) - w\big) > 0$, then the argument works in the same way as before: conditions (i) and (iii) of Lemma 6.6 have to be violated, showing that $w \in \partial\mathcal{K}_{r,a}(\mathcal{W})$ and if $N_w \notin \mathcal{N}_w\big(\mathcal{K}_{r,a}(\mathcal{W})\big)$, we can enlarge the set in the same way as before. Suppose, therefore, that every $w \in \mathcal{C}$ is decomposable only by (a_w, δ_w) with ${N_w}^\top\big(g(a_w) + \delta_w\lambda(a_w) - w\big) = 0$. Since $w \notin \mathcal{S}_r(\mathcal{W})$, the dirft is parallel to $\partial\mathcal{C}$. Thus, there exists v outside of $\mathcal{B}_r(\mathcal{W})$ arbitrarily close to w with ${N_w}^\top\big(g(a_w)+\delta_w\lambda(a_w)-w\big) > 0$. If any such v is in $\mathcal{K}_{r,a}$, we obtain a contradiction in the same way as before. Therefore, $w \in \partial\mathcal{K}_{r,a_w}(\mathcal{W})$. Finally, if $N_w \notin \mathcal{N}_w\big(\mathcal{K}_{r,a_w}(\mathcal{W})\big)$, then there exists a payoff pair $v \in \mathcal{C}$ arbitrarily close to v that are in the interior of $\mathcal{K}_{r,a_w}(\mathcal{W})$. Since there are only finitely many action profiles, this is a contradiction. $\square$

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

This appendix shows that the set $\mathcal{B}_r(\mathcal{W})$ is closed. It also contains the proof of Theorem 6.8 based on the auxiliary results in Appendix C and this appendix.

Lemma D.1. *Suppose that $w \in \mathcal{D}_r(\mathcal{W})$ is either a corner or part of a continuously differentiable line segment in $\mathcal{D}_r(\mathcal{W})$. Then $w \in \mathcal{B}_r(\mathcal{W})$.*

Proof. By Lemma 6.3, any payoff pair $w \in \mathcal{S}_r(\mathcal{W})$ is contained in $\mathcal{B}_r(\mathcal{W})$. Suppose, therefore, that $w \notin \mathcal{S}_r(\mathcal{W})$ and consider first the case where w is a corner of $\mathcal{B}_r(\mathcal{W})$.

Suppose first that there exist an action profile $a \in \mathcal{A}$, a measurable selection $\delta^* : \mathcal{K}_{r,a} \to \Psi_a$ and a time $t_0 > 0$ such that the solution (W, A, β, δ, M) to (2) with $W_0 = w$, $A \equiv a$, $\beta \equiv 0$, $\delta = \delta^*(W)$, and $M \equiv 0$ remains in $\mathcal{B}_r(\mathcal{W}) \cap \mathcal{K}_{r,a}$ on $[\![0, t_0 \wedge \sigma_1)\!)$ and W_{t_0} is in the interior of $\mathcal{B}_r(\mathcal{W})$ on the set $\{t_0 < \sigma_1\}$. Observe that such a solution W to (2) travels on a deterministic path before the arrival of an infrequent event $y \in Y$. In particular, δ is predictable. Since $W_{t_0} \in \mathcal{B}_r(\mathcal{W})$ on the event $\{t_0 < \sigma_1\}$, there exists a solution $\big(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}, \tilde{M}\big)$ attaining W_{t_0} with $(\tilde{\beta}, \tilde{\delta})$ enforcing $\tilde{A}$ such that $\tilde{W} \in \mathcal{B}_r(\mathcal{W})$ up until the arrival of the first infrequent event, at which point W jumps to $\mathcal{W}$. The concatenation $(W, A, \beta, \delta, M)1_{[0,t_0]} + (\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}, \tilde{M})1_{(t_0,\infty)}$ thus satisfies the same conditions on $[0, \infty)$, showing that $w \in \mathcal{B}_r(\mathcal{W})$.

Suppose now that no such a, δ^*, t_0 exist. Then there exists no strategy profile attaining w without using the Brownian information to structure incentives. We can thus obtain a contradiction using an escaping argument similarly to Lemma C.2. Let $\mathcal{C}_\varepsilon$ be a solution to (5) starting in $w - \varepsilon N$ for an outward normal vector N and $\varepsilon > 0$. For ε sufficiently small, $\mathcal{C}_\varepsilon$ intersects $\partial\mathcal{B}_r(\mathcal{W})$ and satisfies the conditions of Lemma C.2 outside a set of measure 0. Such a curve is impossible by Lemma C.2.

Suppose next that $\mathcal{D}_r(\mathcal{W})$ is locally a C^1 line segment $\mathcal{C}$. Since there are only finitely many action profiles, there exists a subsegment $\mathcal{C}'$ of positive length that can be decomposed by some action profile $a \in \mathcal{A}$. If $\mathcal{C}'$ has positive curvature, then $\mathcal{C}' \subseteq \partial\mathcal{K}_{r,a}(\mathcal{W})$ by Proposition 6.7. Since $\mathcal{K}_{r,a}(\mathcal{W})$ is convex, a measurable selector δ^* as in the previous case exists and $\mathcal{C}' \in \mathcal{D}_r(\mathcal{W})$. If $\mathcal{C}'$ is a straight line segment, then any payoff pairs in the relative interior can be attained by public randomization and its endpoints can be attained as before. □

Lemma D.2. *$\mathcal{B}_r(\mathcal{W})$ is closed.*

Proof. By public randomization, a straight line segment is contained in $\mathcal{B}_r(\mathcal{W})$ if both of its end points are contained in $\mathcal{B}_r(\mathcal{W})$. Similarly, Lemma 6.4 shows that curved parts of $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ are contained in $\mathcal{B}_r(\mathcal{W})$ if its end points are. Since Lemma D.1 shows that $\mathcal{D}_r(\mathcal{W}) \subseteq \mathcal{D}_r(\mathcal{W})$, the closure of $\mathcal{B}_r(\mathcal{W})$ is $\mathcal{W}$-relaxed generating, hence it is contained in $\mathcal{B}_r(\mathcal{W})$ by maximality. □

Proof of Theorem 6.8. Lemmas C.3 and C.4 imply that $\partial\mathcal{B}_r(\mathcal{W}) \setminus \mathcal{D}_r(\mathcal{W})$ is a C^1 solution to (5). It follows from Proposition 6.7 that $G_r(\mathcal{W})$ has the desired properties. Finally, $\mathcal{B}_r(\mathcal{W})$ is closed by Lemma D.2. □

E Auxiliary results related to the optimality equation

E.1 Lipschitz continuity of set-valued maps

Consider an arbitrary family $(F_i)_{i\in I}$ of Lipschitz continuous set-valued maps. If their Lipschitz constants $(K_i)_{i\in I}$ are uniformly bounded, then the union $x \mapsto \bigcup_{i\in I} F_i(x)$ is Lipschtiz continuous again. However, the intersection of two Lipschitz continuous maps may fail to be Lipschitz continuous in general. In this appendix, we show Lipschitz continuity of the intersection for two special cases that are relevant in our setting.

Lemma E.1. *The intersection of two convex-valued affine maps is Lipschitz continuous.*

Proof. Let F and G be two convex-valued affine functions. It is sufficient to show that $F \cap G$ is continuous as it is then Lipschitz continuous since both F and G are affine. Suppose towards a contradiction that $F \cap G$ fails to be continuous at some x_0,

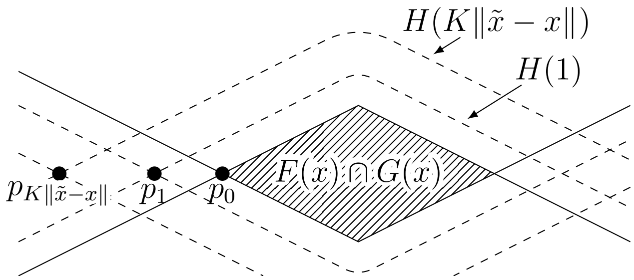

Figure 14: Level sets $H(z)$ of $\partial\big(F(x)\cap G(x)\big)$ containing points p_z with maximal distance from $p_0 \in \partial\big(F(x)\cap G(x)\big)$. Clearly, $\|p_z - p_0\| = z\|p_1 - p_0\|$.

that is, there exists $v \in F(x_0)\cap G(x_0)$ such that $B_\varepsilon(v)\cap F(x)\cap G(x) = \emptyset$ for $\varepsilon > 0$ arbitrarily small and $x \in \operatorname{supp} F\cap G$ arbitrarily close to x_0. Since F and G are affine, this is only possible if $N_F = -N_G$, where N_F and N_G denote the normal vectors to $\partial F(x_0)$ and $\partial G(x_0)$, respectively, at v. It follows from convexity that this is possible only if $F(x)\cap G(x) = \emptyset$ for x arbitrarily close to x_0, contradicting the fact that $x \in \operatorname{supp} F \cap G$. $\square$

Lemma E.2. *Let F and G be Lipschitz continuous maps with bounded support taking values in closed convex polytopes. Denote the outward normal vectors to their hyperfaces by $\pi_i^F(x)$, $i \in I_F$ and $\pi_i^G(x)$, $i \in I_G$, respectively. If for any $J_F \subseteq I_F$, $J_G \subseteq I_G$, the matrix $[(\pi_j^F(x))_{j\in J_F}, (\pi_j^G(x))_{j\in J_G}]$ has constant column rank in a neighbourhood of x, then $F\cap G$ is locally Lipschitz continuous at x. If the ranks of above matrices are constant on the entire support of $F\cap G$, then $F\cap G$ is Lipschitz continuous.*

Proof. Fix x in the support of $F\cap G$ and let K be the maximum of the Lipschitz constants of F and G. Then Lipschitz continuity of the individual maps implies that

$$F(\tilde{x})\cap G(\tilde{x}) \subseteq \big(F(x) + \|\tilde{x}-x\|B_K(0)\big) \cap \big(G(x) + \|\tilde{x}-x\|B_K(0)\big).$$

Observe, however, that the right hand side is larger than $F(x)\cap G(x) + \|\tilde{x}-x\|B_K(0)$. Let $H(z) := \partial\big(F(x)+B_z(0)\big)\cap\big(G(x)+B_z(0)\big)$ be the level sets of $\partial\big(F(x)\cap G(x)\big)$. Let p_1 denote the point in $H(1)$ with maximal distance from $\partial\big(F(x)\cap G(x)\big)$ and let p_0 be the point in $\partial\big(F(x)\cap G(x)\big)$ with minimal distance from p_1 as illustrated in Figure 14. Let $\{\pi_1,\ldots,\pi_n\}$ be a minimal subset of normal vectors to the hyperfaces of $F(x)\cap G(x)$ that intersect at p_0 such that p_1 is the unique point in $H(1)$, which is related to p_0 by

$$\pi_j^\top(p_1-p_0) = 1 \quad \text{for } j = 1,\ldots,m \quad \text{and} \quad p_1 - p_0 \in \operatorname*{span}_{j=1,\ldots,m} \pi_j. \tag{15}$$

By linearity of (15), it follows that $p_{K\|\tilde{x}-x\|} := p_0 + K\|\tilde{x}-x\|(p_1-p_0)$ is a point in $H(K\|\tilde{x}-x\|)$ with maximal distance of $F(x)\cap G(x)$. Its distance from p_0 equals $K\|p_1-p_0\|\|\tilde{x}-x\|$. The statement thus follows once we show that $\|p_1(x)-p_0(x)\|$ is uniformly bounded in x.

By minimality of $\{\pi_1, \dots, \pi_n\}$, the vectors $\pi_1, \dots, \pi_n$ are linearly independent. Thus by assumption, $\pi_1(\tilde{x}), \dots, \pi_n(\tilde{x})$ are linearly independent also for $\tilde{x}$ in a neighbourhood of x. Since F and G are continuous, the norm of the solution is continuous, hence by making the neighbourhood smaller and compact, its maximum is bounded. Because F and G have finitely many hyperfaces, the finite maximum over all possible combinations of normal vectors $\pi_1, \dots, \pi_{\tilde{n}}$ yields a bound for $\|p_1 - p_0\|$ on a sufficiently small neighbourhood of x. Finally, if the rank is constant on $\operatorname{supp} F \cap G$, then $\|p_1 - p_0\|$ is uniformly bounded since $\operatorname{supp} F \cap G$ is compact. □

E.2 Bounds on incentives and solutions of the optimality equation

Denote by $\Phi_a(N, \delta) := \{\beta \mid (\beta, \delta) \text{ enforces } a \text{ and } N^\top \beta = 0\}$ the space of value transfers normal to N that enforce a given value burning δ. Let

$$\Psi_a(w, N, r, \mathcal{W}) := \left\{ \delta \in \mathbb{R}^{2\times|Y|} \;\middle|\; \begin{array}{l} \Phi_a(\delta, N) \neq \emptyset, N^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0, \\ w + r\delta(y) \in \mathcal{W} \text{ for every } y \in Y \end{array} \right\}.$$

Relevant is the optimality equation

$$\kappa_a(w) = \max_{\delta \in \Psi_a(w, N_w, r, \mathcal{W})} \frac{2{N_w}^\top\big(g(a) + \delta\lambda(a) - w\big)}{r\|\phi(a, N_w, \delta)\|^2}, \tag{16}$$

where $\phi(a, N_w, \delta)$ is the shortest vector in $\Phi_a(N_w, \delta)$.

Lemma E.3. *Let $\mathcal{C}$ be a C^1 solution to (16) for fixed a, r, and $\mathcal{W}$ with endpoints v_L, v_R such that the normal vector N_w to $\mathcal{C}$ at w is contained in $E_a^0(r, \mathcal{W}) \setminus \big(\Gamma_a^0(r, \mathcal{W}) \cup \mathcal{P}\big)$ for every $w \in \mathcal{C}$. Then there exists a constant $K > 0$ such that for any $\alpha \geq 0$, for any $w \in \mathcal{C}$, any $(T_w\phi + N_w\chi, \delta)$ enforcing a with ${N_w}^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$ and $w + r\delta(y) \in \mathcal{W}$ satisfies*

$$K\|\chi\| \geq \big\|\hat{\delta}^1 - \delta^1\big\| + \big\|\hat{\delta}^2 - \delta^2\big\|, \qquad \frac{2K + 2\alpha}{\bar{\Psi}}\|\chi\| \geq 1 - \frac{\big(\|\phi\| - \alpha\|\chi\|\big)^2}{\big\|\phi\big(a, N, \hat{\delta}\big)\big\|^2}, \tag{17}$$

where $\hat{\delta}$ is the element of $\Psi_a(w, N_w, r, D)$ that minimizes $\big\|\hat{\delta}^1 - \delta^1\big\| + \big\|\hat{\delta}^2 - \delta^2\big\|$ and $\bar{\Psi} := \inf_{w \in \mathcal{C}} \min_{\delta' \in \Psi_a(w, N_w, r, \mathcal{W})} \|\phi(a, N_w, \delta')\|^2 > 0$.

Proof. We begin the proof by extending $\phi(a, N, \delta)$ to δ with $\Phi_a(N, \delta) = \emptyset$ in a Lipschitz continuous way. We achieve this by introducing the function $\chi(a, N, \delta)$ characterizing the minimal normal component $N^\top\beta$ necessary to enforce a given δ, and setting $\phi(a, N, \delta)$ equal to the minimal tangential component ϕ necessary such that $\big(T\phi + N\chi(a, N, \delta), \delta\big)$ enforces a. Observe that this is indeed an extension of $\phi(a, N, \delta)$ as $\chi(a, N, \delta) \equiv 0$ for δ with $\Phi_a(N, \delta) \neq \emptyset$.

Recall that $\mathcal{I}_a^i(N,\delta^i)$ denotes the set of vectors $\phi \in \mathbb{R}^d$ such that $(T^i\phi,\delta^i)$ satisfies i's incentive compatibility constraints for action profile a, where T is (the clockwise) orthogonal to N. Let $d\big(\mathcal{I}_a^1(N,\delta^1),\mathcal{I}_a^2(N,\delta^2)\big)$ denote the minimum distance between the two sets. Since $\mathcal{I}_a^i(N,\delta)$ are closed for $i=1,2$, the minimum distance is attained between two points $p_i \in \mathcal{I}_a^i(N,\delta^i)$. If the minimal distance is attained for more than one such pair, let (p_1,p_2) be the pair that minimizes the norm of

$$\frac{N^1/T^1 p_1 + N^2/T^2 p_2}{N^1/T^1 + N^2/T^2}.$$

Because changes in N and δ only change the location, but not the direction of the hyperfaces of $\mathcal{I}_a^i(N,\delta)$, p_1 and p_2 are Lipschitz continuous in (N,δ). Therefore, so are

$$\chi(a,N,\delta) := \frac{p_1 - p_2}{N^1/T^1 + N^2/T^2} \quad \text{and} \quad \phi(a,N,\delta) := \frac{N^1/T^1 p_1 + N^2/T^2 p_2}{N^1/T^1 + N^2/T^2}.$$

This definition indeed minimizes first $\|\chi\|$ and then $\|\phi\|$: if $(T\phi + N\chi,\delta)$ enforces a, then $\phi \in \mathcal{I}_a^i(N,\delta^i) - N^i/T^i\chi$ for $i=1,2$ and hence

$$\left(\frac{N^1}{T^1} + \frac{N^2}{T^2}\right)\|\chi\| \geq d\big(\mathcal{I}_a^1(N,\delta^1),\mathcal{I}_a^2(N,\delta^2)\big). \tag{18}$$

Since $\chi(a,N,\delta)$ attains the lower bound in (18), its norm is minimal. Minimality of $\phi(a,N,\delta)$, given $\chi(a,N,\delta)$, follows since $\phi(a,N,\delta)$ is the orthogonal projection of $p1 - p2 - N\chi(a,N,\delta)$ onto T.

Note that $\Psi_a(w,N_w,r,\mathcal{W})$ is non-empty for any $w \in \mathcal{C}$ since $(w,N_w) \in E_a(r,\mathcal{W})$ by assumption. Define the function $D(w,\delta) := \min_{x\in\Psi_a(w,N_w,r,\mathcal{W})} \|x^1-\delta^1\|+\|x^2-\delta^2\|$ and observe that it is well-defined for any $w \in \mathcal{C}$ since $\Psi_a(w,N_w,r,\mathcal{W})$ is non-empty because $(w,N_w) \in E_a(r,\mathcal{W})$ by assumption. Observe also that D is continuous —as it minimizes a convex function over a convex set—and that it is equal to 0 for $\delta \in \Psi_a(w,N_w,r,\mathcal{W})$, which is precisely the case when $\mathcal{I}_a^1(N_w,\delta^1)$ and $\mathcal{I}_a^2(N_w,\delta^2)$ overlap. Since a change in δ corresponds to a linear shift of the hyperfaces of $\mathcal{I}_a^i(N_w,\delta^i)$, $D(w,\delta)$ is piecewise linear in δ. For fixed w, the rate at which $D(w,\delta)$ changes depends on the angle between the closest hyperfaces of $\mathcal{I}_a^1(N_w,\delta^1)$ and $\mathcal{I}_a^2(N_w,\delta^2)$. Since changes in δ do not affect these angles and because there are finitely many hyperfaces, there exists $K_1(w)$ with $|\mathrm{d}D(\delta)/\mathrm{d}\delta| \geq K_1(w) > 0$ outside of $\Psi_a(w,N_w,r,\mathcal{W})$. Changing N_w corresponds to rescaling the sets $\mathcal{I}_a^i(N_w,\delta^i) = \frac{1}{N^i}\mathcal{I}_a^i(\delta)$, but since $\frac{1}{N^i} \geq 1$, it follows that $K_1 := \min_{w\in\mathcal{C}} K_1(w) > 0$. Therefore, $K_1 D(w,\delta) \leq d\big(\mathcal{I}_a^1(N_w,\delta^1),\mathcal{I}_a^2(N_w,\delta^2)\big)$, which implies by virtue of (18) that

$$K_1\big\|\hat{\delta}^1 - \delta^1\big\| + K_1\big\|\hat{\delta}^2 - \delta^2\big\| \leq \sup_{w\in\mathcal{C}}\left(\frac{N_w^1}{T_w^1} + \frac{N_w^2}{T_w^2}\right)\|\chi(a,N,\delta)\|.$$

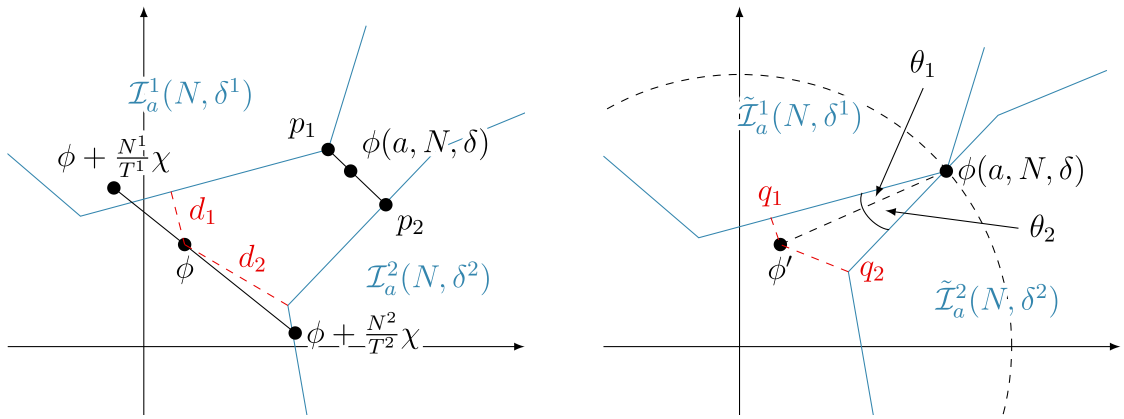

Figure 15: The left panel illustrates the position of $\phi(a,N,\delta)$, p_1 and p_2 relative to $\mathcal{I}_a^1(N,\delta^1)$ and $\mathcal{I}_a^2(N,\delta^2)$. It also shows that necessarily, $d_i \le N^i/T^i\|\chi\|$. The right panel shows that $\theta_1+\theta_2=\gamma$ and hence $\theta_j \ge \gamma/2$.

Because N_w is bounded away from coordinate directions ($(w,N_w)\notin\mathcal{P}$), the supremum on the right-hand side is finite. Since $\|\chi(a,N,\delta)\| \le \|\chi\|$, this proves the first inequality of (17) for constant

$$K = \frac{1}{K_1}\sup_{w\in\mathcal{C}}\left(\frac{N_w^1}{T_w^1}+\frac{N_w^2}{T_w^2}\right).$$

The second inequality follows once we show that there exists $K_3>0$ with

$$\big\|\phi\big(a,N,\hat{\delta}\big)\big\| - \|\phi\| \le K_3 D(\delta) + K_3\|\chi\|. \tag{19}$$

Indeed, the right hand side of (19) is bounded by $K_3(K+1)\|\chi\|$ due to the already established inequality. Thus,

$$1-\frac{\|\phi\|-\alpha\|\chi\|}{\big\|\phi\big(a,N,\hat{\delta}\big)\big\|} \le \frac{\big\|\phi\big(a,N,\hat{\delta}\big)\big\| - \|\phi\| + \alpha\|\chi\|}{\bar{\Psi}} \le \frac{K_3(K+1)+\alpha}{\bar{\Psi}}\|\chi\|. \tag{20}$$

Observe that $\bar{\Psi}>0$ since $\mathcal{N}_{\mathcal{C}}$ is bounded away from $\Gamma(r,D)\cup\mathcal{P}$ by closedness. The second inequality in (17) then follows from (20) in conjunction with $1-x\ge\frac{1}{2}(1-x^2)$.

It remains to show (19). Suppose first that $\|\phi\| \ge \|\phi(a,N,\delta)\|$. Then Lipschitz continuity of ϕ implies that $\|\phi\| \ge \big\|\phi\big(a,N,\hat{\delta}\big)\big\| - K_\phi\big\|\hat{\delta}-\delta\big\|$, which readily implies (19). Suppose therefore that $\|\phi\| < \|\phi(a,N,\delta)\|$. Let d_i denote the distance of ϕ from $\mathcal{I}_a^i(N,\delta)$ for $i=1,2$ and observe that $d_i \le N^i/T^i\|\chi\|$ as illustrated in the left panel of Figure 15. Define the auxiliary sets $\tilde{\mathcal{I}}_a^i(N,\delta^i) := \mathcal{I}_a^i(N,\delta^i) - N^i/T^i\chi(a,N,\delta)$ so that $\phi(a,N,\delta)\in\tilde{\mathcal{I}}_a^1(N,\delta^1)\cap\tilde{\mathcal{I}}_a^2(N,\delta^2)$ as shown in the right panel of Figure 15. Let $\tilde{d}_i$ denote the distance of ϕ from $\tilde{\mathcal{I}}_a^i(N,\delta)$ and observe that $\tilde{d}_i \le d_i$. Let q_i for $i=1,2$ denote the point in $\partial\tilde{\mathcal{I}}_a^i(N,\delta)$ closest to ϕ and let ϕ' be the projection of ϕ

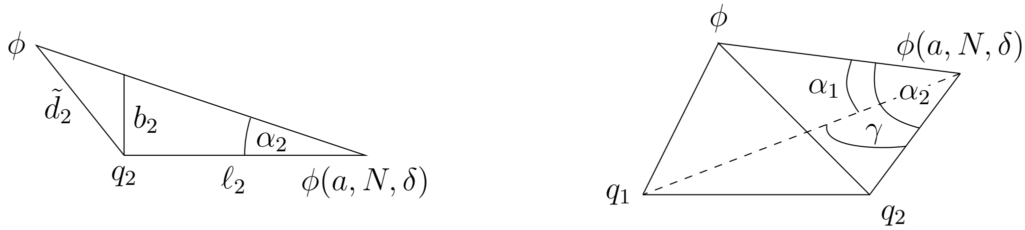

Figure 16: Because q_2 is the projection of ϕ onto $\partial\tilde{\mathcal{I}}_a^2(N, \delta^2)$, the angle at q_2 in the triangle shown to the left is at least $90°$. Therefore, $\tilde{d}_2 \geq b_2 = \ell_2 \tan(\alpha_2)$ and thus the triangle inequality implies $\|\phi(a, N, \delta) - \phi\| \leq \tilde{d}_2 + \ell_2 \leq \tilde{d}_2\big(1 + \frac{1}{\tan(\alpha_2)}\big)$. The right panel illustrates that $\alpha_j \geq \theta_j \geq \gamma/2$.

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 $\theta_j \geq \gamma/2$, where γ is the angle between $\phi(a, N, \delta) - q_1$ and $\phi(a, N, \delta) - q_2$. Let α_i be the angles between $\phi(a, N, \delta) - \phi$ and $\phi(a, N, \delta) - q_i$ and observe that $\alpha_i \geq \theta_i$. Then

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

as illustrated in Figure 16. Observe that it is impossible for γ to be 0 by the definition of $\phi(a, N, \delta)$. Since changes in N and δ do not change the direction of the hyperplanes bounding $\tilde{\mathcal{I}}_a^i(N, \delta)$, a uniform lower bound $\underline{\gamma}$ for γ is given by taking the minimum over all strictly positive angles between the finitely many hyperfaces of $\partial\tilde{\mathcal{I}}_a^1(N, \delta)$ and $\partial\tilde{\mathcal{I}}_a^2(N, \delta)$. Therefore, $\|\phi(a, N, \delta) - \phi\| \leq K_4\|\chi\|$ for

$$K_4 = \left(1 + \frac{1}{\tan(\underline{\gamma}/2)}\right)\sup_{w\in\mathcal{C}}\left(\frac{N_w^1}{T_w^1} + \frac{N_w^2}{T_w^2}\right).$$

(19) now follows from the triangle inequality

$$\big\|\phi\big(a, N, \hat{\delta}\big) - \phi\big\| \leq \big\|\phi\big(a, N, \hat{\delta}\big) - \phi(a, N, \delta)\big\| + \|\phi(a, N, \delta) - \phi\| \leq K_\phi D(\delta) + K_4\|\chi\|. \quad \square$$

Lemma E.4. *Let $\mathcal{C}$ be a C^1 solution to (9) for fixed $r, \mathcal{W}$ oriented by $w \mapsto N_w$ such that $\mathcal{N}_\mathcal{C} \cap \big(\Gamma_a(r, \mathcal{W}) \cup E_a(r, \mathcal{W}) \cup \mathcal{P}\big) = \emptyset$ for some $a \in \mathcal{A}$. Then there exists $K > 0$ such that for any $w \in \mathcal{C}$, any pair $(T_w\phi + N_w\chi, \delta)$ that enforces a with $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and ${N_w}^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$ satisfies $K \leq \|\chi\|$.*

Proof. Let $(T_w\phi + N_w\chi, \delta)$ enforce a with $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and ${N_w}^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$. The condition $(\mathcal{C}, \mathcal{N}_\mathcal{C}) \cap \Gamma(r, \mathcal{W}) = \emptyset$ implies that $T_w\phi + N_w\chi \neq 0$. The condition $(C, \mathcal{N}_\mathcal{C}) \cap E_a(r, \mathcal{W}) = \emptyset$ implies that $\chi \neq 0$ because $\Psi_a(w, N_w, r, \mathcal{W}) = \emptyset$ for any $w \in \mathcal{C}$. This is equivalent to $\mathcal{I}_a^1(N_w, \delta) \cap \mathcal{I}_a^2(N_w, \delta) = \emptyset$, implying that the two sets $\mathcal{I}_a^1(N_w, \delta)$, $\mathcal{I}_a^2(N_w, \delta)$ are strictly separated since they are closed. Let $d(w, \delta)$ denote the minimal distance between the two sets. Because N_w is bounded away from coordinate directions, the map $(N_w, \delta) \mapsto \mathcal{I}_a^i(N_w, \delta)$ is continuous for $i = 1, 2$, hence so is $d(w, \delta)$. Let $\mathcal{J}_a(w)$ denote the set of all δ, for which there exists β such that (β, δ) enforces a with $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and $N_w{}^\top\big(g(a) + \delta\lambda(a) - w\big) \geq 0$. Since $\mathcal{C}$ is C^1, $\mathcal{J}_a$ is continuous in $w \in \mathcal{C}$. The minimum of $\min_{\delta \in \mathcal{J}_a(w)} d(w, \delta)$ over the compact set $\mathcal{C}$ is attained, hence positive. This implies the statement by virtue of (18). $\square$

Lemma E.5. *Let $\varepsilon_1 \leq \varepsilon_2$ and let $D_1(w)$ and $D_2(w)$ be two affine, convex-, and compact-valued maps such that there exists $\varepsilon > 0$ with $D_2(w) \cap H_{\varepsilon_2}(N) \subseteq D_1(w) \cap H_{\varepsilon_1}(N) \subseteq D_2(w) \cap H_{\varepsilon_2}(N) + B_\varepsilon(0)$ for every $w \in \mathcal{V}$ and every $N \in S^1$. Let (w, N) such that (12) with $\big(\varepsilon_1, D_1(w)\big)$ and $\big(\varepsilon_2, D_2(w)\big)$ is Lipschitz continuous in a neighbourhood U of (w, N). Let $\mathcal{C}_1$ and $\mathcal{C}_2$ be two solutions to (12) with $\big(\varepsilon_1, D_1(w)\big)$ and $\big(\varepsilon_2, D_2(w)\big)$, respectively, with initial value (w, N) such that $\mathcal{N}_{\mathcal{C}_1}, \mathcal{N}_{\mathcal{C}_2} \subseteq U$. Then there exist constants K_1, K_2, K_3 such that for any $v \in \mathcal{C}_1$, there exists $v' \in \mathcal{C}_2$ with*

$$\|v - v'\| \leq K_1 \varepsilon \big(\|v - w\|^2 + K_2 \mathrm{e}^{K_3 \|v - w\|}\big).$$

Proof. Let f and h be parametrizations of $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively, in the direction of N. Let w be the origin. Then f and h are solutions to

$$f''(x) = F\big(x, f(x), f'(x)\big), \qquad h''(x) = H\big(x, h(x), h'(x)\big) \tag{21}$$

with $f(0) = h(0) = 0$ and $f'(0) = h'(0) = 0$ for Lipschitz continuous F and H with Lipschitz constants K_F and K_H, respectively. By Lemma B.1, the right hand side of (12) is Lipschitz continuous in δ for $(v, N_v) \in U$ with some Lipschitz constant K. The condition that $D_2(w) \cap H_{\varepsilon_2}(N) \subseteq D_1(w) \cap H_{\varepsilon_1}(N) \subseteq D_2(w) \cap H_{\varepsilon_2}(N) + B_\varepsilon(0)$ implies $0 \leq F(x, d, v) - H(x, d, v) \leq K\sqrt{|Y|}\varepsilon$, hence integrating (21) yields

$$\begin{aligned} f'(x) - h'(x) &= \int_0^x F\big(t, f(t), f'(t)\big) - H\big(t, f(t), f'(t)\big) \\ &\qquad + H\big(t, f(t), f'(t)\big) - H\big(t, h(t), h'(t)\big)\,\mathrm{d}t \\ &\leq K\sqrt{|Y|}\varepsilon x + K_H \int_0^x |f(t) - h(t)| + |f'(t) - h'(t)|\,\mathrm{d}t \end{aligned}$$

Since $H(x, d, v) \leq F(x, d, v)$ in a neighbourhood of 0, we may assume $f'(x) > h'(x)$ and $f(x) > h(x)$ by choosing U small enough. Therefore, $f - h$ satisfies the conditions of Theorem 1.8.1 in Pachpatte [18], which implies that

$$f'(x) - h'(x) \leq K\sqrt{|Y|}\varepsilon\left(x + K_F \int_0^x t + \frac{t^2}{2} + \frac{1}{8K_F^3}\mathrm{e}^{2K_F t}\,\mathrm{d}t\right).$$

Let $c_1 = 2K_F \vee 1$ and $c_2 = c_1\big(K_F + 1/(8K_F^2)\big)$. Using the inequality $t + t^2/2 \leq \mathrm{e}^t$, we obtain $f'(x) - h'(x) \leq K\sqrt{|Y|}\varepsilon(x + c_2\mathrm{e}^{c_1 x})$. Integrating once yields

$$f(x) - h(x) \leq K\sqrt{|Y|}\varepsilon\left(\frac{x^2}{2} + \frac{c_2}{c_1}\mathrm{e}^{c_1 x}\right).$$

For $v = \big(x, h(x)\big)$, let $v' = \big(x, f(x)\big)$, hence the result follows from $x \leq \|v - w\|$. $\square$

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] Dilip Abreu and Yuliy 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] J.-P. Aubin and H. Frankowska: *Set-Valued Analysis*, Springer, Berlin, 1990

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

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

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

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

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

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

[13] 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

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

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

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

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

[18] B.G. Pachpatte: *Inequalities for Differential and Integral Equations*, San Diego: Academic Press, 1998

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

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

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

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

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