CONTINUOUS-TIME GAMES WITH IMPERFECT AND ABRUPT INFORMATION*

Benjamin Bernard†

Forthcoming in the Review of Economic Studies

This paper studies two-player games in continuous time with imperfect public monitoring, in which information may arrive both gradually and continuously, governed by a Brownian motion, and abruptly and discontinuously, according to Poisson processes. For this general class of two-player games, we characterize the equilibrium payoff set via a convergent sequence of differential equations. The differential equations characterize the optimal trade-off between value burnt through incentives related to Poisson information and the noisiness of incentives related to Brownian information. In the presence of abrupt information, the boundary of the equilibrium payoff set may not be smooth outside the set of static Nash payoffs. Equilibrium strategies that attain extremal payoff pairs as well as their intertemporal incentives are elicitable from the limiting solution. The characterization is new even when information arrives through Poisson events only.

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

1 Introduction

A major goal of the theory of repeated games is to explain how the structure of intertemporal incentives in extremal equilibria depends on features of the underlying information structure. For two-player games with public monitoring, there are two qualitative features of monitoring structures whose impact on the optimal provision of incentives is reasonably well understood: (i) whether deviations by a player from a given profile can be statistically distinguished from deviations by the other player and (ii) how much overall noise is involved in the latter inference. The literature on discrete-time repeated games has shed a lot of light on the impact of (i) and (ii). However, monitoring structures can also differ in other important attributes, chief among them the nature of the *flow of information*: whether information arrives gradually and continuously over time or

*I would like to thank Christoph Frei, Fabien Gensbittel, Tilman Klumpp, Yuliy Sannikov, participants of the 11th World Congress of the Econometric Society, the 5th World Congress of the Game Theory Society, the Asian Meeting of the Econometric Society 2016, the 2017 Midwest Economic Theory Conference, and seminar participants at the University of Amsterdam, the University of Bern, the University of Bonn, Carnegie Mellon University, National Taiwan University, and Academia Sinica for valuable comments and interesting discussions. Financial support by the Natural Sciences and Engineering Research Council of Canada through grant RGPIN/402585-2011, the Swiss National Science Foundation through grant P2SKP1_171737, the Ministry of Science and Technology in Taiwan through grant 108-2410-H-002-249, and the Center for Research in Econometric Theory and Applications by the Ministry of Education in Taiwan through grant 109L900203 is gratefully acknowledged.

†Department of Economics, University of Wisconsin–Madison, 1180 Observatory Dr, Madison, WI 53706; bbernard3@wisc.edu.

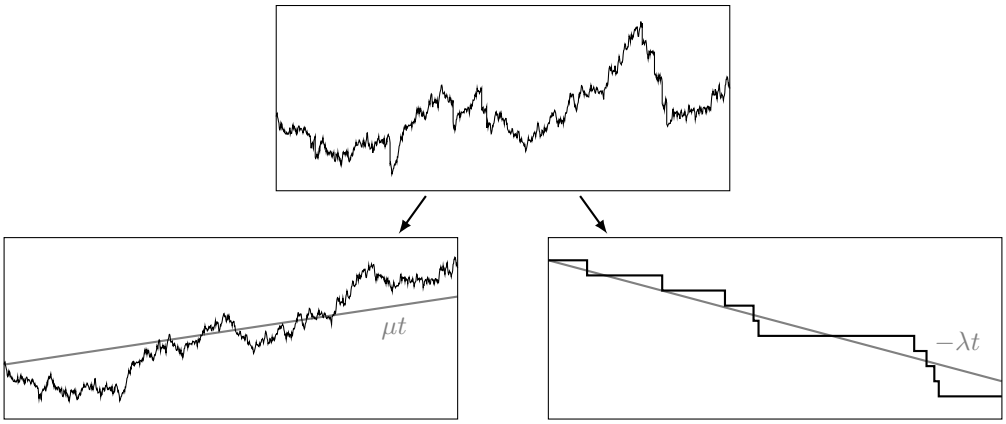

Figure 1: The total demand for the product (top panel) responds to normal market fluctuations (bottom left panel) and shocks due to negative media coverage (bottom right panel). Both components carry important information about past play: the expected rise in continuous revenue μ and the frequency λ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

abruptly and discontinuously. In order to study the impact of the flow of information on the structure of intertemporal incentives, it is natural to consider a continuous-time model so that information may arrive at *any* time and not just at multiples of a fixed time interval.

We model information that arrives gradually by a drifted Brownian motion: the players' actions determine the expected change of the signal, but the signal is distorted by Brownian noise. Players receive information intermittently when a discrete event occurs, whose frequency depends on the chosen actions. Arrival of these events is modeled through Poisson processes, hence these events cannot be anticipated by the players.[1] In many applications, there is a clear distinction between the two types of information arrival. Consider, for example, a partnership between two firms, in which the firms lack the expertise to directly evaluate the quality of the component contributed by the partnering firm. Each firm has an incentive to produce a lower-quality component at a lower cost as it may do so undetectedly. The partnering firm only receives noisy information from the performance of the product on the market. Demand moves continuously due to day-to-day fluctuations and it is subject to demand shocks when the product receives negative media coverage. As illustrated in Figure 1, a decomposition of the total demand leads to two separate signals that are both indicative of the quality of the partner's component: the continuous increase in demand without the impact of demand shocks and the frequency at which those demand shocks occur.

The two types of information arrival are able to sustain different kinds of intertemporal incentives on the efficient frontier of the equilibrium payoff set. Abreu, Pearce, and Stacchetti (1986) have shown that mutual punishments after undesirable signal realizations as in Green and Porter (1984) can be an efficient way of providing equilibrium incentives for impatient players. These incentives, to which we will refer as *value burning*, cannot be sustained on the efficient frontier using only continuous information: because Brownian information is so noisy, the chance of wrongfully entering a punishment phase is prohibitively large. Burning value after undesirable outcomes of Brownian

[1]While the lack of events may provide information about past actions, the memoryless property of the exponential distribution implies that the absence of events provides no information about when the next event will occur.

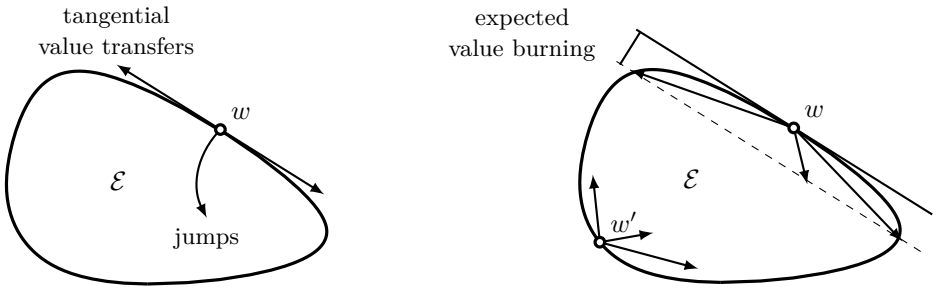

Figure 2: Because Brownian information is used to transfer value between players along tangents, only local information about the geometry of the equilibrium payoff set $\mathcal{E}$ is needed to structure incentives at a payoff pair w on the boundary. In contrast, abrupt information may be used to transfer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

information thus destroys too much value in expectation.[2] Brownian information is therefore used only to transfer value tangentially between players as in Fudenberg, Levine, and Maskin (1994). The use of Poisson information is not restricted in this way and can be used to transfer value tangentially and/or to burn value. On the efficient frontier, value burning corresponds to an inward movement and value transfers to a tangential movement of the continuation value. For ease of reference, we say that value is burnt/transferred whenever the continuation moves inwards/tangentially also for inefficient extremal equilibrium payoffs. See Figure 2 for an illustration.

Our main result characterizes the perfect public equilibrium (PPE) payoff set for an arbitrary discount rate. Its boundary is characterized through a linear program and two ordinary differential equations (ODEs), depending on which kind of information is needed to provide incentives:

(i) At so-called *stationary payoffs*, it is sufficient for players to condition their actions on the occurrence of discrete events alone. Stationary payoffs are characterized by a linear program similar to discrete time.

(ii) If, additionally, players must condition on the absence of discrete events, the boundary is characterized by a first-order ODE, called the *abrupt-information optimality equation*.

(iii) If players must use the continuous signal as well, the boundary is characterized by a second-order ODE that extends the *optimality equation* in Sannikov (2007).

The similarity to discrete-time games in (i) stems from the fact that Poisson events are inherently discrete. While the lack of discrete events or the observation of a continuous signal reveals an infinitesimal amount of information in each instant, the occurrence of a discrete event reveals a positive amount—the same as any signal in a discrete-time setting. Thus, stationary payoffs are characterized through a fixed point of a linear program similar to Abreu, Pearce, and Stacchetti (1990). We call those payoffs stationary because they can be attained by a PPE whose continuation value is locally constant until the arrival of the next event. If players do not condition on the absence of events, the strategy profile and the provided incentives cannot depend on time, i.e., they are constant until the next event occurs. Examples of stationary payoffs are static Nash payoffs

[2]Sannikov and Skrzypacz (2010) provide an excellent discussion of this point.

and, in the partnership game, payoffs that can be attained by a grim-trigger profile.

If players also condition on the absence of events as in (ii), strategy profiles and incentives may change over time. Because an event occurs with probability 0 at any given instant, rewards or punishments related to the absence of events must be deterministic and infinitesimal, i.e., the continuation value is absolutely continuous with respect to time. This gives rise to the novel abrupt-information optimality equation, which minimizes the amount of value that has to be burnt to provide incentives over all enforceable action profiles. Solutions to this ODE are differentiable, hence the boundary of the equilibrium payoff set is smooth, except at stationary payoffs and at payoffs, where the optimal action profile is enforceable only by extremal punishments/rewards. Let us illustrate the latter set of payoffs in the partnership example: exerting effort can be enforced by attaching a minimal punishment of, say, x after bad news arrives. Thus, effort by player i can be enforced credibly only if the current continuation value is at least x payoff units above player i's minimal equilibrium payoff $\underline{w}^i$. We say that incentives at payoff pairs w with $w^i = \underline{w}^i + x$ are *binding and extremal*—no lower punishment enforces effort and no larger punishment is credible. Where incentives are binding and extremal, the optimal action profile and the necessary incentives may change discontinuously, in which case the boundary of the equilibrium payoff set may not be smooth.

Finally, if incentives through the Brownian signal are required as in (iii), the volatility of the tangential transfers makes the boundary smooth. While the use of tangential transfers is efficient in the limit as players get arbitrarily patient, their use may be inefficient for impatient players because Brownian information is very noisy. The optimality equation trades off value burning with the reduction of volatility of the tangential transfers. The two ODEs are related in the sense the abrupt-information optimality equation is the limit of the optimality equation as the tangential transfers go to zero. Despite this limit property, the techniques used in the derivation of the two ODEs are quite different since the smoothness property from the Brownian signal—on which the techniques in Sannikov (2007) build—is missing for the abrupt-information optimality equation.

The equilibrium payoff set has corners only at stationary payoffs and at payoffs, where the maximal/minimal credible incentives bind for at least one player. This property is reminiscent of the discrete-time PPE payoff set with perfect monitoring characterized in Abreu and Sannikov (2014). Unlike in their setting, the PPE payoff set in the current framework does not have to be a polygon even when the Brownian signal is completely uninformative; see Section 4.5 for an example. In addition to the quantitative characterization, we can use our main result to gain qualitative insights about the trade-offs between the two kinds of information arrival. The ability to burn value upon the arrival of a discrete event may give impatient players the option to enforce a larger set of action profiles than might be possible with Brownian information exclusively. Thus, Poisson information generates a "core component" of equilibrium payoffs, which is then enlarged through Brownian information by exchanging some of the value that has to be burnt with tangential transfers. This is formalized in a result as follows: if the PPE payoff set $\mathcal{E}_P$ with Poisson information exclusively is a polygon, then the extremal payoffs of $\mathcal{E}_P$ are still extremal in the game with both kinds of information as long as the signal-to-noise ratio of the Brownian signal is below a certain threshold.

If information is entirely Brownian, Sannikov (2007) shows that the optimality equation is an *explicit* ODE, expressing the curvature of the equilibrium payoff set in terms of the unique equilibrium incentives at that point. Crucial for the explicit form is that equilibrium incentives at the boundary consist of tangential value transfers only and, in particular, they depend on the geometry of the equilibrium payoff set only through the direction of the tangent. Therefore, only local information about the boundary of the equilibrium payoff set is needed to characterize the set of equilibrium payoffs, in the same way that any ODE uses only local information of a function to encode its global properties. In contrast, when information arrives according to Poisson processes, any punishments or rewards are consistent with equilibrium behavior as long as the continuation value remains within the equilibrium payoff set. This restriction on incentives, however, depends on the entire geometry of the equilibrium payoff set. To describe the boundary locally through equilibrium incentives, global information about the set of equilibrium payoffs is required. A differential equation describing the boundary of the equilibrium payoff set via equilibrium incentives is thus self-referential.

We show that this self-referential differential equation can be approximated by a convergent sequence of explicit ODEs that are obtained from an iterative procedure over the arrival times of discrete events. This requires the introduction of the new concept of *relaxed self-generating* payoff sets $\mathcal{B}(\mathcal{W})$, which are self-generating only up to the time of the first discrete event, at which time the continuation value has to come from some fixed payoff set $\mathcal{W}$. 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. Because we iterate over the occurrence of discrete events, this algorithm is similar in spirit to the discrete-time algorithm in Abreu et al. (1990). Contrary to the APS operator, the fact that $\mathcal{B}(\mathcal{W})$ is relaxed self-generation before the first event yields a differential description of its boundary through the players' incentive constraints. Consequently, the sets in each step of the continuous-time algorithm can be computed efficiently as the numerical solution to an explicit ODE. This paper thus also contributes to the literature on computing equilibrium payoffs, providing an alternative to Judd, Şevin Yeltekin, and Conklin (2003) for two-player games and an extension of Abreu and Sannikov (2014) to imperfect monitoring.

The construction of equilibria through relaxed self-generating payoff sets does not crucially rely on the fact that information arrives according to Poisson processes. At a fundamental level, the developed methodology works for any purely discontinuous information process with totally inaccessible jump times, i.e., any information that is "abrupt" in the sense that its arrival cannot be anticipated by the players.[3] We restrict abrupt information to be Poissonian in this paper for notational simplicity and because they arise naturally in many context. Nevertheless, the techniques could be applied to other abrupt information processes such as, say, gamma processes or other Lévy processes at the cost of additional notational complexity.

In the partnership example, we compare numerically the share of total incentives that are provided through Poisson information for various monitoring structures. Because punishments/rewards after a discrete event lead to a discontinuous change in the continuation value, the continuation value lies

[3]A stopping time is totally inaccessible if it cannot be approximated by a sequence of predictable times.

below the tangent where the boundary of the equilibrium payoff set is curved; see the right panel of Figure 2. Thus, where the boundary is curved, incentives provided through abrupt information necessarily burn a positive amount of value. As a result, the continuation value after a discrete event is chosen close to tangentially to reduce the amount of value burnt—unless enforceability requires a minimum value be burnt for both players. This is reminiscent of the proof of the folk theorem: to attain asymptotic efficiency, players cannot rely on incentives that burn *any* value. This paper does not prove a folk theorem but instead characterizes the equilibrium payoff set for any positive discount rate. Therefore, some value burning is possible even in extremal equilibria, but used sparingly unless incentives cannot be provided through value transfers.

While the information structure in this paper is identical to Sannikov and Skrzypacz (2010), the focus, the techniques, and the results of the two papers are rather different. In a discrete-time framework with increasingly frequent actions, Sannikov and Skrzypacz (2010) establish a bound for the equilibrium payoff set, which is generically tight as players get arbitrarily patient. Their focus is on the generality of monitoring structures at the cost of an asymptotic result. The limit argument partially obscures the trade-off between the noisiness of Brownian information and value burning because infinitely patient players do not care about noise in a signal and they can be incentivized by an infinitesimal amount of value burnt.[4] The main result of this paper characterizes the PPE payoff set in a continuous-time game for *impatient* players—at the cost of assumptions that guarantee that the optimality equation is well behaved.[5] Under these assumptions, the payoff bound in Sannikov and Skrzypacz (2010) is trivial: it is simply equal to the convex hull of the stage-game payoffs of all enforceable action profiles. Moreover, in the limit as players get patient, the players could rely on tangential transfers through Brownian information alone, similarly to the folk theorem. Nevertheless, our results show that for impatient players, abrupt information may have a qualitative impact on equilibrium payoffs: the destruction of value may be the only way to credibly enforce a given action profile when the Brownian signal is too noisy.

The impact of information arrival on equilibrium payoffs has been analyzed by several papers using a discrete-time framework with increasingly frequent actions. Abreu, Milgrom, and Pearce (1991) demonstrate that intertemporal incentives may be destroyed in games with Poisson signals if the signals are not sufficiently informative. Sannikov and Skrzypacz (2007) show that intertemporal incentives are destroyed in a Cournot duopoly with Brownian information. In a model where players observe only cumulative signals, Fudenberg and Levine (2007) demonstrate that Brownian information supports a larger (smaller) set of equilibrium behavior than Poisson information if actions affect the volatility (drift) of the Brownian motion. Our paper deviates from this stream of literature by directly studying the continuous-time game. This allows us to address questions not only on asymptotic efficiency, but also on equilibrium behavior for any level of discounting: the

[4]The order of limits is important here. Sannikov and Skrzypacz (2010) characterize the limit $\lim_{r\to 0}\lim_{\Delta\to 0}\mathcal{E}(r,\Delta)$, where Δ is the length of the time period. Thus, information is fixed when the discount rate r is taken to 0.

[5]Unless the Brownian information is completely uninformative, the main result requires that the Brownian signal is pairwise identifiable, i.e., deviations by the two players can be statistically distinguished, and that any enforceable action profile can be enforced using Brownian information exclusively.

solution to the approximating sequence of ODEs reveals the unique action profiles and intertemporal incentives that are used at any given extremal equilibrium payoff. This allows us to show with relative ease that any equilibrium payoff in the partnership game with sufficiently informative demand shocks can be attained by play of one-sided effort, followed by a forgiving grim trigger profile or permanent Nash reversion. We can elicit equilibria in this way not just in the partnership game, but in any finite-action two-player game.

Two earlier working papers have analyzed continuous-time partnerships, in which players receive information about the chosen effort through separate Poisson signals. The main result in Kalesnik (2005) locally describes the boundary of the PPE payoff set under an endogenous condition that optimal rewards/punishments at the boundary are close to tangential. While this assumption is satisfied in the limit as the signal becomes arbitrarily informative and frequent, our examples show that it is not satisfied in general. Hauser and Hopenhayn (2008) analyze the exchange of favors between two players if the opportunities to provide favors arise according to Poisson events. They characterize the efficient frontier of the PPE payoff set and show that the efficient frontier is self-generating, i.e., efficient equilibria are renegotiation proof. In our paper, the partnership game is simply used as an illustrative example. Stage-game utilities and Poisson signals in this paper are arbitrary, and players may also observe Brownian signals. This paper covers all games from Kalesnik (2005), but not those by Hauser and Hopenhayn (2008) because the size of transfers upon the arrival of a Poisson event are bounded in their model by the utility of a single favor.

The remainder of the paper is organized as follows. Section 2 introduces the model and previews our results in the partnership example. We present the methodology to construct PPE in this framework in Section 3. Sections 4 and 5 explain how this methodology allows us characterize the boundary of the equilibrium payoff set. Section 4 does so in the more digestible partnership example, whereas Section 5 fills in many additional details that are needed in the general case. A description of how to implement the numerical solution of our algorithm is presented in Section 6 and we conclude in Section 7. The vast majority of the proofs are contained in Appendices A–C.

2 The model

Consider a game in which two players $i = 1, 2$ continuously choose actions from finite sets $\mathcal{A}^i$ at each point in time $t \in [0, \infty)$. The set of all pure action profiles $a = (a^1, a^2)$ is denoted by $\mathcal{A} = \mathcal{A}^1 \times \mathcal{A}^2$. Players cannot directly observe each other's actions and instead see only the impact of the chosen actions on the distribution of a public signal. The signal contains gradual and continuous information modeled by a d-dimensional Brownian motion Z, as well as abrupt and discontinuous information through the observation of infrequent events. There are finitely many (possibly zero) different types of infrequent events, indexed by $y \in Y$. Events arrive according to Poisson processes $(J^y)_{y \in Y}$ that are independent of each other and independent of the Brownian motion Z. An event of type y leads to a jump in the public signal of size $h(y)$ so that the public signal is given by $X = Z + \sum_{y \in Y} h(y)\, J^y$.

The public information at time t is a σ-algebra $\mathcal{F}_t$ that contains the history of the processes Z and $(J^y)_{y \in Y}$ up to time t, as well as orthogonal information that players may use for public

randomization. Because we study perfect public equilibria, a player's choice of action at time t must be based solely on information in $\mathcal{F}_t$, formalized in the following definition.

Definition 2.1. A *(public) pure strategy* A^i for player i is an $\mathcal{A}^i$-valued $(\mathcal{F}_t)_{t\geq 0}$-predictable process or the limit of a sequence of such processes.

Remark 2.1. Call a strategy A^i that is not predictable a *limit strategy*. Because jump times of Poisson processes are totally inaccessible, they cannot be approximated by a sequence of predictable times. In particular, limit strategies cannot anticipate discrete events. Because Brownian information arrives continuously, limit strategies differ from a predictable process only in that they may use public randomization in a non-predictable way. Allowing non-predictable public randomization simplifies the characterization because it ensures that the strategy space is closed; see Section 5.3 for details.[6] Non-predictable public randomization is standard in continuous-time models with Brownian information, where strategies are only required to be progressively measurable.

The game primitives $\mu : \mathcal{A} \to \mathbb{R}^d$ and $\lambda(y \,|\, \cdot\,) : \mathcal{A} \to (0, \infty)$ determine the impact of a chosen action profile on the drift rate of the continuous signal and the intensity of events of type $y \in Y$, respectively. We denote by $\lambda(a) := \big(\lambda(y_1 \,|\, a), \dots, \lambda(y_{|Y|} \,|\, a)\big)^\top$ the column vector of all intensities. As is standard, we model the impact of the chosen actions on the distribution of the public signal through a change of probability measure. Specifically, play of a strategy profile $A = (A^1, A^2)$ induces a family of probability measures $Q^A = (Q^A_t)_{t\geq 0}$, such that the public signal takes the form

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

on $[0, T]$ under Q^A_T for any $T > 0$. Here $Z^A = Z - \int \mu(A_s)\, \mathrm{d}s$ is a Q^A_T-Brownian motion describing noise in the continuous component and J^y has instantaneous intensity $\lambda(y \,|\, A)$ under Q^A_T.[7]

Remark 2.2. We normalize the covariance matrix to be the identity matrix. With the techniques in this paper, it is possible to study 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 a fixed 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 $X = \sigma\tilde{X}$.[8]

[6] In general, the limit of a sequence of predictable processes need not be predictable; see Remark VI.5.14 in Jacod and Shiryaev (2002) for a counterexample.

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

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

[8] Contrary to discrete-time models with frequent actions (c.f. Fudenberg and Levine (2007)), the volatility of the

Remark 2.3. Anderson (1984) and Simon and Stinchcombe (1989) demonstrate that seemingly simple strategies need not necessarily lead to unique outcomes in continuous-time games of perfect monitoring. The problem arises when strategies depend on actions chosen in the immediate past by either player. This is not a problem in our model because we restrict attention to public strategies. Indeed, one can identify the probability space with the path space of all publicly observable processes, hence a realized path $\omega \in \Omega$ leads to the unique outcome $A(\omega)$.

Definition 2.2. Let $r > 0$ be a discount rate common to both players. Each player i receives an unobservable expected flow payoff $g^i : \mathcal{A} \to \mathbb{R}$.[9]

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

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

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

$$W^i(A) \geq W^i(\tilde{A}^i, A^{-i}) \text{ a.e.,}^{10}$$

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

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

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

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 payoffs $\mathcal{V} := \operatorname{conv}\{g(a) \mid a \in \mathcal{A}\}$ is thus given by the convex hull of stage-game payoffs. Because we restrict attention to pure strategies, by deviating to their strategy of myopic best responses, each player i can ensure that their payoff in equilibrium dominates their pure-action 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 stage-game Nash equilibria

continuous 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.

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

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

	0	1
0	0, 0	4, −1
1	−1, 4	2, 2

$g(a)$

	0	1
0	21	17
1	17	1

$\lambda_1(a)$

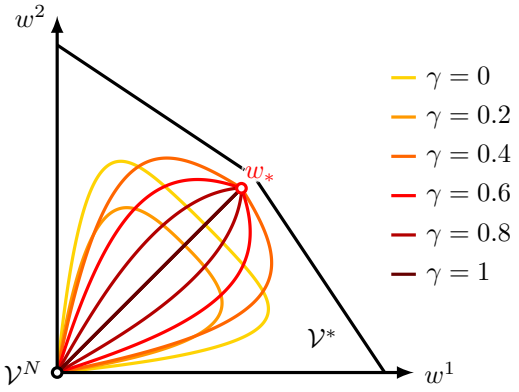

Figure 3: The left panel shows the matrix of stage game payoffs and the intensity of natural disasters. The right panel shows the equilibrium payoff set $\mathcal{E}_\gamma(5)$ as a subset of $\mathcal{V}^*$ for γ in increments of 0.2.

and denote by $\mathcal{V}^N := \operatorname{conv}\{g(a) \mid a \in \mathcal{A}^N\}$ 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 $\mathbb{E}[W_0(A)]$ and $\mathbb{E}[W_0(A')]$, respectively, any payoff pair $\nu\mathbb{E}[W_0(A)] + (1-\nu)\mathbb{E}[W_0(A')]$ 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.

2.1 Example of a partnership game

Consider a partnership game between two players, where each player continuously chooses a binary effort level from the set $\mathcal{A}^i = \{0,1\}$. Players enjoy output but dislike effort, captured by the expected flow payoffs $g^i(a) = 4(a^1 + a^2) - a^1a^2 - 5a^i$. For specificity, suppose that the partnership is a climate agreement between two countries that exert effort to reduce greenhouse gas emissions. While country i's effort is not directly observable, it affects the distribution of a publicly observable climate index $\mathrm{d}X_t^i = \mu_\gamma^i(A_t)\,\mathrm{d}t + \mathrm{d}Z_t^A$, where $\mu_\gamma^i(a) = 2.4(1-\gamma)a^i$. Additionally, there may be natural disaster as a consequence of climate change that occur at rate $\lambda_\gamma(a) = \gamma(21 - 4(a^1 + a^2) - 12a^1a^2)$. The parameter γ captures the relative informativeness of the two kinds of information arrival. For low values of γ, the climate indices are at their most informative, whereas natural disasters become more frequent and more informative as γ increases. Alternatively, one could envision a partnership between two firms, where X^i is firm i's stream of revenue and the discrete events are coverage of the product in the news. Figure 3 shows g and λ_1 in matrix form.

Figure 3 shows the equilibrium payoff set for discount rate $r = 5$ and different levels of relative informativeness between the two types of signals. As the figure shows, neither type of information arrival dominates the other in terms of impact on equilibrium payoffs, but they serve different purposes. If natural disasters are relatively uninformative ($\gamma < \frac{1}{3}$), mutual effort cannot be credibly enforced, hence the PPE payoff set is entirely contained below the negative diagonal that connects the stage-game payoffs of one-sided effort $g(0,1)$ and $g(1,0)$.[11] The presence of (sufficiently informative)

[11] Sannikov and Skrzypacz (2010) show that in a partnership game with one-dimensional Brownian monitoring, mutual effort is not enforceable, hence $\mathcal{E}(r)$ lies below this diagonal for any discount rate r. The mechanism in our

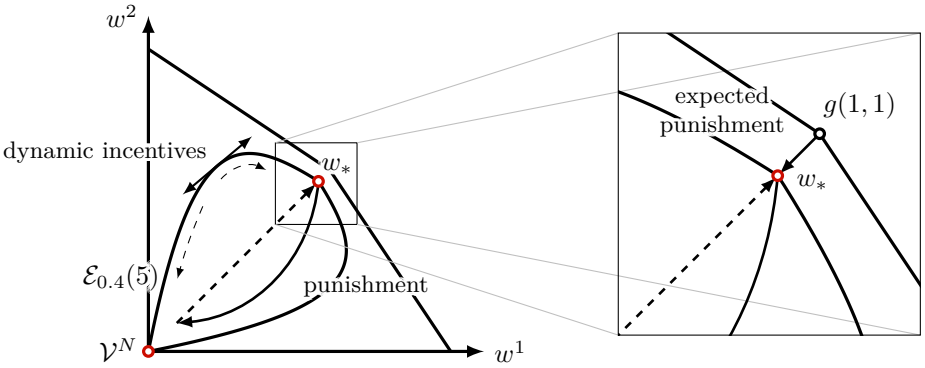

Figure 4: The equilibrium payoff set $\mathcal{E}_{0.4}(5)$ is shown together with the path of the continuation value of a forgiving grim-trigger profile attaining w_*. The solid arrow indicates the minimal amount of value that has to be burnt to enforce mutual effort. As the punishment phase progresses, the countries' continuation value increases along the dashed arrow until it is equal to w_* again. The zoom-in in shows that w_* lies below the stage-game payoff pair $g(1,1)$ precisely by the minimum expected value burnt by the punishment. On curved parts of the boundary, one-sided effort is enforced mostly by rewarding the country for high realizations of the climate indices. The dashed arrows indicate the drift rate of the continuation value.

natural disasters thus enlarges the set of enforceable action profiles for $r = 5$, whereas the climate indices serve to bulge out the equilibrium payoff set and make it fuller. This qualitative difference between the two kinds of information holds more generally in a sense formalized by Lemma 5.12.

The easiest way to provide incentives for effort is through a forgiving grim trigger profile, in which the countries temporarily suspend their effort after the occurrence of a natural disaster. If natural disasters are sufficiently informative relative to the players' patience ($\gamma \geq \frac{1}{3}$), mutual effort can be enforced. The highest symmetric payoff pair w_* is bounded away from the efficient frontier by the expected punishment that is required to enforce effort; see Figure 4. The natural disasters provide a second kind of information: their absence is good news, which can be used to reward the countries. In a discrete-time setting, the two kinds of information are symmetric in the sense that both are positive-probability events. In continuous time, a disaster occurs at any given time with probability 0, hence any rewards attached to the absence of disasters must be deterministic and infinitesimal. Suppose we try to attain asymmetric equilibria, in which some country i starts by exerting effort alone. Because effort is costly and the maximal credible reward for exerting effort is bounded, country i can be incentivized to exert effort only within a deterministic time frame before a change in strategy must occur. For discount rate $r = 5$, this time frame is too short to provide meaningful incentives and the equilibrium payoff set $\mathcal{E}_1(r)$, when only Poisson information is informative, is contained entirely on the diagonal. In particular, every equilibrium payoff pair w is attainable by a forgiving grim trigger profile (starting with punishment if w is in the relative interior). Section 4.5 shows an example, in which deviations by the effort-exerting country are easier to detect and asymmetric equilibria exist in the Poisson-only case.

Brownian information allows non-deterministic incentive provision, as long as it is not completely uninformative. As a consequence, a country can be dynamically rewarded for one-sided effort through

example differs because mutual effort is, in principle, enforceable through the pairwise identifiable Brownian signal. However, for impatient players the signal involves too much noise to be used efficiently.

the performance of its climate index. Specifically, the length of the period, during which country i has to exert effort alone, can be shortened if its climate index X^i increases, and it can be extended otherwise. Thus, if the length of the time period is adjusted appropriately, during which one country exerts effort alone, one-sided effort can be part of an equilibrium with Brownian information. Figure 3 shows that the equilibrium payoff set gets wider as the Brownian information gets more informative. For discount rate $r = 5$, however, the Brownian signal is not informative enough to enforce mutual effort. Figure 3 may give a false sense that the shape of $\mathcal{E}_\gamma(5)$ changes monotonically in γ. For this specific parametrization, $\mathcal{E}_{0.3}(r)$ coincides with the set of static Nash payoffs: neither gradual nor abrupt information are sufficiently informative to provide any intertemporal incentives.

3 Construction of equilibria

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

Lemma 3.1. *For a two-dimensional process W and a pure strategy profile A, equivalent are:*

(i) W is the continuation value of A

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

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

for a martingale M^i (strongly) orthogonal to Z and all J^y with $M_0^i = 0$ as well as predictable and square-integrable processes β^i and $\delta^i(y)$ for $y \in Y$.[12]

The first term in (2) describes the expected movement of the continuation value. It points away from the expected flow payoff: if $g^i(A) > W^i$, then player i extracts an instantaneous payoff rate that exceeds their continuation value and, therefore, doing so has to decrease their future payoff in expectation. The second and third term in (2) characterize the exposure of the continuation value to the public signal. The exposure to the continuous component is linear in the difference between the realized change $\mathrm{d}Z_t$ of the signal and its expected change $\mu(A)\,\mathrm{d}t$, captured by the parameter β^i. Arrival of abrupt information changes the continuation value discontinuously, where $r\delta^i(y)$ are the impacts on player i's continuation value when an event of type y occurs. The absence of an event of type y impacts player i's continuation value continuously and in the opposite direction through the term $-r\delta^i(y)\lambda(y\,|\,A)\,\mathrm{d}t$. The second and third term are martingale increments, which means that their impact on the continuation value averages out in expectation. Nevertheless, the exposure to the public signal is important for the provision of incentives as we shall see below. Finally, M captures moves in the continuation value due to public randomization: they average out in expectation and M is 0 if players do not use public randomization.

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

In discrete-time games, incentives are provided by a continuation promise that maps the outcome of the public signal to a promised continuation payoff for every player; see, for example, Abreu et al. (1990). The representation in (2) shows that in continuous-time games, such a continuation promise has to be provided via the sensitivities β and δ to the public signal.

Definition 3.2. An action profile $a \in \mathcal{A}$ is *enforceable* if there exists a *continuation promise* (β, δ) with $\beta = (\beta^1, \beta^2)^\top \in \mathbb{R}^{2\times d}$ and $\delta = (\delta^1, \delta^2)^\top \in \mathbb{R}^{2\times m}$ such that for every player i, and every deviation $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$,

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

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

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

Lemma 3.3. *A strategy profile A 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 3.1 and 3.3 motivate how we construct equilibrium profiles in continuous time—as the solution to (2) subject to the enforceability constraint in (3). To do so, we use the fact that repeated games are time homogeneous: since the continuation profile of a PPE after any time t is also an equilibrium of the entire game, the continuation value has to remain within the equilibrium payoff set at all times. This property is known as self generation. In our setting, it is formalized as follows.

Definition 3.4. A payoff set $\mathcal{W} \subset \mathbb{R}^2$ is called *self-generating* if for every payoff pair $w \in \mathcal{W}$, there exists a solution (W, A, β, δ, M) to (2) such that (β, δ) enforces A a.e., $W_0 = w$ a.s., and $W_\tau(A) \in \mathcal{W}$ a.s. for every stopping time τ.

Remark 3.1. The stochastic differential equation (2) does not in general admit strong solutions, that is, it may not be possible to solve (2) for a fixed Brownian motion Z and a fixed set $(J^y)_{y\in Y}$ of Poisson processes. In the proofs that are contained in the appendices, we use weak solutions to (2), in which the Brownian motion, the Poisson process, and the probability space are part of the solution. To keep the notation simple, we do not make this distinction in the main text. See Appendix A.2 for details.

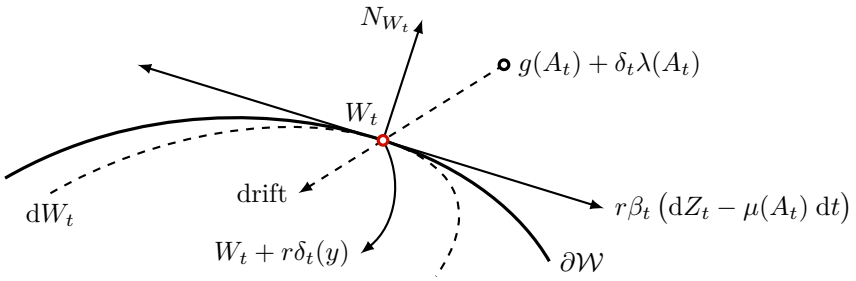

Figure 5: For continuation value W to remain in $\mathcal{W}$, it is necessary that intertemporal incentives satisfy informational restrictions (I1)–(I3) at the boundary.

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

Lemma 3.5. *The set $\mathcal{E}(r)$ is the largest bounded self-generating payoff set.*

3.2 Iteration over Poisson events

The characterization of the equilibrium payoff set as the largest bounded self-generating set allows us to construct equilibria using a stochastic control approach. Because the equilibrium payoff set is self-generating, the continuation value of a PPE has to remain within the set at all times. At the boundary, the law of motion given in (2) thus places certain restrictions on admissible continuation promises; see Figure 5. For ease of reference, let us summarize these informational restrictions here for an arbitrary compact, convex, self-generating payoff set $\mathcal{W}$.

For 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 in $\mathbb{R}^2$. 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 a boundary payoff pair w of a self-generating set $\mathcal{W}$ are the following:

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

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

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

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

When information is entirely Brownian, only local information about the boundary is necessary to describe the equilibrium payoff set. Crucial in this regard is that Brownian information arrives continuously, i.e., only the informational restrictions (I1) and (I2) are observed. These restrictions depend on the geometry of the equilibrium payoff set only through the normal vector N_w at w, which gives rise to an explicit description of the boundary through an ordinary differential equation in the state (w, N_w). When information arrives discontinuously as well, such a local description is no longer possible because (I3) is a global restriction involving the precise shape of the equilibrium

payoff set. A local description of $\mathcal{E}(r)$ using restrictions (I1)–(I3) thus involves $\mathcal{E}(r)$ itself, creating a non-trivial fixed-point problem. We solve it with an iteration over the arrival times of discrete events, where (I3) is relaxed so that continuation payoffs after an event come from a fixed payoff set $\mathcal{W}$.

Definition 3.6. Fix a payoff set $\mathcal{W} \subseteq \mathbb{R}^2$ and let σ denote the arrival time of the first event.

(i) We say that a solution (W, A, β, δ, M) to (2) is $\mathcal{W}$-*enforceable* if $(\beta_t(\omega), \delta_t(\omega))$ enforces $A_t(\omega)$ for almost every (ω, t) with $0 \leq t < \sigma(\omega)$ and $W_\sigma \in \mathcal{W}$ a.s.[13]

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

(iii) We denote by $\mathcal{B}(\mathcal{W})$ the largest bounded $\mathcal{W}$-relaxed self-generating payoff set.[14]

The operator $\mathcal{B}$ shares some similarity with the standard set operator in Abreu et al. (1990): payoff pairs in $\mathcal{B}(\mathcal{W})$ can be attained by an enforceable strategy profile with continuation promise W_σ in $\mathcal{W}$. The (strong) additional requirement in continuous time is that before a discrete event occurs, the continuation value has to remain in $\mathcal{B}(\mathcal{W})$. We can recover self-generating payoff sets from relaxed self-generating payoff sets with the following lemma.

Lemma 3.7. *Let $\mathcal{W} \subseteq \mathcal{V}$. If $\mathcal{W}$ is self-generating, then $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$. Conversely, if $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$, then $\mathcal{B}(\mathcal{W})$ is self-generating.*

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

Proposition 3.8. *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 $\bigcap_{n\geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*[15]

At first glance, the algorithm looks similar to the algorithm in Abreu et al. (1990) because the operator $\mathcal{B}$ serves to concatenate enforceable solutions at the arrival times of Poisson events. The similarity in the notation may be deceiving: while the incentives provided at the time of discrete events are the only incentives in a discrete-time setting, they are only a fraction of the incentives provided in continuous time. In our model, players need to be incentivized between the arrival of discrete events as well, subject to informational restrictions (I1)–(I3). As we argue in the next two

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

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

[15] If there are no pure-strategy PPE, then the algorithm described in Proposition 3.8 must converge to the empty set: if it converged to a non-empty fixed point of $\mathcal{B}$, that fixed point would be self-generating by Lemma 3.7, contradicting the fact that there are no pure-strategy PPE.

sections, those restriction give rise to an explicit characterization of the boundary in each step. To illustrate just how strong the additional incentive constraints are in each step of Proposition 3.8 compared to discrete time, note that $\mathcal{B}(\mathcal{V}^*) = \mathcal{E}(r)$ if information is entirely Brownian.

4 Characterization of $\partial\mathcal{B}(\mathcal{V}^*)$ in the partnership game

In this section, we motivate the key techniques of the equilibrium construction by illustrating how we characterize $\mathcal{B}(\mathcal{V}^*)$ in the partnership example of Section 2.1. The characterization of $\mathcal{B}(\mathcal{W})$ for general games and arbitrary compact, convex payoff sets $\mathcal{W}$ is deferred to Section 5. This separation allows us to avoid some technical complications in the present section. Since the definitions and results presented in this section are used in the formal proof, we state those for the general case. We construct $\mathcal{V}^*$-enforceable strategy profiles that rely on increasing informational complexity to provide incentives, starting with strategies that use punishments only when a natural disaster occurs, followed by strategies that also reward countries in the absence of natural disasters, and finally strategies that also take into account the performance of the climate indices.

4.1 Stationary payoffs

Let us start with strategy profiles that condition only on the occurrence of natural disasters. Denote by Ψ_a the set of all rewards/punishments δ that incentivize play of action profile a, i.e., those δ that solve (3) for both countries with $\beta = 0$. For example, the enforceability constraint for mutual effort is $2+\gamma\delta^i \geq 4+17\gamma\delta^i$, which is equivalent to $\delta^i \leq -\frac{1}{8\gamma}$ for $i = 1, 2$. To be incentivized to exert effort, both countries need to be punished if a disaster occurs, and the minimal strength of the punishment decreases with the informativeness γ of the disasters. If a continuation promise $(0, \delta)$ is kept and countries do not use public randomization, the continuation value evolves according to

$$\mathrm{d}W_t = r(W_t - g(A_t) - \delta_t\lambda(A_t))\,\mathrm{d}t + r\sum_{y\in Y}\delta_t(y)\,\mathrm{d}J_t^y. \tag{4}$$

Because jumps in the continuation value are proportional to the discount rate, the actual punishment needed to enforce mutual effort is $\frac{r}{8\gamma}$. The first term in (4) shows that the good news from the lack of natural disasters has a continuous impact on he continuation value: if $\delta_t^i \leq -\frac{1}{8\gamma}$, then $-r\delta_t\lambda(A_t)$ is positive, increasing the drift of the continuation value. If the drift rate is 0 and the continuation value W is locally constant, then even though the absence of natural disasters affects W through (4), it may not affect the strategy profile. Indeed, since the continuation value is locally constant, the same continuation strategy profile attains W_t for any $t < \sigma$. In particular, A can be chosen constant before the first natural disaster, hence countries condition only on the occurrence of these events. Note that $W_t = w$ implies $\delta_t = \delta_0$, i.e., the threatened punishments are constant until a disaster occurs. We call such payoff pairs stationary payoffs because the provided incentives, the strategies, and the continuation value are locally constant until the arrival of the next event. Static Nash payoffs are special cases of stationary payoffs: they are stationary for $\delta_0 = 0$.

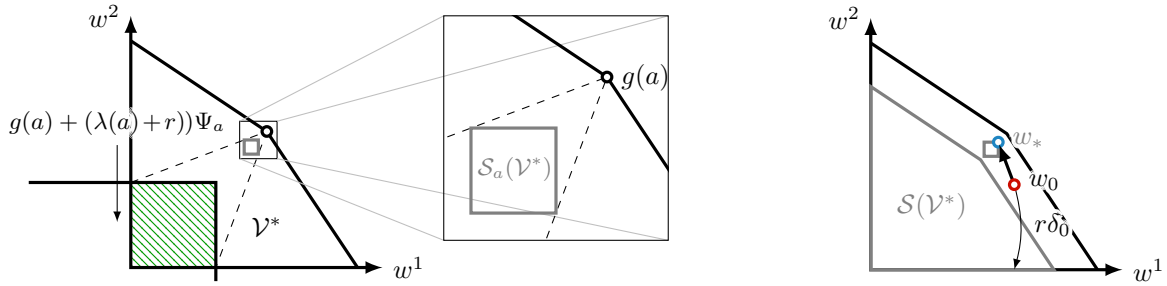

Figure 6: The left panel shows the construction of payoffs $\mathcal{S}_a(\mathcal{V}^*)$ that can be attained with local play of $a = (1, 1)$ in the partnership example for $r = 5$ and $\gamma = 1$. The right panel shows all $\mathcal{V}^*$-stationary payoffs, as well as a payoff pair w_0 that can be attained by local play of $(0, 1)$, enforced by δ_0. The drift of the continuation value points away from $g(0, 1) + \delta_0\lambda(0, 1) = (4, -5.25)$ and towards w_*.

Definition 4.1. A payoff pair w is $\mathcal{W}$-*stationary* for action profile a if there exists $\delta_0 \in \Psi_a$ such that $w = g(a) + \delta_0\lambda(a)$ and $w + r\delta_0(y) \in \mathcal{W}$ for every $y \in Y$. We denote by $\mathcal{S}_a(\mathcal{W})$ the set of all payoff pairs w that are $\mathcal{W}$-stationary for a and by $\mathcal{S}(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \mathcal{S}_a(\mathcal{W})$ the set of all $\mathcal{W}$-stationary payoffs.

The set of stationary payoffs can be computed fairly easily as intersections and scalings of convex sets. The stationarity condition implies that $w + r\delta_0 = g(a) + \delta_0\lambda(a) + r\delta_0$ must lie in $\mathcal{V}^*$ for some $\delta_0 \in \Psi_a$. Thus, the promised continuation values after a natural disaster come from the set

$$\big(g(a) + (\lambda(a) + r)\Psi_a\big) \cap \mathcal{V}^* \tag{5}$$

for some action profile a. Following from our computations above, for mutual effort this is the set of all those payoff pairs v with $0 \le v^i \le 2 - \frac{r+1}{8\gamma}$. The set $\mathcal{S}_a(\mathcal{V}^*)$ is then the expected value of these promised continuation payoffs and the current flow payoff. Geometrically, this corresponds to shrinking $\big(g(a) + (\lambda(a) + r)\Psi_a\big) \cap \mathcal{V}^*$ towards $g(a)$ by a factor of $\frac{\lambda(a)}{\lambda(a)+r}$; see the left panel of Figure 6. If the countries are promised a continuation value in the green shaded area (5) when a natural disaster occurs, they are willing to exert effort regardless of how long it takes for a natural disaster to occur. The strength of the punishment determines their expected payoff in $\mathcal{S}_a(\mathcal{V}^*)$.

Repeating this construction for all action profiles a yields the set of stationary payoffs. It turns out that there are no $\mathcal{V}^*$-stationary payoffs supported by one-sided effort in this partnership game, but there are $\mathcal{V}^*$-stationary payoffs supported by the static Nash profile; see the right panel of Figure 6. Because the continuation value is locally constant, the singleton $\{w\}$ for any stationary payoff w is self-generating up to the time of the first disaster, which implies $\operatorname{conv}\mathcal{S}(\mathcal{V}^*) \subseteq \mathcal{B}(\mathcal{V}^*)$.

4.2 Incentives through Poisson information

Suppose now that countries condition also on the lack of natural disasters, i.e., the drift rate in (4) is not necessarily zero. Let us aim to construct an enforceable strategy profile, in which country 2 is rewarded for being the sole country to exert effort when no disaster occurs. By (3), action profile $(0, 1)$ is enforced by $\delta^1 \ge -\frac{1}{8\gamma}$ and $\delta^2 \le -\frac{1}{4\gamma}$, imposing a minimal punishment for country 2 and a maximal punishment for country 1 when a disaster occurs. As a first attempt, consider a strategy profile A starting with action profile $(0, 1)$ enforced with minimal punishments $\delta \equiv \delta_0 = (0, -\frac{1}{4\gamma})$. Its

continuation value W, given by (4), has the explicit solution $W_t = w_0 + (\mathrm{e}^{rt} - 1)\big(w_0 - g(a) + \delta_0\lambda(a)\big)$, conditional on no disaster occurring. We observe that W remains on the straight line through w_0 and $g(a) + \delta_0\lambda(a)$, drifting away from $g(a) + \delta_0\lambda(a)$; see the right panel of Figure 6. If the path of W reaches some stationary payoff w_* before a disaster occurs, we can concatenate the solution with the stationary profile that attains w_*. Specifically, if $w_* \in \mathcal{S}_{1,1}(\mathcal{V}^*)$, there exists δ_* such that $w_* = g(1,1) + \delta_*\lambda(1,1)$ by stationarity, hence the concatenation of A and δ before the first event are

$$A_t = \begin{cases} (0,1) & \text{if } t < t_0, \\ (1,1) & \text{if } t \geq t_0, \end{cases} \qquad \delta_t = \begin{cases} \delta_0 & \text{if } t < t_0, \\ \delta_* & \text{if } t \geq t_0, \end{cases}$$

where t_0 is the deterministic time, at which w_* is reached, conditional that no disaster occurs. By construction, the continuation value after a disaster comes from $\mathcal{V}^*$ regardless of whether it occurs before or after w_* is reached. We have thus successfully attained such w_0 with a $\mathcal{V}^*$-enforceable strategy profile. Note that country 2 is rewarded and its continuation value increases if no natural disasters occur before t_0 because it receives the help of country 1 at time t_0. Conversely, country 1's continuation value decreases because its period of free-riding comes to an end at time t_0.

A quick inspection of the above argument reveals that we could have chosen different values for δ_0. While country 1 does not need to be incentivized to shirk, the choice of $\delta_0^1 = 0$ is inefficient. The total amount of value destroyed upon the occurrence of a natural disaster can be reduced if country 1 is rewarded. The efficient solution is to reward country 1 by the maximal amount in $\mathcal{V}^*$ that is consistent with country 2's punishment. Generally, the efficient δ will vary with time, but the two main steps in constructing enforceable strategy profiles are the same as above:

(i) Construct a $\mathcal{W}$-enforceable solution W to (2) that remains on some curve $\mathcal{C}$ until either an event occurs or one of the end points of $\mathcal{C}$ is reached.

(ii) If an end point v of $\mathcal{C}$ lies in $\mathcal{B}(\mathcal{W})$ and v is reached before an event occurs, we can concatenate the solution with a $\mathcal{W}$-enforceable solution that attains v.

We generalize the construction in two steps. First, we describe for which continuation values sufficient incentives can be provided through Poisson information exclusively. Second, we derive an ODE that characterizes the efficient use of incentives along such curves $\mathcal{C}$.

In the above example, not every payoff pair on the line through w_* and $g(0,1) + \delta_0\lambda(0,1)$ can be attained by an enforceable strategy profile in the same way. Because country 2 has to be punished by at least $\frac{r}{4\gamma}$ payoff units upon the occurrence of a disaster and the continuation value after a punishment is at least 0, this construction works only if $w^2 \geq \frac{r}{4\gamma}$. Denote by $\Psi_a(w, \mathcal{W}) := \{\delta \in \Psi_a \mid w + r\delta(y) \in \mathcal{W} \text{ for every } y \in Y\}$ the set of continuation promises in $\mathcal{W}$ that enforce action profile a. Sufficient incentives for action profile a can be provided only if the continuation value is in the set $\mathcal{K}_a(\mathcal{W}) := \{w \mid \Psi_a(w, \mathcal{W}) \neq \emptyset\}$. Let us compute the sets $\mathcal{K}_a(\mathcal{W})$ in the partnership example. Country 1 can be punished by at most $\frac{r}{8\gamma}$ payoff units in order to incentivize $(0,1)$. Since the highest-possible continuation value in $\mathcal{V}^*$ is $\frac{10}{3}$, it follows that $w^1 \leq \frac{10}{3} + \frac{r}{8\gamma}$. If w^1 is any higher, the punishment of any feasible and individually rational continuation value is too strong. At payoff

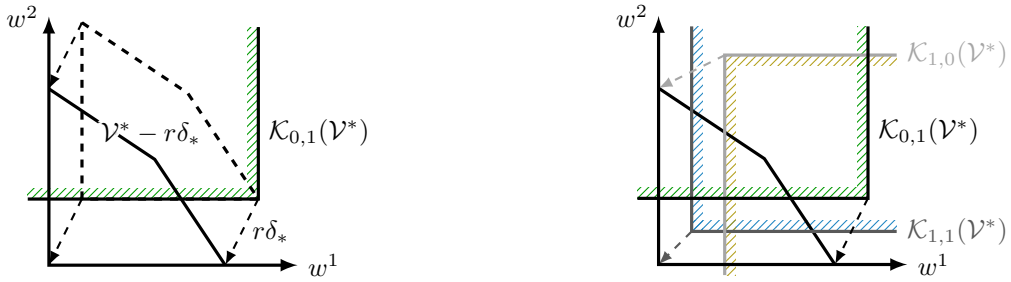

Figure 7: In the partnership game, mutual effort is enforceable by punishing both countries by at least $r/8\gamma$ payoff units. One-sided effort is enforceable by punishing the effort-exerting country by at least $r/4\gamma$ and punishing the shirking country by at most $r/8\gamma$. The arrows show these extremal punishments for $r = 5\gamma$. The sets $\mathcal{K}_a(\mathcal{V}^*)$ are infinite extensions of translates of $\mathcal{V}^*$ by those extremal punishments.

pairs w with $w^1 = \frac{10}{3} + \frac{r}{8\gamma}$, incentives bind for country 1 and they are extremal in $\mathcal{V}^*$. Carrying out the same argument for country 2, we obtain $\mathcal{K}_{0,1}(\mathcal{V}^*) = \left(-\infty, \frac{10}{3} + \frac{r}{8\gamma}\right] \times \left[\frac{r}{4\gamma}, \infty\right)$; see the left panel of Figure 7. A visual way of constructing $\mathcal{K}_a(\mathcal{V}^*)$ is to translate $\mathcal{V}^*$ by $-r\delta_*$, where δ_* are the extremal incentives that enforce a, and then extending $\mathcal{V}^* - r\delta_*$ towards $+\infty$ in dimension i if δ_*^i is a minimal punishment and towards $-\infty$ if δ_*^i is a maximal punishment; see Figure 7.

Our first result about the provision of incentives in absence of discrete events states that at the boundary of $\mathcal{B}(\mathcal{W})$, the drift must be tangential to $\mathcal{B}(\mathcal{W})$ unless incentives are both binding and extremal in $\mathcal{W}$. Incentives are both binding and extremal in the set $\mathcal{K}(\mathcal{W}) := \bigcup_{a\in\mathcal{A}} \partial\mathcal{K}_a(\mathcal{W})$.

Lemma 4.2. *For any payoff pair $w \in \partial\mathcal{B}(\mathcal{W}) \setminus \mathcal{K}(\mathcal{W})$ and any normal vector $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$, it is impossible that there exist a and $\delta \in \Psi_a(w, \mathcal{W})$ with $N^\top(g(a) + \delta\lambda(a) - w) > 0$.*

Proof. Suppose towards a contradiction that such a and δ exist. Lemma B.5 in the appendix establishes that the map $w \mapsto \Psi_a(w, \mathcal{W})$ is continuous, hence for any w' sufficiently close to w, there exists $\delta' \in \Psi_a(w', \mathcal{W})$ with $N^\top(g(a) + \delta'\lambda(a) - w') > 0$. Choose $w' \notin \operatorname{cl}\mathcal{B}(\mathcal{W})$ close enough to w such that the straight line segment through $x := g(a) + \delta'\lambda(a)$ and w' contains a point $v \in \operatorname{int}\mathcal{B}(\mathcal{W})$; see Figure 8. Let W be a solution to (2) starting at $W_0 = w'$ with $A \equiv a$, $\beta \equiv 0$, $\delta \equiv \delta'$, and $\mathrm{d}M_t = (v - w')(\mathrm{d}J'_t - \lambda' \,\mathrm{d}t)$, where J' is a Poisson process independent of the public signal with intensity $\lambda' = r\|w' - x\|/\|v - w'\|$. The public randomization device M replaces the gradual inward drift $r(W_t - g(a) - \delta'\lambda(a)) \,\mathrm{d}t$ with an occasional jump from w' to v. Thus, W remains in w' until it jumps either to v at the first jump time of J' or to $\mathcal{W}$ when an event y occurs. If the jump to v occurs first, then, because $v \in \mathcal{B}(\mathcal{W})$, there exists a $\mathcal{W}$-enforceable solution W^v to (2) that starts at v. The concatenation of W and W^v thus attains w' and remains in $\{w'\} \cup \mathcal{B}(\mathcal{W})$, showing that $\{w'\} \cup \mathcal{B}(\mathcal{W})$ is relaxed self-generating. This contradicts maximality of $\mathcal{B}(\mathcal{W})$. Lemma A.2 in the appendix contains the details to the concatenation procedure. □

Lemma 4.2 has two important implications. The first implication is that $\mathcal{K}(\mathcal{W})$ must contain all non-stationary corners of $\mathcal{B}(\mathcal{W})$. Indeed, if there are three distinct normal vectors N_1, N_2, N_3 at w, for which there exists $\delta \in \Psi_a(w, \mathcal{W})$ with ${N_k}^\top(g(a) + \delta\lambda(a) - w) = 0$, convexity of $\Psi_a(w, \mathcal{W})$ implies that $w \in g(a) + \Psi_a(w, \mathcal{W})\lambda(a)$, i.e., w is stationary; see the right panel of Figure 8. The second implication is that if incentives are provided trough Poisson information exclusively, the boundary

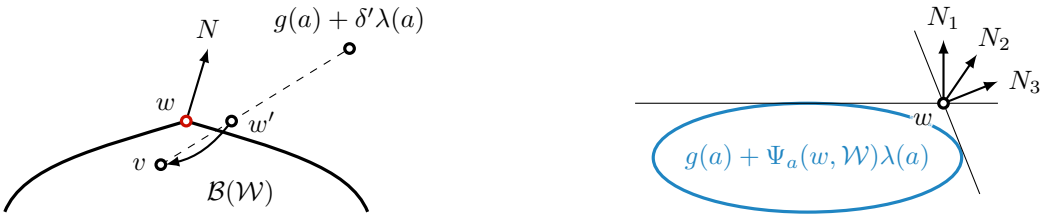

Figure 8: Payoff pair w' in the left panel can be attained by a locally constant solution to (2), in which the drift towards v is offset by a compensated Poisson public randomization device. The resulting solution remains at w' until it jumps to $\mathcal{W}$ at the time of an event or to v through public randomization. The right panel illustrates that convexity of $\Psi_a(w,\mathcal{W})$ and Lemma 4.2 imply that all corners of $\mathcal{B}(\mathcal{W})$ are in $\mathcal{K}(\mathcal{W})\cup\mathcal{S}(\mathcal{W})$.

must be a continuously differentiable solution to the *abrupt-information optimality equation*

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

at any payoff pair w outside of $\mathcal{K}(\mathcal{W})\cup\mathcal{S}(\mathcal{W})$. The maximization over δ corresponds to minimizing the expected value that has to be burnt in direction N_w to enforce any given action profile. The maximization over action profiles then selects the action profile that requires the least amount of value to be burnt. Equation (6) is an implicit first-order ODE, involving both the state w and its first derivative N_w. We discuss existence and uniqueness of solutions to (6) in Section 5.1. For now, let us anticipate one consequence of that discussion, which is that the maximizers a_w and δ_w on the right-hand side of (6) are generically unique unless the solution to (6) locally is a straight line segment.

Remark 4.1. One can bring (6) into a more standard form by considering a local parametrization f of a solution $\mathcal{C}$. For simplicity of notation, suppose that that parametrization is in the direction of player 2's payoffs so that $w\in\mathcal{C}$ is parametrized as $(w^1, f(w^1))$, but any other parametrization will work. Then the normal vector at w is $N_w = (-f'(w^1),1)^{\top}/\sqrt{1+(f'(w^1))^2}$ and (10) takes the form

$$f'(w^1) = \frac{f(w^1) - g^2(a_w) - \delta_w^2\lambda(a_w)}{w^1 - g^1(a_w) - \delta_w^1\lambda(a_w)}, \tag{7}$$

i.e., the slope of the curve at w is the slope of the line through $g(a_w)+\delta_w\lambda(a_w)$ and w.[16]

Let us return to the partnership game. The left panel in Figure 9 shows the set of stationary payoffs and a solution $\mathcal{C}$ to (6) starting at the stationary payoff w_* that maximizes the sum of the countries' payoffs. Along the entire segment, the abrupt-information optimality equation is maximized at the action profile $(0,1)$ of one-sided effort by country 2. The solution ends at $\hat{w}\in\partial\mathcal{K}_{0,1}(\mathcal{V}^*)$ because incentives cannot be provided through disasters alone outside of $\mathcal{K}_{0,1}(\mathcal{V}^*)$. Incentives $r\delta_w$ that maximize (6) at a payoff pair w minimize the expected value burnt. In this particular example, δ_w involves the minimal punishment $-r\delta_w^2 = \frac{r}{4\gamma}$ necessary to enforce effort by country 2 and the maximal reward $r\delta_w^1$ in $\mathcal{V}^*$ that is consistent with $\delta_w^2 = -\frac{1}{4\gamma}$. Optimal incentives δ_* and $\hat{\delta}$ at the end points w_* and $\hat{w}$ of $\mathcal{C}$ are shown in Figure 9. Note that $\mathcal{C}$ is strictly curved because $\delta_*^1 > \hat{\delta}^1$. As before, any payoff pair on $\mathcal{C}$ can be attained by the strategy profile of one-sided effort, followed by the stationary profile once w_* is reached.

[16]Note that (7) will have to be solved with numerical methods for implicit ODEs since δ_w depends on $f'(w^1)$.

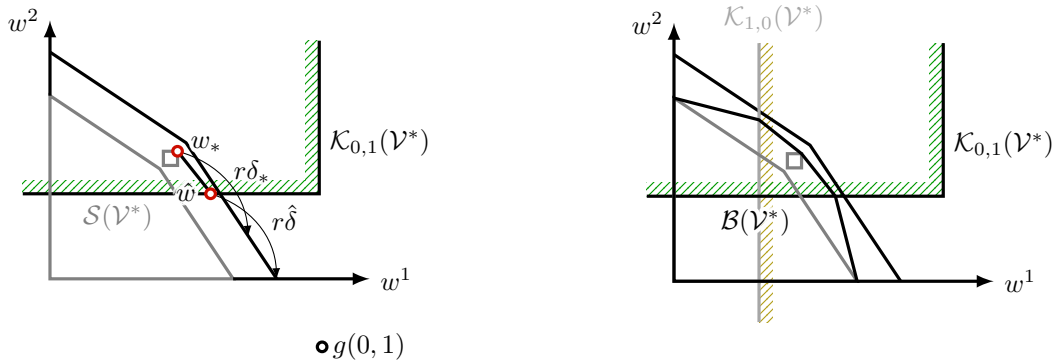

Figure 9: The left panel shows a solution to (6), starting from w_* for $r = 5$ in the partnership game when $\gamma = 1$, i.e., when only the natural disasters are informative. Along the entire segment, (6) is maximized by action profile $(0, 1)$. The right panel shows $\mathcal{B}(\mathcal{V}^*)$ for the same parameters.

The following lemma formalizes this method of constructing $\mathcal{W}$-enforceable strategy profiles.

Lemma 4.3. *Let $\mathcal{C}$ be a differentiable solution to (6) such that $T_w{}^\top(g(a_w) + \delta_w\lambda(a_w) - w)$ is either non-positive or non-negative for every $w \in \mathcal{C}$, where a_w and δ_w denote the maximizers in (6) at w. Then the solution W to (4) for $A = a_{W-}$ and $\delta = \delta_{W-}$ starting at $w_0 \in \mathcal{C}$ remains on $\mathcal{C}$ until either an event occurs or an end point of $\mathcal{C}$ is reached. If the end point is in $\mathcal{B}(\mathcal{W})$, then $\mathcal{C} \subseteq \mathcal{B}(\mathcal{W})$.*

The condition that $T_w{}^\top(g(a_w) + \delta_w\lambda(a_w) - w)$ always has the same sign means that the continuation value of the constructed strategy profile always drifts in the same direction. This is the case for our solution $\mathcal{C}$ in the partnership game and we discuss the general case in Section 5.3.

Because the partnership game is symmetric, there is a solution $\mathcal{C}'$ to (6) symmetric to $\mathcal{C}$, along which country 1 is the sole country to exert effort. One may wonder whether there other solutions to (6) that are relevant in the characterization of $\mathcal{B}(\mathcal{V}^*)$. If natural disasters are the only informative signals, i.e., if $\gamma = 1$, then the answer is negative. Let $\mathcal{C}_\varepsilon$ and $\mathcal{C}'_\varepsilon$ solve (6) below and above the diagonal, respectively, for an initial payoff pair $w_\varepsilon = w_* + \varepsilon(1, 1)^\top$. As the initial value moves closer to the efficient frontier of $\mathcal{V}^*$, the solution becomes more parallel to the frontier because the optimal incentives δ must become more parallel.[17] As a consequence, the angle between solutions $\mathcal{C}_\varepsilon$ and $\mathcal{C}'_\varepsilon$ increases with ε. However, Lemma 4.2 implies that there cannot be any corners outside of $\mathcal{S}(\mathcal{V}^*) \cup \mathcal{K}(\mathcal{V}^*)$, hence no such w_ε is contained in $\partial\mathcal{B}(\mathcal{V}^*)$ when $\gamma = 1$. Figure 9 shows $\mathcal{B}(\mathcal{V}^*)$ when $\gamma = 1$, given by the convex hull of $\mathcal{S}(\mathcal{V}^*)$ and the solutions $\mathcal{C}$ and $\mathcal{C}'$ to (6).

4.3 Incentives through Brownian information

We now add incentives from the performance of the climate indices. Restricting attention to the partnership example does not simplify this step. We will thus motivate the general procedure and then interpret the solutions in the context of the partnership game. At the boundary of $\mathcal{B}(\mathcal{W})$, informational restrictions (I1)–(I3) must hold, giving rise to the following definition.

[17] In the limit as w_ε approaches the frontier, w_ε, $w_\varepsilon + r\delta$, and $g(a) + \delta\lambda(a)$ are collinear and $\mathcal{C}_\varepsilon$ and $\mathcal{C}'_\varepsilon$ are straight line segments.

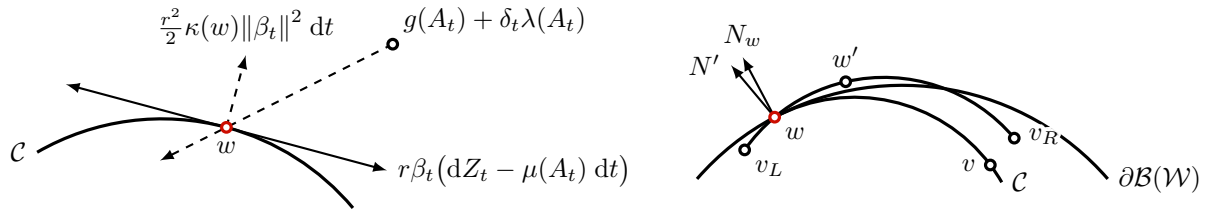

Figure 10: The left panel visualizes that the infinitesimally quick value transfers cause the continuation value to drift away from the curve in orthogonal direction. The right panel shows the perturbation argument that shows that a solution to (9) cannot fall into the interior of $\mathcal{B}(\mathcal{W})$.

Definition 4.4. For a payoff pair $w \in \mathbb{R}^2$, a direction $N \in S^1$, and a payoff set $\mathcal{W}$, we say that a continuation promise (β, δ) from the set

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

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

As a motivation, suppose that a $\mathcal{W}$-enforceable solution W to (2) remains on a continuously differentiable curve $\mathcal{C}$ with curvature κ before the first infrequent event occurs. For W to remain on $\mathcal{C}$, incentives β related to Brownian information must be tangential, i.e., β are tangential transfers between the players. Because Brownian motion has unbounded variation, players transfer value very rapidly. The infinitesimally quick tangential oscillation created by those transfers causes the continuation value to drift away from the curve in orthogonal direction, similar in spirit to the centrifugal force in physics. The magnitude of the outward drift is given by Itō's formula and equals $\frac{r^2}{2}\kappa(w)\|\beta_t\|^2$; see the left panel of Figure 10. For the continuation value to remain on the curve, this outward drift is precisely counteracted by the inward-pointing drift so that at w,

$$\frac{r^2}{2}\kappa(w)\|\beta_t\|^2 = -rN_w{}^\top(w - g(A_t) - \delta_t\lambda(A_t)). \tag{8}$$

Equation (8) shows that for impatient players, the use of noisy information to structure incentives in an equilibrium necessarily causes inefficiency. Self-generation imposes that the inward drift is positive, i.e., the continuation value w lies strictly below the sum of flow payoffs and expected continuation payoffs due to rewards/punishments related to Poisson information.

Observe that the expression in (8) is similar to the abrupt-information optimality equation (6): if incentives through the continuous component of the public signal are not required, i.e., if $\beta = 0$, equation (8) reduces to $N_w{}^\top(g(A_t) + \delta_t\lambda(A_t) - w) = 0$. If incentives through the continuous component are required, we can solve (8) for κ and obtain an expression for the curvature. As we will shortly motivate, the curvature of the boundary of $\mathcal{B}(\mathcal{W})$ is the largest-possible curvature among all restricted-enforceable incentives, that is, it is a solution to the *optimality equation*

$$\kappa(w) = \max_{a\in\mathcal{A}} \max_{(\beta,\delta)\in\Xi_a(w,N_w,\mathcal{W})} \frac{2N_w{}^\top(g(a)+\delta\lambda(a)-w)}{r\|\beta\|^2}, \tag{9}$$

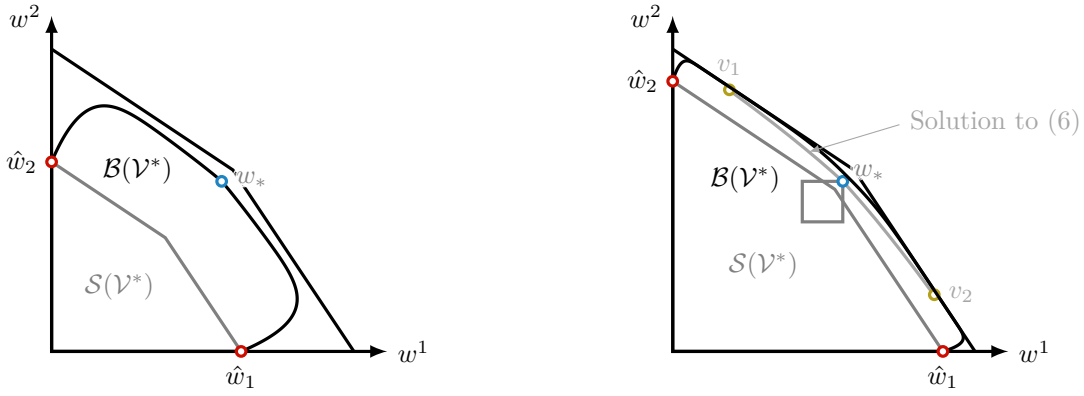

Figure 11: The two panels show $\mathcal{B}(\mathcal{V}^*)$ in the partnership example for $\gamma = 0.4$ and $r = 5$ (left) and $r = 1$ (right). In both cases, the boundary consists of stationary points and solutions to (9) exclusively.

where we set $\kappa(w) = 0$ if the maxima are taken over empty sets.[18] Note that the expected inward movement $-N_w{}^\top \delta\lambda(a)$ of the continuation value due to incentives from Poisson information is the expected value burnt through the use of such incentives. The optimality equation is decreasing both in the expected value burnt and the magnitude $\|\beta\|$ of the tangential transfers provided through Brownian information. Thus, (9) characterizes the optimal trade-off between two sources of inefficiency in the incentive provision—value burning and noisiness of information. To see why the curvature has to be maximal at the boundary, we establish the analogue of Lemma 4.3.

Lemma 4.5. *Let $\mathcal{C}$ be a continuously differentiable solution to* (9) *and let W be a solution to* (2) *starting at $w_0 \in \mathcal{C}$ for $A = a_{W-}$, $\beta = \beta_{W-}$, $\delta = \delta_{W-}$, and $M \equiv 0$, where a_w, β_w, and δ_w are maximizers in* (9) *at w. Then W remains on $\mathcal{C}$ until either an event occurs or an end point of $\mathcal{C}$ is reached. If both end points are in $\mathcal{B}(\mathcal{W})$ or $\mathcal{C}$ is a closed curve, then $\mathcal{C} \subseteq \mathcal{B}(\mathcal{W})$.*

Consider a boundary payoff $w \in \mathcal{B}(\mathcal{W})$ and suppose that the curvature of $\partial\mathcal{B}(\mathcal{W})$ is not maximized among all restricted-enforceable incentives at (w, N_w). Then a solution $\mathcal{C}$ to (9) starting at (w, N_w) falls into the interior of $\mathcal{B}(\mathcal{W})$. Consequently, a solution $\mathcal{C}'$ with initial conditions (w, N') for a slight rotation N' of N_w would leave and re-enter $\mathcal{B}(\mathcal{W})$; see the right panel of Figure 10. That is, it would attain a payoff pair w' strictly outside of $\mathcal{B}(\mathcal{W})$. But Lemma 4.5 implies that $w' \in \mathcal{B}(\mathcal{W})$, which is a contradiction. This argument requires continuity of solutions to (9) in initial conditions, which we address in Section 5.1. On the other hand, a solution $\mathcal{C}$ to (9) cannot escape $\mathcal{B}(\mathcal{W})$ either because $\mathcal{C}$ maximizes the curvature over all restricted-enforceable action profiles and their incentives. Any other enforceable strategy profile that is played has to involve non-tangential value transfers, which would cause the continuation value to escape with positive probability. Details of this argument are contained in Appendix C.1.

Let us now assemble the various pieces in the partnership game for $\gamma = 0.4$, i.e., when both climate indices and natural disasters are informative. By virtue of Lemma 4.2, all corners of $\mathcal{B}(\mathcal{V}^*)$ are either stationary or end points v_i of the solutions to (6), along which only country i exerts effort. The boundary of $\mathcal{B}(\mathcal{V}^*)$ is thus a solution to the optimality equation (9) that connects the

[18] We show in Appendix B that $\kappa(w) = 0$ only if the continuous signal is completely uninformative.

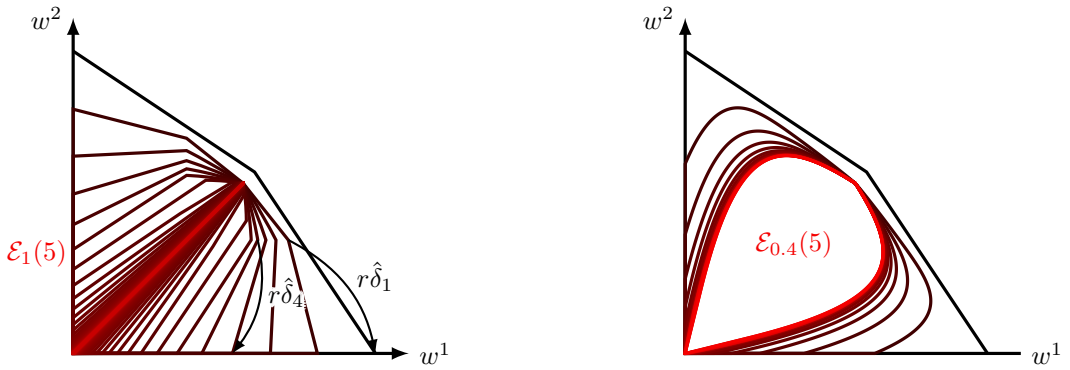

Figure 12: Convergence of the algorithm in Proposition 3.8 for $r = 5$ and $\gamma = 1$ (left) and $\gamma = 0.4$ (right).

extremal stationary payoffs $\hat{w}_i$ on coordinate axis i either to the highest symmetric stationary payoff pair w_*, to the payoff pairs v_{-i}, or directly to each other; see Figure 11. Since $\gamma = 0.4$, natural disasters occur 40% less frequently than in the examples of the previous subsections. If a country is incentivized to exert effort through the threat of punishments, the magnitude of the punishments have to be scaled up accordingly so that the expected punishments remain the same. For discount rate $r = 5$ (left panel), such punishments cannot be delivered in an individually rational way, hence there are no solutions to the abrupt-information optimality equation (6). Natural disasters are still used to provide incentives on the boundary, but they are supplemented by incentives via the climate indices. For discount rate $r = 1$ (right panel), it is possible to provide incentives for one-sided effort by relying on disasters alone, but doing so is not optimal.

4.4 Convergence to the equilibrium payoff set

Any payoff pair $w \in \mathcal{B}(\mathcal{V}^*)$ can be attained with a continuation promise W_σ from $\mathcal{V}^*$ when a natural disaster occurs. However, W_σ may not be credible. The algorithm from Proposition 3.8 iteratively eliminates these non-credible continuations; see Figure 12.

If the continuous information is completely uninformative, i.e., if $\gamma = 1$, then the equilibrium payoff set is entirely contained on the positive diagonal between the static Nash payoff and the highest symmetric stationary payoff. Even though $\mathcal{B}(\mathcal{V}^*)$ contains asymmetric payoff pairs, at which one-sided effort is $\mathcal{V}^*$-restricted enforceable, those disappear as $\mathcal{B}$ is applied repeatedly. At the boundary of $\mathcal{B}(\mathcal{W}_n)$ for small n, it is optimal to punish country 2 by the minimal punishment necessary to incentivize effort. For small n, it is possible to reward country 1 at the same time, but for $n \geq 3$, any continuation value in $\mathcal{W}_n$ consistent with country 2's punishment must also punish country 1; see the left panel of Figure 12. As n increases, the required punishments to country 1 become so severe that country 1 will find it beneficial to deviate in order to reduce the frequency of disasters. Crucial for the collapse to the diagonal is the fact that from action profile $(0, 1)$, a deviation by country 1 has a much stronger impact on the frequency of natural disasters than a deviation by country 2. Thus, country 1 is not willing to tolerate the punishment needed to incentivize country 2.[19]

[19] More formally, any boundary payoff of $\mathcal{B}(\mathcal{W}_n)$, at which $(0, 1)$ is played, has to lie in $\mathcal{K}_{0,1}(\mathcal{W}_n)$. As illustrated in Figure 7, the south-easternmost point of $\mathcal{K}_{0,1}(\mathcal{W}_n)$ is a translate of the south-easternmost point of $\mathcal{W}_n$ by $-r\delta_*$,

The right panel of Figure 12 shows the convergence to the equilibrium payoff set when $r = 5$ and $\gamma = 0.4$. At all curved parts of the boundary, the optimality equation is maximized at the action profiles of one-sided effort. Any PPE that attains a non-stationary boundary payoff thus starts with one-sided effort until the continuation value either reaches w_*, in which case countries adopt a grim-trigger profile, or it reaches the static Nash payoff and countries stop cooperating forever. When a natural disaster occurs, the continuation value jumps along the same segment of the boundary to a point closer to the static Nash payoff; see the left panel of Figure 15. The closer the current continuation value is to the static Nash payoff, the smaller are the credible incentives that can be provided in this way. The remaining incentives for the effort-inducing country i are provided tangentially to the set, depending on the performance of the climate index X^i. Because all jumps point away from w_*, it is possible to reach the grim-trigger phase of the strategy profile only after consistently high values of the climate index X^i. As previewed in Section 2.1, country i is rewarded for effort by shortening the time frame, during which it is the sole actor to exert effort. Conversely, country i is punished for low values of X^i by entering the static Nash phase. However, the equilibrium does not admit thresholds $\underline{x}(J)$ and $\bar{x}(J)$ such that full cooperation starts if $X^i \geq \bar{x}(J)$ and cooperation stops if $X^i \leq \underline{x}(J)$ because there is non-zero drift along the boundary; see Figure 4. To implement the PPE $A(w_0)$ that attains w_0, players have to calculate W_t by solving the SDE (2) based on the realizations of the public signals, then select their action depending on whether or not w_* or $(0, 0)$ have been reached before. An initial value closer to w_* will impose lower expectations on the effort-inducing country, whereas an initial value closer to the static Nash payoff requires a much better performance of X^i for the grim trigger profile to be reached. While such a strategy profile is more complicated than those that attain the abrupt-information optimality equation, it is still simple within the class of all PPE because it is Markovian in the two-dimensional state W.

Payoff pairs in the interior of $\mathcal{E}_\gamma(r)$ can be attained by many different equilibria. In fact, since any action profile is enforceable with $\delta = 0$ in the partnership game, any payoff pair in the interior can be attained by an enforceable solution (W, A, β, δ) to (2) such that W is continuous until the boundary is reached for arbitrary A. In particular, any PPE payoff can be attained by a strategy profile of one-sided effort, followed by a grim trigger profile once continuation value reaches the diagonal (starting with the punishment phase if a payoff in the relative interior of the diagonal is reached).

4.5 Alternative information structures

The optimality equation (9) trades off incentives provided via the two channels of information. Because the boundary is strictly curved where it satisfies (9), the continuation value after a discrete event has to lie strictly below the tangent. Incentives provided via abrupt information thus necessarily burn value, but their use may be efficient if it sufficiently reduces the volatility of the tangential transfers. In this section, we analyze the trade-off between the two types of information for two more monitoring structures in the partnership example.

where $r\delta_* = -(r/8, r/4)^\top$ are the extremal incentives that enforce $(0, 1)$. As a consequence, $\mathcal{K}_{0,1}(\mathcal{W}_n)$ moves $r/\sqrt{8}$ payoff units north west in every iteration of the algorithm until, eventually, no payoff pairs below the diagonal are left.

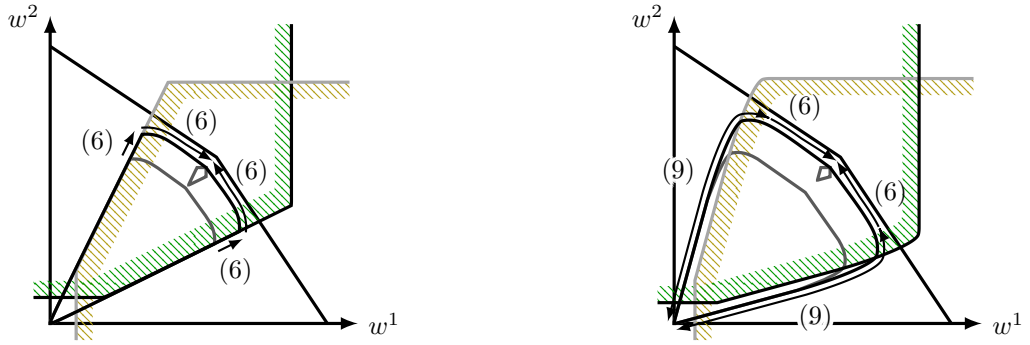

Figure 13: The equilibrium payoff set $\mathcal{E}_\gamma(r)$, the set of stationary equilibrium payoffs $\mathcal{S}(\mathcal{E}_\gamma(r))$, as well as $\mathcal{K}_{0,1}(\mathcal{E}_\gamma(r))$ and $\mathcal{K}_{1,0}(\mathcal{E}_\gamma(r))$ in monitoring structure 2 are shown for $r = 5$ and $\gamma = 1$ in the left and $\gamma = 0.8$ in the right panel. On any segment of the boundary that satisfies (9), the continuation value has infinite variation and it may reach either end point of that segment. On segments that satisfy (6), the continuation value of any equilibrium moves deterministically in the direction of the arrows until a natural disaster occurs.

Let us consider a second monitoring structure, in which the natural disasters arrive according to $\lambda_\gamma(a) = \gamma(33 - 16(a^1 + a^2))$ and μ is defined as in Section 2.1. Natural disasters arrive with the same frequency in both monitoring structures when at least one country exerts effort, but natural disasters are more frequent in monitoring structure 2 when neither country exerts effort. Since this makes it easier to detect deviations by country 2 from $(0, 1)$, smaller punishments are sufficient to deter deviations, which do not violate country 1's incentive constraint.[20] The left panel of Figure 13 shows that $\mathcal{E}_1(5)$ indeed contains asymmetric payoffs under this monitoring structure. The right panel of Figure 13 shows that for $\gamma = 0.8$, the boundary of the equilibrium payoff set changes in a differentiable way from a solution to (9) to a solution to (6) in the interior of $\mathcal{K}_{0,1}(\mathcal{E}_{0.8}(5))$ and $\mathcal{K}_{1,0}(\mathcal{E}_{0.8}(5))$.

Both in this monitoring structure and the monitoring structure of Section 2.1, the matrices $M^i(a)$ for $i = 1, 2$ have rank $|\mathcal{A}^i| - 1$ for any action profile a, that is, any action profile has *individual full rank*. Thus, for any incentives δ provided via Poisson information, there exists β such that the enforceability constraints (3) for both players can be solved with equality. Additionally, the actions by the two countries affect separate climate indices, hence deviations by the two countries, are so-called *pairwise identifiable*, i.e., they can be statistically distinguished. Together with the individual full rank condition, pairwise identifiability guarantees that for any δ, there exists β such that (δ, β) enforces a with $N^\top \beta = 0$ for any non-coordinate direction N. One may wonder whether an individual full rank condition or a pairwise identifiability condition is equally powerful if it is imposed for matrices $\Lambda^i(a)$ consisting of column vectors $\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)$ for $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$. For the theoretical derivation, the answer is negative because informational restriction (I3) does not require the provision of tangential rewards/punishments. Nevertheless, from a quantitative perspective, the ability for player-specific punishments upon the arrival of a discrete event can increase the equilibrium payoff set drastically as we illustrate with the following example.

Consider a third monitoring structure, in which there are two types of natural disasters y_1, y_2 with instantaneous intensities $\lambda_\gamma(y_i, a) = \gamma(16.5 - 16a^i)$ for $i = 1, 2$. The frequency that any event

[20] Arguing as in Footnote 19, the extremal punishments $r\delta_*$ that enforce $(0, 1)$ are now $-(r/8, r/16)^\top$ so that $\mathcal{K}_{0,1}(\mathcal{W})$ need not lie further northwest than $\mathcal{W}$.

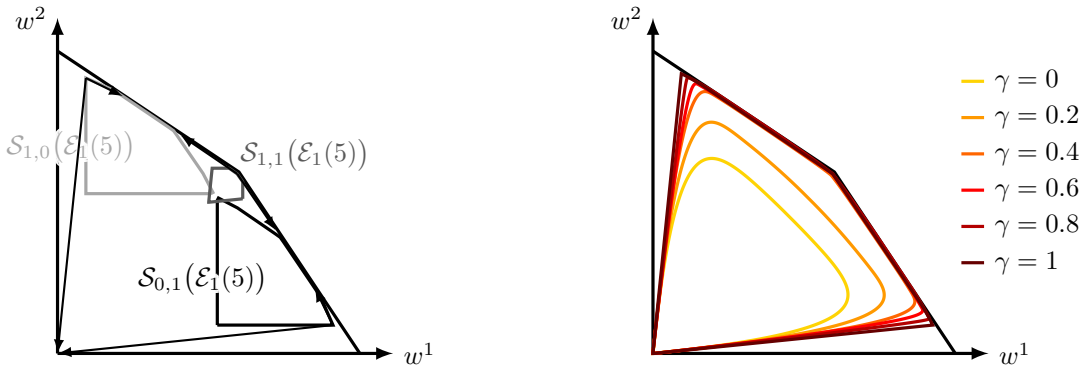

Figure 14: The left panel shows the stationary equilibrium payoffs that can be attained by local play of the three non-static Nash profiles in the partnership game with monitoring structure 3 and $\gamma = 1$. The arrows indicate the rewards/punishments after the two different types of natural disasters for three extremal equilibrium payoff pairs. The right panel shows the equilibrium payoff sets $\mathcal{E}_\gamma(5)$ for γ between 0 and 1.

occurs is the same as in the second monitoring structure, but deviations by the two players can be statistically distinguished. Mutual effort is still enforced by punishing the likely deviator i with $5r/(8\gamma)$ payoff units after any natural disaster y_i occurs. However, since rewards/punishments after event y_i do not affect $-i$'s incentives, it is possible to reward player $-i$ after event y_i so that less total value is destroyed upon the arrival of a natural disaster. As a consequence, the highest symmetric equilibrium payoff is much closer to the efficient frontier than when there is only one type of disaster; see Figure 14. Moreover, there are stationary equilibrium payoffs that support one-sided effort due to the possibility of attaching different continuation values after the two types of natural disasters. In the case $\gamma = 1$, when only the natural disasters are informative, the equilibrium payoff set is the convex hull of stationary payoffs: outside of $\mathcal{S}(\mathcal{E}_1(r))$, the boundary is a trivial solution to (6), i.e., those are straight lines through stationary payoffs.

Figure 15 illustrates the optimal incentives provided via natural disasters in the three different monitoring structures. We see that rewards/punishments $\delta(y)$ after event y are provided close to tangentially in order to minimize the expected value that is burnt through these incentives, unless strong punishments are needed to enforce mutual effort. Consider a boundary payoff w, where an asymmetric action profile a is played in equilibrium. Let $\beta(w)$ and $\delta(w)$ denote the maximizers in (9) at w and define by $\rho(w) := \delta^i(w)\Lambda^i(a)/(\delta^i(w)\Lambda^i(a) + \beta^i(w)M^i(a))$ the ratio of incentives that are provided to the effort-exerting player i through natural disasters. Figure 15 shows the ratio ρ for the three monitoring structures. We observe that in all monitoring structures, $\rho(w)$ decreases as w approaches the static Nash payoff because the maximal credible punishment decreases. In monitoring structures 1 and 2, the continuation value after an event moves almost tangentially to the set. Thus, ρ is largest where the curvature is small: the magnitude of the credible punishments is high relative to the value that has to be burnt. Incentives provided in this way are stronger in monitoring structure 2 because mutual shirking is easier to distinguish from one-sided effort than it is in monitoring structure 1. In monitoring structure 3, mutual effort can be enforced more efficiently because player-specific punishments are available. At boundary payoffs w, where one-sided effort by player i is optimal, incentives $\delta(y_i)$ are provided close to tangentially except where

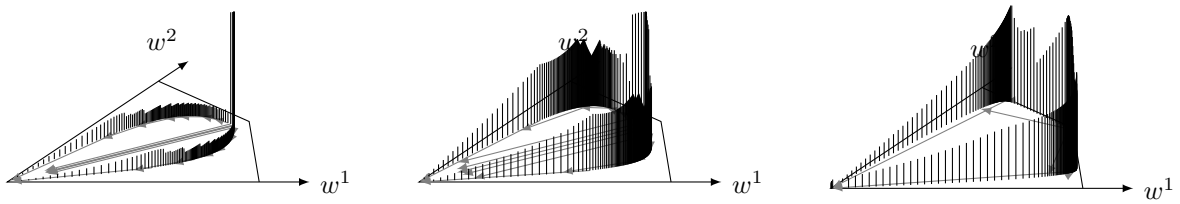

Figure 15: The grey arrows in the (w^1, w^2)-plane indicate the equilibrium rewards/punishments provided upon the realization of a natural disaster in monitoring structure 1 (left), 2 (middle), and 3 (right) for $r = 5$ and $\gamma = 0.4$. The vertical axis shows the fraction ρ of total incentives that are provided via Poisson information. For most payoffs outside of $\Gamma(\mathcal{W})$, it is in the range of 8%–15% in monitoring structure 1, 31%–61% in monitoring structure 2, and 45%–86% in monitoring structure 3.

the curvature is very large. Where the curvature is large, only small tangential transfers can be provided due to (9), hence incentives via Poisson information are simply necessary even if they burn a large amount of value. Moreover, since changes in continuation values after an event y_{-i} have no impact on i's incentives, at those points we have $\delta(w; y_{-i}) = 0$ to minimize the value burnt.

5 Characterization of $\mathcal{B}(\mathcal{W})$ in the general case

Let us turn to the general case, providing additional details to the arguments in the previous section where needed. We begin by discussing existence and uniqueness of solutions to the optimality equations (6) and (9) in Section 5.1. Additionally, we state sufficient conditions on the game primitives to ensure that solutions to (9) are continuous in initial conditions, which is required for the argument highlighted in Figure 10. Section 5.2 expands on boundary points in $\mathcal{K}(\mathcal{W})$, at which extremal incentives in $\mathcal{W}$ bind for at least one player. In Section 5.3 we discuss the extent, to which limit strategies are needed in the derivation, and we present the main result in Section 5.4.

5.1 Regularity of the optimality equations

We begin with existence and uniqueness of solutions to the abrupt-information optimality equation (6). As a motivation, fix an action profile $a \in \mathcal{A}$ and consider a curve $\mathcal{C}$ that satisfies

$$N_w{}^\top w = \max_{\delta \in \Psi_a(w,\mathcal{W})} N_w{}^\top \big(g(a) + \delta\lambda(a)\big), \tag{10}$$

where N_w is the most clockwise normal vector to $\mathcal{C}$ at w. The set $\Psi_a(w, \mathcal{W})$ of all admissible incentives is the intersection of the closed convex set Ψ_a and the $|Y|$-fold product of the compact convex set $(\mathcal{W} - w)/r$. Therefore, both $\Psi_a(w, \mathcal{W})$ and $\mathcal{D}_a(w) := g(a) + \Psi_a(w, \mathcal{W})\lambda(a)$ are compact and convex. At w, the drift rate of an enforceable solution to (4) with $A_0 = a$ thus points in the direction $w - d$ for some $d \in \mathcal{D}_a(w)$. In particular, w is stationary for action profile a if and only if $w \in \mathcal{D}_a(w)$. This yields a geometric interpretation of (10) for non-stationary payoffs $w \notin \mathcal{S}_a(\mathcal{W})$: a solution to ODE (10) evolves in the direction of a tangent to $\mathcal{D}_a(w)$ through w; see Figure 16. It is immediate that there are two possible solutions to (6) starting at any $w \notin \mathcal{S}_a(\mathcal{W})$. Call a solution $\mathcal{C}$ to (10) an *oriented solution* if N_w always points towards the same side of $\mathcal{C}$.

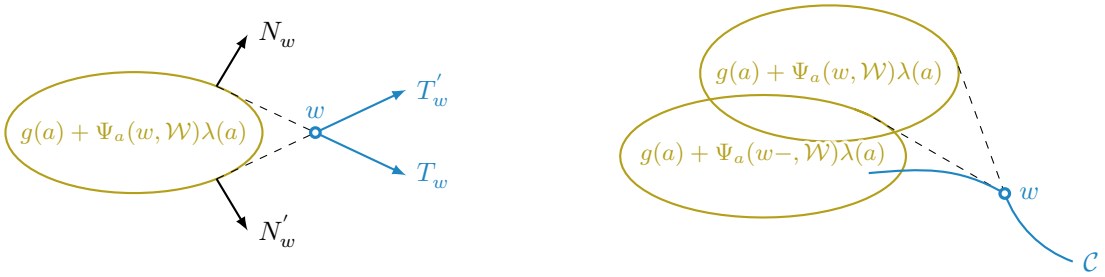

Figure 16: The left panel provides a geometric interpretation of (10). We see that (10) has two solutions for any non-stationary w. The right panel illustrates that an oriented solution is continuously differentiable.

Lemma 5.1.

(i) Any oriented solution to (10) *outside of* $\mathcal{S}_a(\mathcal{W})$ *is continuously differentiable.*

(ii) Within $\operatorname{int}\mathcal{K}_a(\mathcal{W}) \setminus \mathcal{S}_a(\mathcal{W})$*, oriented solutions to* (10) *are unique and continuous in initial conditions.*

Initial conditions in part (ii) consist of an initial state and a binary variable that selects among the two solutions. Because oriented solutions are the only relevant solutions for characterizing the boundary of a convex set, we will simply call them solutions in what follows.

Proof of Lemma 5.1. We begin by showing that $w \mapsto N_w$ is continuous for an oriented solution to (10). Let $(w_n)_{n\geq 0} \subseteq \mathcal{C}$ be a sequence approximating $w \notin \mathcal{S}_a(\mathcal{W})$ from the counterclockwise direction. Lemma B.5 in the appendix shows that $w \mapsto \Psi_a(w, \mathcal{W})$ is continuous, hence $\mathcal{D}_a(w_n)$ converges to $\mathcal{D}_a(w)$. Because w is not stationary and $\mathcal{D}_a(w)$ is closed, $\mathcal{D}_a(w)$ is bounded away from w. Therefore, $\mathcal{D}_a(w_n)$ is bounded away from w_n for n sufficiently large, showing that $N_{w_n} \mapsto N_w$. For the second statement, Lemma B.5 establishes that $w \mapsto \Psi_a(w, \mathcal{W})$ is locally Lipschitz continuous in the interior of $\mathcal{K}_a(\mathcal{W})$. Together with the fact that $\mathcal{D}_a$ is bounded, this shows that that $w \mapsto N_w$ is locally Lipschitz continuous. Finally, the maximum of a locally Lipschitz continuous function over a locally Lipschitz continuous set of parameters is locally Lipschitz continuous by Lemma B.2, which implies uniqueness and continuity in initial conditions by the Picard-Lindelöf theorem. □

Because the maximum over finitely many locally Lipschitz continuous functions is again locally Lipschitz continuous, uniqueness and continuity in initial conditions of solutions to (6) are an immediate corollary to Lemma 5.1 at payoffs where a solution exists. If players additionally maximize over action profiles, the drift rate at w may go in direction $w - d$ for d from any of the sets $\mathcal{D}_a(w)$. Let $\bar{\mathcal{S}}(\mathcal{W})$ denote the set of all payoff pairs w, for which w lies in $\mathcal{D}(w) := \operatorname{conv}\big(\bigcup_{a\in\mathcal{A}} \mathcal{D}_a(w)\big)$. Payoffs in $\bar{\mathcal{S}}(\mathcal{W})$ can be considered stationary for a suitable limit strategy profile, hence the notation $\bar{\mathcal{S}}(\mathcal{W})$; see Section 5.3 for details. The geometric interpretation of the abrupt-information optimality equation (6) is thus identical to (10) with $\mathcal{D}(w)$ and $\bar{\mathcal{S}}(\mathcal{W})$ taking on the roles of $\mathcal{D}_a(w)$ and $\mathcal{S}_a(\mathcal{W})$. Consequently, for any w outside of $\bar{\mathcal{S}}(\mathcal{W})$, there are two oriented solutions starting at any initial point w_0.

Corollary 5.2. *Outside of* $\mathcal{K}(\mathcal{W}) \cup \bar{\mathcal{S}}(\mathcal{W})$*, oriented solutions to* (6) *are continuously differentiable and they are unique and continuous in initial conditions.*

Remark 5.1. Alternatively, we can parametrize a solution to (6) via $\mathrm{d}w_t = (w_t - g(a_{w_t}) - \delta_{w_t}\lambda(a_{w_t}))\,\mathrm{d}t$ for maximizers a_w and δ_w on the right-hand side of (6). Uniqueness in initial conditions implies that $g(a_{w_t}) + \delta_{w_t}\lambda(a_{w_t})$ is generically unique where the solution is strictly curved.

If the public signal has no continuous component or if the continuous component is completely uninformative, stationary payoffs and the abrupt-information optimality equation characterize the boundary of $\mathcal{B}(\mathcal{W})$ entirely. If the continuous component of the public signal is informative, incentives can be provided through abrupt information exclusively for (w, N_w) in the set

$$\Gamma(\mathcal{W}) := \Big\{(w,N) \;\Big|\; \text{there exist } a \in \mathcal{A} \text{ and } \delta \in \Psi_a(w,\mathcal{W}) \text{ with } N^\top(g(a) + \delta\lambda(a) - w) \geq 0\Big\}.$$

Outside of this set, the boundary of $\mathcal{B}(\mathcal{W})$ is characterized by the optimality equation (9). To state the conditions, under which solutions are unique and continuous in initial conditions, let us write the optimality equation as a function of two variables

$$\kappa(w,N) = \max_{a\in\mathcal{A}} \max_{(\beta,\delta)\in\Xi_a(w,N,\mathcal{W})} \frac{2N^\top(g(a) + \delta\lambda(a) - w)}{r\|\beta\|^2}. \tag{11}$$

We will impose conditions such that κ is locally Lipschitz continuous in both arguments. A sufficient condition for local Lipschitz continuity in non-coordinate directions N is a pairwise identifiability condition, stating that deviations by the two players can be statistically distinguished.

Definition 5.3. For any $a \in \mathcal{A}$, let $M^i(a)$ denote the $d \times (|\mathcal{A}^i| - 1)$-dimensional matrix, whose column vectors are given by $\mu(\tilde{a}^i, a^{-i}) - \mu(a)$, for $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$. An action profile a is *pairwise identifiable* if $\operatorname{span} M^1(a) \cap \operatorname{span} M^2(a) = \{0\}$.

Assumption 1. Every enforceable action profile is pairwise identifiable.

To ensure that solutions to (9) are sufficiently well-behaved in coordinate directions N, we impose the following condition.

Assumption 2. For $i = 1, 2$, the minmax profile $\underline{a}_i$ against player i and the global maximizer $\bar{a}_i$ of g^i are enforceable. Moreover, at least one of the following conditions is satisfied:

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

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

Remark 5.2. Assumption 2 differs slightly from the analogous assumption in Sannikov (2007), which requires that (i) or (ii) holds for all enforceable action profiles without imposing enforceability specifically for $\underline{a}_i$ and $\bar{a}_i$. While such an assumption is sufficient to establish global Lipschitz continuity of (9) in games with Brownian information, it is not in our model.[21] Instead of establishing

[21] Restricted-enforceability of an action profile a in a coordinate direction $N = \pm e_i$ requires $\beta^i = 0$. Thus, in games with Brownian information, such an action profile must involve a best response by player i. In our model, an action profile may be restricted-enforceable in a coordinate direction even if it does not involve a best response, as long as sufficient incentives δ^i are provided through Poisson information. Whether such incentives are feasible depends on the location w, hence $\Xi_a(w, N, \mathcal{W})$ may be empty for some w but not others, leaving no hope that the maximum in (11) is continuous, unless we impose the (strong) condition that each $M^i(a)$ has individual full rank.

local Lipschitz continuity in coordinate directions, Assumption 2 is sufficient to establish that the maximization in (11) is always taken over non-empty sets and that coordinate payoffs in $\partial\mathcal{B}(\mathcal{W})$ are isolated, i.e., there are at most four such boundary payoffs.

The following assumption ensures that (11) is locally Lipschitz continuous in w.

Assumption 3. For any enforceable action profile a, there exists β such that $(0,\beta)$ strictly enforces a.

Because the set of continuation promises that enforce any given action profile are convex, Assumption 3 provides a condition, under which local changes in incentives from Poisson signals can be compensated with incentives from the Brownian signal. In particular, Assumption 3 rules out the possibility that an action profile is restricted-enforceable at a payoff pair w, but not at w' arbitrarily close to w. Together with Assumptions 1 and 2, this guarantees local Lipschitz continuity of (11) at almost every (w,N), in a notion made precise in Lemmas B.9, B.10 and C.4.

5.2 Binding and extremal incentives

If it is possible to provide incentives through Poisson information exclusively at some payoff pair w, then—in analogy to discrete time—we will say that w is decomposable.

Definition 5.4. Consider $w \in \partial\mathcal{B}(\mathcal{W})$ and a set of action profiles $\mathcal{A}_w \subseteq \mathcal{A}$.

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

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

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

Recall that the set $\mathcal{K}(\mathcal{W}) = \bigcup_{a\in\mathcal{A}} \partial\mathcal{K}_a(\mathcal{W})$ is the set of payoff pairs, at which only extremal rewards/punishments from $\mathcal{W}$ after the occurrence of some event can deter deviations by at least one player. The following result characterizes all such payoffs on the boundary of $\mathcal{B}(\mathcal{W})$.

Proposition 5.5. *Let $w \in \mathcal{K}(\mathcal{W})$ be a non-stationary boundary payoff of $\mathcal{B}(\mathcal{W})$ such that $(w,N) \in \Gamma(\mathcal{W})$ for some $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$. Then exactly one of the following conditions holds:*

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

$$\mathcal{N}_w(\mathcal{B}(\mathcal{W})) \subseteq \mathcal{N}_w(\mathcal{K}_{\mathcal{A}_w}(\mathcal{W})), \tag{12}$$

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

(ii) There exists $a_* \in \mathcal{A}$ such that $w \in \partial\mathcal{K}_{a_*}(\mathcal{W})$, w is decomposed by $\{a_*\}$, and a_* attains (6) for an extremal normal vector $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$.

If a boundary payoff w is strictly decomposable by some set of action profiles $\mathcal{A}_w$, case (i) establishes that $\mathcal{K}_{\mathcal{A}_w}(\mathcal{W})$ locally is a subset of $\mathcal{B}(\mathcal{W})$; see the left panel of Figure 17. If w is minimally

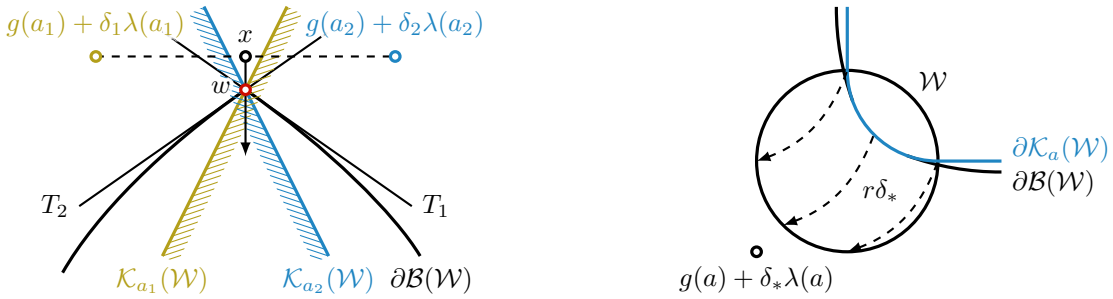

Figure 17: Both panels illustrates case (i) of Proposition 5.5. In the left panel, w is minimally decomposed by $\mathcal{A}_w = \{a_1, a_2\}$ with $\Psi_{a_k}(w) = \{\delta_k\}$ for $k = 1, 2$. Since $g(a_k) + \delta_k\lambda(a_k)$ lies below tangent T_k for $k = 1, 2$, neither $\{a_1\}$ nor $\{a_2\}$ decomposes w. While w may not be a corner of either $\mathcal{K}_{a_1}(\mathcal{W})$ or $\mathcal{K}_{a_2}(\mathcal{W})$, it is a corner of $\mathcal{K}_{\mathcal{A}_w}(\mathcal{W})$ satisfying (12). The right panel shows an example of smooth segment of $\partial\mathcal{B}(\mathcal{W})$ that is strictly decomposed by $\{a\}$, where a is enforceable by $\delta^i \leq \delta^i_*$. Such a segment must lie in $\partial\mathcal{K}_a(\mathcal{W})$ due to (12).

decomposed by a singleton $\mathcal{A}_w = \{a_*\}$, then strict decomposition implies that an enforceable solution to (2) with $A \equiv a_*$ enters the interior of $\mathcal{B}(\mathcal{W})$, where it can be concatenated with a solution that remains in $\mathcal{B}(\mathcal{W})$ until the arrival of the first event. If the boundary is smooth at w and, hence, $\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ is a singleton at w, then minimal decomposition implies that w is necessarily decomposed by a single action profile. Thus, any smooth segments in $\mathcal{K}(\mathcal{W})$ of positive length must be contained in $\partial\mathcal{K}_a(\mathcal{W})$ for a single action profile a; see the right panel of Figure 17. If $\mathcal{A}_w$ is not a singleton, we show in the next subsection how w can be attained by a limit strategy profile.

If a boundary payoff $w \in \mathcal{K}(\mathcal{W})$ is not strictly decomposable, case (ii) establishes that w is the starting point or end point of a solution to the abrupt-information optimality equation (6). This is the case we encountered in the partnership example; see Figure 9. If such w is a corner of $\mathcal{B}(\mathcal{W})$, the fact that w is decomposed by $\{a_*\}$ implies that the drift rate points towards the interior of $\mathcal{K}_{a_*}(\mathcal{W})$, i.e., the continuation value of the strategy profile constructed in Lemma 4.3 moves along $\partial\mathcal{B}(\mathcal{W})$.

While Proposition 5.5 is all we can say for general $\mathcal{W}$ and general signal structures, there are quite a few special signal structures, in which we can rule out some or all of these cases.

Lemma 5.6. *If $|Y| = 1$, there are no curved segments of positive length in $\partial\mathcal{E}(r) \cap \mathcal{K}(\mathcal{E}(r))$.*

Proof. Suppose towards a contradiction that there is a curved segment $\mathcal{C} \subseteq \partial\mathcal{E}(r) \cap \mathcal{K}(\mathcal{E}(r))$. By shortening the segment, we may assume that $\mathcal{C} \subseteq \partial\mathcal{K}_a(\mathcal{E}(r))$ for some action profile a that minimally decomposes w and that each $w \in \mathcal{C}$ is non-coordinate. Because $|Y| = 1$, any such segment in $\partial\mathcal{K}_a(\mathcal{W})$ must be a translate of $\partial\mathcal{W}$ by a constant amount $-r\delta_*$, where δ_* enforces a. Since $\mathcal{B}(\mathcal{W}) = \mathcal{W}$ for $\mathcal{W} = \mathcal{E}(r)$, it follows that $\delta_* = 0$, hence a is a static Nash profile. In particular, $g(a) \in \mathcal{E}(r)$, which contradicts the fact that $\{a\}$ strictly decomposes w. □

Lemma 5.7. *Suppose that the Poisson signals have a* product structure*, i.e., there exists a decomposition of Y into disjoint sets Y_1 and Y_2 with $Y_1 \cup Y_2 = Y$ such that actions by player i affect only the intensity of events in Y_i. Then each $\mathcal{K}_a(\mathcal{W})$ is a (possibly unbounded) rectangle. In particular, there are no curved segments of positive length in $\partial\mathcal{B}(\mathcal{W}) \cap \mathcal{K}(\mathcal{W})$.*

The proof relies on the following bound for $\mathcal{K}_a(\mathcal{W})$, which is tight for games with a product structure. Any $\delta \in \Psi_a(w, \mathcal{W})$ must satisfy $w^i + r \max_{y\in\mathcal{Y}} \delta^i(y) \leq \max_{v\in\mathcal{W}} v^i$ for $i = 1, 2$. A necessary

condition for $w \in \mathcal{K}_a(\mathcal{W})$ is thus that $w^i \leq \max_{v\in\mathcal{W}} v^i - r \min_{\delta^i\in\Psi_a^i} \max_{y\in\mathcal{Y}} \delta^i(y)$. Otherwise, there is no way to provide sufficient incentives for player i with continuations in $\mathcal{W}$. Inspired by this observation, let us define the minmax and maxmin incentives required to enforce a by

$$\underline{\delta}_a^i := \min_{\delta\in\Psi_a^i} \max_{y\in Y} \delta^i(y), \qquad\qquad \bar{\delta}_a^i := \max_{\delta\in\Psi_a^i} \min_{y\in Y} \delta^i(y). \tag{13}$$

An upper bound for $\mathcal{K}_a(\mathcal{W})$ is thus the (possibly unbounded) rectangle $\bar{\mathcal{K}}_a(\mathcal{W})$ of all payoff pairs w such that for both players $i = 1, 2$, the following inequalities are satisfied:

$$\underline{w}_a^i(\mathcal{W}) := \min_{v\in\mathcal{W}} v^i - r\bar{\delta}_a^i \leq w^i \leq \max_{v\in\mathcal{W}} v^i - r\underline{\delta}_a^i =: \bar{w}_a^i(\mathcal{W}).$$

Proof of Lemma 5.7. In games with a product structure, incentives are minimized/maximized over disjoint events Y_1 and Y_2 in (13). Thus, $\underline{w}_a^i(\mathcal{W})$ and $\bar{w}_a^i(\mathcal{W})$ can be attained without affecting incentives for player $-i$. In particular, $\mathcal{K}_a(\mathcal{W}) = \bar{\mathcal{K}}_a(\mathcal{W})$. □

Lemma 5.8. *Suppose that the total intensity of events is constant, i.e., there exists $c > 0$ such that $\sum_{y\in Y} \lambda(y \,|\, a) = c$ for any action profile a. Then $\mathcal{K}(\mathcal{W}) = \emptyset$ for any $\mathcal{W}$.*

Proof. Fix an action profile a with non-empty $\mathcal{K}_a(\mathcal{W})$. Select $w \in \mathcal{K}_a(\mathcal{W})$ and $\delta_w \in \Psi_a(w)$ and define $\delta_v := \delta_w + \frac{w-v}{r}\mathbf{1}^\top$ for any $v \in \mathbb{R}^2$, where $\mathbf{1}$ is the column vector in $\mathbb{R}^{|Y|}$ that contains all ones. By assumption, $\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)$ sums to 0 for any deviation $\tilde{a}^i$, hence $\delta_v^i(\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)) = \delta_w^i(\lambda(\tilde{a}^i, a^{-i}) - \lambda(a))$. In particular, δ_v enforces a. Since $v + r\delta_v(y) = w + r\delta(y)$ for each $y \in Y$ by construction, we have shown $\delta_v \in \Psi_a(v)$. Since v was arbitrary, it follows that $\mathcal{K}_a(\mathcal{W}) = \mathbb{R}^2$. □

5.3 Limit strategies

Let us start by giving two examples of limit strategies and how they may arise in the characterization of $\mathcal{B}(\mathcal{W})$. We then proceed to show that these two examples are the only cases in which limit strategies are necessary to attain payoff pairs in $\mathcal{B}(\mathcal{W})$.

Consider a non-stationary corner w of $\mathcal{B}(\mathcal{W})$ that is strictly and minimally decomposed by a set $\mathcal{A}_w = \{a_1, a_2\}$ as in the left panel of Figure 17. Suppose that $\Psi_{a_k}(w) = \{\delta_k\}$ is a singleton for $k = 1, 2$. Regardless of whether a_1 or a_2 is played initially, an enforceable solution to (2) must immediately escape $\mathcal{B}(\mathcal{W})$ because $g(a_k) + \delta_k\lambda(a_k)$ lies below tangent T_k. If the players, however, publicly randomize infinitesimally quickly between a_1 and a_2 with equal probability, then an exact law of large numbers implies that the drift points in direction $w - x$, where x is the center between $g(a_1) + \delta_1\lambda(a_1)$ and $g(a_2) + \delta_2\lambda(a_2)$. In particular, the continuation value of such a publicly randomized strategy profile enters the interior of $\mathcal{B}(\mathcal{W})$, where it can be concatenated with an enforceable solution to (2). The left panel of Figure 18 illustrates that such infinitesimally quick randomization can indeed arise as a limit strategy profile. Let N be the normal vector to $\mathcal{B}(\mathcal{W})$ in direction $x - w$ so that $w_\varepsilon = w - \varepsilon N$ lies in the interior of $\mathcal{K}_{\mathcal{A}_w}(\mathcal{W})$ for any $\varepsilon > 0$. Any payoff pair w_ε can be attained by alternating between a_1 and a_2 on intervals $[0, \Delta t], [\Delta t, 3\Delta t], [3\Delta t, 5\Delta t] \ldots$ In the limit as $\varepsilon \to 0$, $\Delta t \to 0$ and the continuation value of the limit profile evolves along the straight line in direction $-N$.

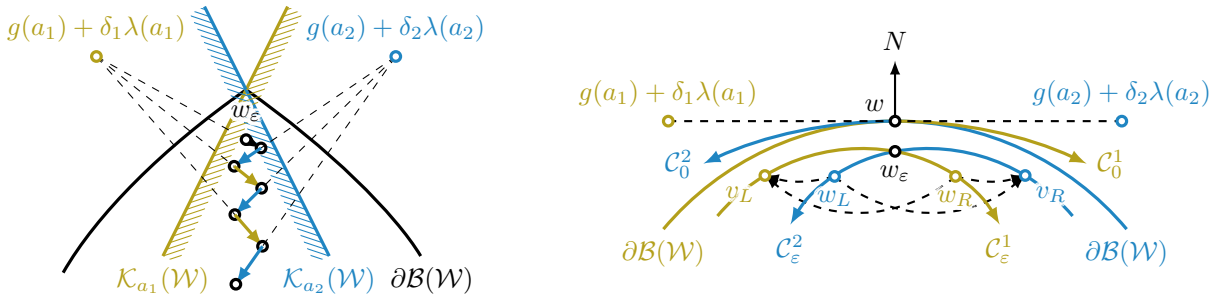

Figure 18: Construction of limit strategies that attain boundary payoffs. In the right panel, the arrows point in the direction, in which a $\mathcal{W}$-enforceable solution to (2) evolves.

As a second example, consider a non-stationary payoff pair $w \notin \mathcal{K}(\mathcal{W})$, at which the maximizing action profile in the abrupt-information optimality equation (6) changes from a_1 to a_2 as in the right panel of Figure 18. If the drift rate along $\partial\mathcal{B}(\mathcal{W})$ points away from w on at least one side of w, then w can be attained by an enforceable strategy profile that remains on $\partial\mathcal{B}(\mathcal{W})$ via the construction in Lemma 4.3. If the drift rate on both sides of w point towards w as in Figure 18, attaining w with an enforceable strategy profile may be a bit trickier. Because w does not lie in $\mathcal{K}(\mathcal{W})$, it lies in the interior of $\mathcal{K}_{a_1}(\mathcal{W})$ and $\mathcal{K}_{a_2}(\mathcal{W})$, hence (10) has a solution for either action profile in a neighborhood of w. Let $\mathcal{C}_\varepsilon^k$ denote the solution for a_k starting at $w-\varepsilon N_w$ for $\varepsilon \geq 0$ and $k = 1, 2$; see the right panel of Figure 18. If either $\mathcal{C}_0^1$ or $\mathcal{C}_0^2$ falls into the interior of $\mathcal{B}(\mathcal{W})$, then w can be attained without the need for a limit strategy. If both $\mathcal{C}_0^1$ and $\mathcal{C}_0^2$ escape $\mathcal{B}(\mathcal{W})$ at w, then w can be attained by an infinitesimally quick public randomization between (a_1, δ_1) and (a_2, δ_2), playing (a_k, δ_k) with probability p_k such that $w = p_1\big(g(a_1)+\delta_1\lambda(a_1)\big)+p_2\big(g(a_2)+\delta_2\lambda(a_2)\big)$. By an exact law of large numbers, the drift rate of the resulting continuation value is 0, suggesting that $w \in \mathcal{B}(\mathcal{W})$. Indeed, such infinitesimally quick public randomization can be attained by the following limit strategy profile. We show in Lemma C.10 that for any $\varepsilon > 0$, the solution $\mathcal{C}_\varepsilon^1$ enters the convex set circumscribed by $\mathcal{C}_\varepsilon^2$ and vice versa; as illustrated in Figure 18. Thus, w_ε can be attained by the strategy profile that alternates between a_1 and a_2 when the continuation value reaches payoff pairs w_L and w_R, and then publicly randomizes between v_L and v_R. In the limit as $\varepsilon \to 0$, we obtain the previously described strategy profile with infinitesimally quick public randomization. In this second example, the continuation value of the limit strategy profile is locally constant. We may think of such payoff pairs as stationary in limit strategies and, indeed, the payoff pair w in the second example lies in $\bar{\mathcal{S}}(\mathcal{W})$ defined in the previous subsection.

In the remainder of this section, we argue that these are the only cases, in which limit strategies are relevant for the characterization of $\mathcal{B}(\mathcal{W})$. Infinitesimally quick public randomization does not change the fact that each action profile in the support of the randomization and its associated continuation promise are subject to informational restrictions (I1)–(I3). Thus, public randomization cannot enlarge the set of action profiles and continuation promises over which we maximize in the optimality equations. Lemmas 4.3 and 4.5 show that limit strategies are not needed along solutions to (6) and (9) except for the case discussed in the previous paragraph. Note that such cases must be isolated points: if $\mathcal{C}_0^1$ and $\mathcal{C}_0^2$ coincide for a segment $\mathcal{C}_*$ of positive length, the entire segment can be attained without limit strategies by alternating action profiles when end points of $\mathcal{C}_*$ are reached.

Moreover, $\mathcal{B}(\mathcal{W})$ cannot have any corners at stationary limit payoffs $\bar{\mathcal{S}}(\mathcal{W})$ outside of $\mathcal{K}(\mathcal{W}) \cup \mathcal{S}(\mathcal{W})$ because at least one (a, δ) in the support would violate Lemma 4.2. Therefore, limit strategies only serve to ensure that $\mathcal{B}(\mathcal{W})$ is closed, but they do not allow us to attain payoffs bounded away from those that are possible without infinitesimally quick public randomization (if $\mathcal{B}(\mathcal{W})$ has non-empty interior). This discussion is formalized in the following two results.

Lemma 5.9. *$\mathcal{B}(\mathcal{W})$ is closed.*

Proposition 5.10. *Any payoff pair in the interior of $\mathcal{B}(\mathcal{W})$ can be attained without limit strategies. Moreover, every $w \in \mathcal{B}(\mathcal{W})$ can be attained without limit strategies if and only if:*

(i) For every non-stationary corner of $\mathcal{B}(\mathcal{W})$ that is strictly and minimally decomposed by a non-singleton set $\mathcal{A}_w$, there exists $a \in \mathcal{A}_w$ such that the solution to (10) *enters the interior of $\mathcal{B}(\mathcal{W})$.*

(ii) For every payoff pair w, at which the maximizer in (6) *changes outside of $\mathcal{K}(\mathcal{W})$, the solution to* (10) *for one maximizer a starting at w remains in $\mathcal{B}(\mathcal{W})$ for a segment of positive length.*

Because $\mathcal{B}(\mathcal{E}(r)) = \mathcal{E}(r)$, Proposition 5.10 gives us necessary and sufficient conditions whether limit strategies are needed to attain extremal equilibrium payoffs.

Remark 5.3. If $\mathcal{B}(\mathcal{W})$ has non-empty interior, Proposition 5.10 implies that the characterization of $\partial\mathcal{B}(\mathcal{W})$ is identical regardless of whether or not we allow limit strategies. The reason why we allow them is that closedness is important for the algorithm in Proposition 3.8. Suppose that $\mathcal{B}'(\mathcal{W})$ is defined analogously to $\mathcal{B}(\mathcal{W})$ without limit strategies and we set $\mathcal{W}_0 = \mathcal{V}^*$ and $\mathcal{W}_n = \operatorname{cl} \mathcal{B}'(\mathcal{W}_{n-1})$ for $n \geq 1$. The proof of Proposition 3.8 would then show only that the interior of $\mathcal{W}_\infty$ is $\mathcal{W}_\infty$-relaxed self-generating, which is not sufficient to conclude that $\operatorname{int} \mathcal{W}_\infty$ is, in fact, self-generating.

5.4 Main results

Aggregating all of our results, the following is our main result, providing a complete characterization of $\mathcal{B}(\mathcal{W})$. To state it efficiently, we denote the normal bundle of a convex set $\mathcal{X}$ by

$$\mathcal{N}_\mathcal{X} := \Big\{(w, N) \in \partial\mathcal{X} \times S^1 \;\Big|\; N \in \mathcal{N}_w(\mathcal{X})\Big\}.$$

Theorem 5.11. *Suppose that either the Brownian information is completely uninformative or Assumptions 1–3 hold. Let $\mathcal{W} \subseteq \mathcal{V}^*$ be convex and compact with $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$. Then $\mathcal{B}(\mathcal{W})$ is the largest closed convex subset $\mathcal{X}$ of $\mathcal{V}^*$ that contains $\mathcal{S}(\mathcal{W})$ such that:*

(i) Outside of $\mathcal{S}(\mathcal{W}) \cup \mathcal{K}(\mathcal{W})$, the boundary $\partial\mathcal{X}$ is continuously differentiable and $(w, N_w) \in \mathcal{N}_\mathcal{X}$ solves (6) *inside of $\Gamma(\mathcal{W})$ and* (9) *outside of $\Gamma(\mathcal{W})$,*

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

Moreover, if Assumptions 1–3 hold (which precludes the Brownian signal from being completely uninformative), straight line segments in $\partial\mathcal{B}(\mathcal{W})$ are contained in $\partial\mathcal{V}^* \cap \operatorname{conv} \mathcal{S}(\mathcal{W})$.

In addition to characterizing the largest bounded $\mathcal{W}$-relaxed self-generating set, Theorem 5.11 provides a fixed-point characterization of the PPE payoff set $\mathcal{E}(r)$ for $\mathcal{W} = \mathcal{E}(r)$. The entirety of the discussion in Sections 4 and 5 thus applies to $\mathcal{E}(r)$ as well. We summarize the main points here. If extremal incentives can be provided without the use the Brownian signal, the abrupt-information optimality equation (6) imposes that players play the action profile that is enforceable by the smallest amount of value burning. If Brownian information is used at the boundary, the optimality equation (9) optimizes the trade-off between the expected value burnt through abrupt information and the volatility of tangential transfers related to the Brownian signal. Both of these optimizations give rise to a smooth boundary except at stationary payoffs and where the extremal credible rewards/punishments bind for at least one player.

While the optimality equation characterizes the quantitative trade-off between the two types of information, the following lemma describes a qualitative trade-off.

Lemma 5.12. *For any game parametrized by g, μ, λ and any $\gamma \in [0,1]$, let $\mathcal{E}_\gamma(r)$ denote the equilibrium payoff set of the game parametrized by g, $\gamma\mu$, and λ. If $\mathcal{E}_0(r)$ is a polygon, then there exists $\gamma_0 > 0$ such that $\operatorname{ext}\mathcal{E}_0(r) \subseteq \operatorname{ext}\mathcal{E}_\gamma(r)$ for any $\gamma < \gamma_0$.*

PPE that attain extremal payoffs when the Brownian signal is completely uninformative still attain extremal payoffs if the Brownian signal is somewhat informative. Fundamentally, Poisson information serves to enlarge the set of enforceable action profiles, whereas Brownian information improves efficiency incrementally by optimizing the trade-off in the optimality equation.

Our last result concerns the use of public randomization. In addition to its use in limit strategies, which is discussed in Proposition 5.10, public randomization has simplified the analysis by allowing us to conclude early on that the equilibrium payoff set is convex. Despite these two important uses of public randomization in the derivation of our result, it can easily be verified whether the equilibrium payoff set $\mathcal{E}(r)$ with public randomization coincides with the equilibrium payoff set $\mathcal{E}_*(r)$ without public randomization. Since $\mathcal{E}_*(r) \subseteq \mathcal{E}(r)$, it is sufficient to find conditions such that any payoff pair in $\mathcal{E}(r)$ can be attained without public randomization.

Lemma 5.13. *Suppose that the conditions of Proposition 5.10 are satisfied. If either,*

(i) Assumptions 1–3 hold and $\partial\mathcal{V}^* \cap \mathcal{E}(r) \subseteq \mathcal{S}(\mathcal{E}(r))$, or

(ii) there is a static Nash profile in pure actions, and the infinite extensions of all straight line segments of $\partial\mathcal{E}(r)$ contain a static Nash payoff,

then every payoff pair $w \in \mathcal{E}(r)$ can be attained with a PPE without the use of public randomization.

For all monitoring structures of the partnership game considered in Section 4.5, the above conditions are satisfied and hence $\mathcal{E}(r) = \mathcal{E}_*(r)$.

6 Computation

In this section, we illustrate how to implement Theorem 5.11 numerically. We discuss how to compute $\mathcal{S}(\mathcal{W})$ and $\mathcal{K}(\mathcal{W})$ when there is more than one type of events in Sections 6.1 and 6.2, before

illustrating how to solve (9) and (6) numerically in Section 6.3. Finally, we provide an improvement over the algorithm in Proposition 3.8 for numerical implementation in Section 6.4.

6.1 Computing $\mathcal{S}(\mathcal{W})$ for arbitrary signal structures

If there is more than one type of events, the steps in Section 4.1 have to be carried out in the incentive space rather than the payoff space. Let $e_y \in \mathbb{R}^{|Y|}$ denote the unit vector corresponding to event y and define $f_y(\delta) := \delta(\lambda(a) + re_y)$ for any $a \in \mathcal{A}$. For a stationary payoff $w = g(a) + \delta\lambda(a)$, the continuation value after event y is given by $g(a) + f_y(\delta) = w + r\delta(y)$. Incentives at stationary payoffs in $\mathcal{S}_a(\mathcal{W})$ thus have to come from the set $\mathcal{X}_a(\mathcal{W}) := \Psi_a \cap \bigcap_{y\in Y} f_y^{-1}(\mathcal{W} - g(a))$, 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 $\pm g(a)$. Denote by $\mathcal{X}_a(\mathcal{W})\lambda(a) := \{w \in \mathbb{R}^2 \mid \exists\ \delta \in \mathcal{X}_a(\mathcal{W}) \text{ with } \delta\lambda(a) = w\}$ the projection of $\mathcal{X}_a(\mathcal{W})$ onto $\mathbb{R}^2$ in the direction $\lambda(a)$. It follows that

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

Note that Ψ_a is a convex polyhedron characterized by the affine inequalities in (3). For a discretization $\mathcal{W}'$ of $\mathcal{W}$ with extremal points $w_1, \dots, w_n$ and corresponding normal vectors $N_1, \dots, N_n$, the set $\mathcal{X}_a(\mathcal{W}')$ is a convex polyhedron again, characterized by the affine inequalities in (3) and

$$N_j{}^\top \delta(\lambda(a) + re_y) \le N_j{}^\top (w_j - g(a)), \qquad j = 1, \dots, n,\ y \in Y.$$

If $|Y|$ is sufficiently small to compute the extremal points of $\mathcal{X}_a(\mathcal{W}')$, one can compute $\mathcal{X}_a(\mathcal{W}')\lambda(a)$ by taking the convex hull of the projection $\delta\lambda(a)$ of those extremal points δ. If $|Y|$ is large, one can maximize $N^\top\delta\lambda(a)$ over $\delta \in \mathcal{X}_a(\mathcal{W}')$ for a grid of normal vectors N.

6.2 Computing $\mathcal{K}(\mathcal{W})$ for arbitrary signal structures

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

$$\mathcal{K}_a(\mathcal{W}) = \bigcap_{j=1}^m \mathcal{K}_a^j(\mathcal{W}) \qquad\qquad \mathcal{K}_a^j(\mathcal{W}) := \bigcup_{\delta\in\partial\Psi_a(Y_j)} \bigcap_{y\in Y_j} \big(\mathcal{W} - r\delta(y)\big). \tag{14}$$

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

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

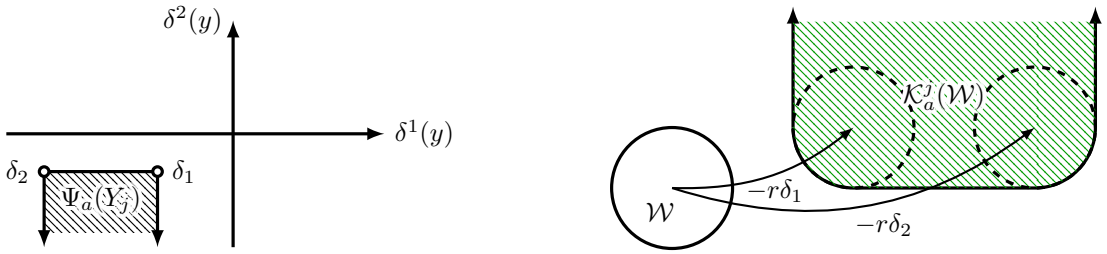

Figure 19: Construction of $\mathcal{K}_a^j(\mathcal{W})$ if $Y_j = \{y\}$ is a singleton. The arrows indicate that $\Psi_a(Y_j)$ and $\mathcal{K}_a^j(\mathcal{W})$ continue indefinitely in these directions, i.e., two "corners" of $\Psi_a(Y_j)$ are degenerate.

of event of type y that enforce a. The set $\mathcal{K}_a^j(\mathcal{W})$ is then simply the convex hull of $\mathcal{W} - r\delta_k$ for $k = 1, \dots, 4$, where δ_k denote the four (possibly degenerate) corners of $\Psi_a(Y_j)$; see Figure 19.

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

$$\min_{v \in \mathcal{W}} v^i - r\bar{\delta}^i_{a,j} \leq w^i \leq \max_{v \in \mathcal{W}} v^i - r\underline{\delta}^i_{a,j}$$

where $\underline{\delta}^i_{a,j} := \min_{\delta \in \Psi_a} \max_{y \in Y_j(a)} \delta^i(y)$ and $\bar{\delta}^i := \max_{\delta \in \Psi_a} \min_{y \in Y_j(a)} \delta^i_{a,j}(y)$ are the extremal punishments/rewards for player i that can be carried out after the arrival of any event in Y_j such that i is incentivized to play a^i against a^{-i}. Note that the decomposition of Y into such sets $\mathcal{Y}(a)$ allows us to compute $\mathcal{K}_a(\mathcal{W})$ with these simple steps even if $|Y| > 1$ and the game is not of a product structure, as long as the restriction to each $Y_j(a)$ exhibits one of these two forms.

The computation of $\mathcal{K}_a(\mathcal{W})$ is cumbersome in the most general case. However, we show in the next subsection that in order to compute $\mathcal{B}(\mathcal{W})$, it is sufficient to verify whether a given payoff pair is in $\mathcal{K}_a(\mathcal{W})$. For a polygon approximation $\mathcal{W}'$ of $\mathcal{W}$ with extremal points $w_1, \dots, w_n$ and corresponding normal vectors $N_1, \dots, N_n$, this corresponds to verifying whether or not the system of affine inequalities given by (3) and $rN_j{}^\top \delta(y) \leq N_j{}^\top (x_j - w)$ for $j = 1, \dots, n$ and $y \in Y$ is feasible.

6.3 Computing $\mathcal{B}(\mathcal{W})$ for arbitrary signal structures

Since $\mathcal{B}(\mathcal{W})$ is convex, we parametrize its boundary via tangential angle θ. Let $w(\theta)$ denote the set of payoff pairs in $\partial\mathcal{B}(\mathcal{W})$ with normal vector $N(\theta) = \big(\cos(\theta), \sin(\theta)\big)^\top$. If Assumptions 1–3 are satisfied, the boundary is strictly curved, hence $w(\theta)$ is a singleton. Outside of $\Gamma(\mathcal{W})$, one can thus solve

$$\frac{\mathrm{d}w(\theta)}{\mathrm{d}\theta} = \frac{T(\theta)}{\kappa(\theta)} \tag{15}$$

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 (9). If the continuous signal is completely uninformative, solutions to (9) are simply straight line segments in the direction of $T(\theta)$. Within $\Gamma(\mathcal{W})$, one can solve (6) numerically by

letting $w(\theta_{k+1})$ be the payoff w furthest from $w(\theta_k)$ in the direction of $T(\theta_k)$, for which

$$\max_{a\in\mathcal{A}} \max_{\delta\in\Psi_a(w,\mathcal{W})} N(\theta_{k+1})^\top (g(a)+\delta\lambda(a)-w) = 0. \tag{16}$$

In order to compute $\mathcal{B}(\mathcal{W})$, we first find the set of stationary payoffs as in Section 4.1 or 6.1. Since $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$, we know that any stationary payoff w_0 in $\partial\mathcal{W}$ must be a contained in $\partial\mathcal{B}(\mathcal{W})$. Therefore, we can search for the maximal set $\mathcal{X}$ that satisfies Theorem 5.11 by searching for the initial angle θ_0 at w_0, with which to solve (6) and (9). If no extremal payoff pair from $\mathcal{B}(\mathcal{W})$ is known, one can use an iterative procedure as described in Section 8 of Sannikov (2007). One can find solutions to (6) and (9) for some initial value (w_0, θ_0) with the following steps:

(i) Compute $q_a(\theta_k) := \max_{\delta\in\Psi_a(w_k,\mathcal{W})} N(\theta_k)^\top (g(a)+\delta\lambda(a)-w_k)$ and $q(\theta_k) = \max_{a\in\mathcal{A}} q_a(\theta_k)$.

(ii) If $q(\theta_k) < 0$, then $(w_k, N(\theta_k)) \notin \Gamma(\mathcal{W})$, hence we obtain w_{k+1} by solving (15).

(iii) If $q(\theta_k) = 0$ (within numerical tolerance), test whether w_k lies on the boundary of $\mathcal{K}_a(\mathcal{W})$ for some action profile a, for which $q_a(\theta_k) = 0$. For example, test whether $\Psi_a(w_k + \varepsilon T(\theta_k), \mathcal{W})$ is empty for $\varepsilon > 0$. If it is not, w_k satisfies case (ii) of Proposition 5.5. Set $w_{k+1} = w_k + \varepsilon T(\theta_k)$ and $\theta_{k+1} = \theta_k$ so that in the next step, the solution outside of $\mathcal{K}_a(\mathcal{W})$ will be found. If w_k lies in the interior of $\mathcal{K}_a(\mathcal{W})$, the boundary must solve (6), hence we obtain w_{k+1} from (16).

(iv) If $q_a(\theta_k) > 0$ for a single action profile a:

  (a) If $w_k \in \operatorname{int}\mathcal{K}_a(\mathcal{W})$, then w_k cannot be on the boundary of $\partial\mathcal{B}(\mathcal{W})$ and we stop the solution.

  (b) If w_k is on the boundary of $\mathcal{K}_a(\mathcal{W})$ and $N(\theta_{k+1})$ is a normal vector to $\mathcal{K}_a(\mathcal{W})$ at w_k, then $\partial\mathcal{B}(\mathcal{W})$ may have a corner at w_k and we set $w_{k+1} = w_k$.

  (c) If w_k is on the boundary of $\mathcal{K}_a(\mathcal{W})$ and $N(\theta_{k+1})$ is not a normal vector to $\mathcal{K}_a(\mathcal{W})$ at w_k, then $\partial\mathcal{B}(\mathcal{W})$ may evolve along $\partial\mathcal{K}_a(\mathcal{W})$ as in the right panel figure of Figure 17. Let w_{k+1} be the payoff pair w closest to w_k that maximizes $N(\theta_{k+1})^\top w$ among $w \in \mathcal{K}_a(\mathcal{W})$.

(v) If $q_a(\theta_k) > 0$ for multiple action profiles a:

  (a) If $w_k \in \operatorname{int}\mathcal{K}_a(\mathcal{W})$ for some a with $q_a(\theta_k) > 0$, then w_k cannot be on the boundary of $\partial\mathcal{B}(\mathcal{W})$ and we stop the solution.

  (b) If $w_k \in \partial\mathcal{K}_a(\mathcal{W})$ for all a with $q_a(\theta_k) > 0$, then $\partial\mathcal{B}(\mathcal{W})$ may have a corner at w_k as in case (i) of Proposition 5.5. We thus search over the largest solutions to (6) and (9) that connect to $\mathcal{S}(\mathcal{W})$ or w_0, starting at $(w_k + \varepsilon T(\theta_j), \theta_j)$ for $j \geq k$.

If the solution falls into the interior of $\operatorname{conv}(\mathcal{S}(\mathcal{W}))$ or we have stopped the solution in cases (iv.a) or (v.a), then the entire solution lies inside $\mathcal{B}(\mathcal{W})$ and we adjust the initial angle to enclose a larger area. If the solution escapes $\mathcal{W}$ at any point, then the entire solution lies outside of $\mathcal{B}(\mathcal{W})$, hence we adjust the initial angle in the opposite direction. We illustrate this procedure by showing how $\mathcal{B}(\mathcal{V}^*)$ can computed in the partnership example for the second monitoring structure with $r = 5$ and $\gamma = 0.8$. We first compute the set of stationary payoffs as described in Section 4.1. Since $\mathcal{S}(\mathcal{V}^*)$ overlaps with $\partial\mathcal{V}^*$, it is necessary that $\mathcal{S}(\mathcal{V}^*) \cap \partial\mathcal{V}^*$ lies on the boundary of $\mathcal{B}(\mathcal{V}^*)$. Payoff pairs

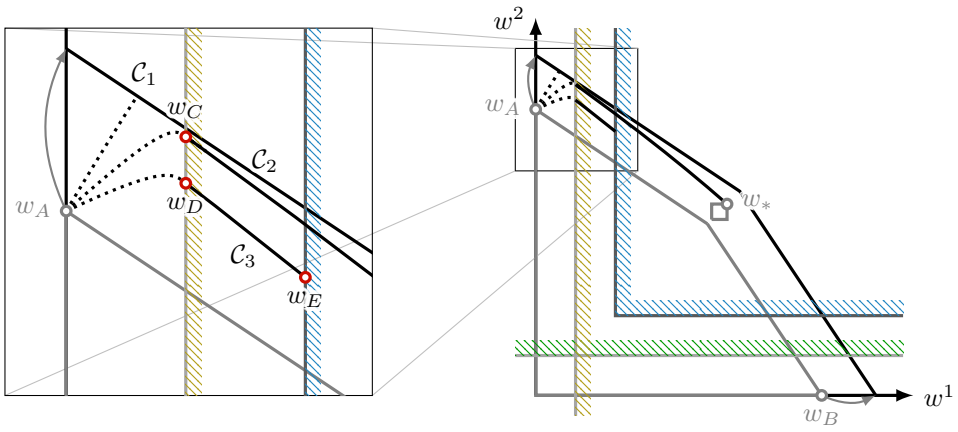

Figure 20: Shown are the set $\mathcal{S}(\mathcal{V}^*)$ of stationary payoffs, the sets $\mathcal{K}_{1,1}(\mathcal{V}^*)$, $\mathcal{K}_{0,1}(\mathcal{V}^*)$, and $\mathcal{K}_{1,0}(\mathcal{V}^*)$, where incentives are binding and extremal, as well as three solutions $\mathcal{C}_i$ to (9) (dotted) and (6) (solid) for $i = 1, 2, 3$ with initial conditions $(w_A, N(\theta_i))$ and $\theta_1 > \theta_2 > \theta_3$. The light gray arrows indicate the rewards/punishments that attain stationary payoff pairs w_A and w_B.

w_A and w_B in Figure 20 thus serve as starting points for solving (6) and (9). We search for the maximal angle θ at w_A, for which a solution connects to either w_B or w_*. Solution $\mathcal{C}_1$ for initial condition (w_A, θ_1) escapes $\mathcal{V}^*$. It follows that θ_1 is too large. Solution $\mathcal{C}_3$ for initial condition (w_A, θ_3) reaches a payoff pair w_D, for which action profile $(0, 1)$ satisfies (6) at (w_D, N_{w_D}). We extend $\mathcal{C}_3$ by solving (6) with initial state w_D. The curve eventually reaches a payoff pair $w_E \in \partial\mathcal{K}_{1,1}(\mathcal{V}^*)$, for which $\max_\delta {N_{w_E}}^\top(g(1,1) + \delta\lambda(1,1) - w_E) > 0$. Since (12) is violated, $\mathcal{C}_3$ cannot be a part of the boundary of $\mathcal{B}(\mathcal{W})$, hence θ_3 is too small. Solution $\mathcal{C}_2$ for initial condition (w_A, θ_2) changes from a solution to (9) to a solution to (6) at w_C and it connects to w_*. After verifying that any solution with initial conditions (w_A, θ) for $\theta > \theta_2$ escapes $\mathcal{V}^*$, we conclude that $\mathcal{C}_2$ coincides with $\partial\mathcal{B}(\mathcal{V}^*)$.[22]

While it is helpful to know $\mathcal{K}_a(\mathcal{W})$ explicitly for steps (iii)–(v), one can also use the verification procedure at the end of Section 6.2. Specifically, $N(\theta)$ is a normal vector to $\mathcal{K}_{\mathcal{A}_\theta}(\mathcal{W})$ at $w(\theta)$ if and only if $w(\theta) \in \mathcal{K}_a(\mathcal{W})$ for every $a \in \mathcal{A}_\theta$ and $w(\theta) + \varepsilon T'$ lies outside of $\mathcal{K}_a(\mathcal{W})$ for any $a \in \mathcal{A}_\theta$, any sufficiently small $\varepsilon > 0$, and sufficiently slight outward rotations T' of $\pm T(\theta)$. To find the payoff pair $w \in \mathcal{K}_a(\mathcal{W})$ that maximizes $N(\theta_{k+1})^\top w$, one can perform a binary search between $w(\theta_k)$ and the intersection point $\hat{w}$ of $\partial\mathcal{K}_a(\mathcal{W})$ and the straight line through $w(\theta_k)$ in direction $T(\theta_{k+1})$ as follows. Let $\tilde{w}$ denote the projection of the mid point between $w(\theta_k)$ and $\hat{w}$ onto $\partial\mathcal{K}_a(\mathcal{W})$ in direction $N(\theta_{k+1})$. If $N(\theta_{k+1})$ is a normal vector of $\mathcal{K}_a(\mathcal{W})$ at $\tilde{w}$, then $w(\theta_{k+1})$ is the payoff pair on the tangent through $\tilde{w}$ closest to $w(\theta_k)$. If $N(\theta_{k+1})$ is not a normal vector to $\mathcal{K}_a(\mathcal{W})$ at $\tilde{w}$, then we continue the binary search between $\tilde{w}$ and $\hat{w}$ if $\tilde{w} + \varepsilon T(\theta_{k+1}) \in \mathcal{K}_a(\mathcal{W})$ and between $\tilde{w}$ and $w(\theta_k)$ if $\tilde{w} - \varepsilon T(\theta_{k+1}) \in \mathcal{K}_a(\mathcal{W})$. The computational procedure can be simplified with Lemmas 5.7 and 5.8 if the game has a special signal structure by terminating non-maximal solutions early. For example, if the conditions of Lemma 5.8 are met, steps (iv) and (v) cannot occur on $\partial\mathcal{B}(\mathcal{W})$, hence we can stop the solution in either case.

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

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

Observe that the set of stationary payoffs shrinks with every iteration: extremal payoff pairs in $\mathcal{S}(\mathcal{W}_n) \cap \partial\mathcal{V}^*$ are not contained in the next step of the iteration if they are decomposed using outward jumps, i.e., jumps $\delta(y)$ with $N^\top \delta(y) > 0$ for some event y. Examples are the payoff pairs w_A and w_B in Figure 20. Those are decomposed by the static Nash profile and incentives that give countries 2 and 1, respectively, their highest feasible and individually rational payoff upon the occurrence of a natural disaster. In the limit as $\mathcal{B}$ is applied infinitely often, incentives using outward jumps cannot support equilibrium behavior since $\mathcal{E}(r)$ is convex. This motivates a refined algorithm that excludes such incentives from the beginning. We begin by defining the set of payoff pairs that can be decomposed using inward jumps. For any action profile $a \in \mathcal{A}$ and any direction $N \in S^1$, let

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

be the half-space of payoff pairs that can be decomposed with inward-jumps with respect to N. Similarly to Fudenberg and Levine (1994) and Sannikov and Skrzypacz (2010), $g(a)$ cannot be in the interior of $H_a(N)$. Let $D_a := \{N \in S^1 \mid H_a(N) \neq \emptyset\}$ denote the set of directions, with respect to which a can be decomposed using inward jumps only. The set $\mathcal{Q}_a := \bigcap_{N \in D_a} H_a(N)$ is a bound for all payoff pairs that are decomposable by a using only inward-pointing jumps. Since a static Nash profile $a \in \mathcal{A}^N$ is decomposable without any jumps at all, $g(a) \in \partial H_a(N)$ for all $N \in S^1$ and hence $\mathcal{Q}_a = \{g(a)\}$. The following lemma follows straight from the definition of $\mathcal{Q}_a$.

Lemma 6.1. *Suppose that $\mathcal{A}_w \subseteq \mathcal{A}$ decomposes $w \in \partial\mathcal{B}(\mathcal{W})$. For any $a \in \mathcal{A}_w$, either $w \in \mathcal{Q}_a$ or any δ with (a, δ) decomposing w satisfies $N^\top \delta(y) > 0$ for some $y \in Y$ and some $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$.*

We obtain the following fixed-point characterization of $\mathcal{E}(r)$ as an immediate consequence to Theorem 5.11 and Lemma 6.1.

Corollary 6.2. *$\mathcal{E}(r)$ is the largest closed convex subset $\mathcal{X}$ of $\mathcal{V}^*$ that satisfies Conditions (i) and (ii) of Theorem 5.11 for $\mathcal{W} = \mathcal{X}$ such that every corner w of $\mathcal{X}$ lies in $\mathcal{S}_a(\mathcal{W}) \cap \mathcal{Q}_a$ or in $\mathcal{K}_a(\mathcal{W}) \cap \mathcal{Q}_a$ for every a in $\mathcal{A}_w$ that minimally decomposes w.*

The refinement in Corollary 6.2 allows us to exclude many points that would have to be considered as potential corners and extremal points of $\mathcal{E}(r)$. The following algorithm clarifies that we can exclude all of these points straight from the beginning, leading to a faster computation of $\mathcal{E}(r)$.

Definition 6.3. For a compact and convex set $\mathcal{W} \subseteq \mathcal{V}^*$, let $\tilde{\mathcal{B}}(\mathcal{W})$ denote the largest closed, convex subset $\mathcal{X}$ of $\mathcal{V}^*$ such that:

(i) Outside of $\mathcal{Q} := \bigcup_{a \in \mathcal{A}} (\mathcal{S}_a(\mathcal{W}) \cup \partial\mathcal{K}_a(\mathcal{W})) \cap \mathcal{Q}_a$, the boundary $\partial\mathcal{X}$ is a solution to (6) within $\Gamma(\mathcal{W})$ and to (9) outside of $\Gamma(\mathcal{W})$, where admissible incentives in (9) and (6) additionally have to satisfy $N^\top \delta(y) \leq 0$ for every $y \in Y$.

(ii) Within $\mathcal{Q}$, the boundary $\partial\mathcal{X}$ satisfies Condition (ii) of Theorem 5.11.

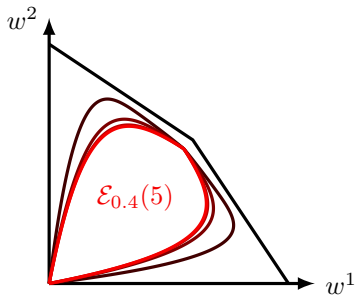

Figure 21: Convergence of the algorithm in Proposition 6.4.

Proposition 6.4. *Let $\mathcal{W}_0 = \mathcal{V}^*$ and $\mathcal{W}_n := \tilde{\mathcal{B}}(\mathcal{W}_{n-1})$ for $n \geq 1$. Then $(\mathcal{W}_n)_{n\geq 0}$ is decreasing in the set-inclusion sense with $\bigcap_{n\geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*

By comparing Figures 11 and 21, we observe that the algorithm in Proposition 6.4 converges significantly faster than the algorithm in Proposition 3.8.

7 Conclusion

This paper provides a methodology for computing the PPE payoff set $\mathcal{E}(r)$ and for constructing extremal PPE in continuous-time repeated games, in which information may arrive gradually and continuously or abruptly and discontinuously. The characterization uncovers that there are three fundamental ways, in which incentives are provided, based on the flow of information. First, incentives are provided discontinuously when a discrete event occurs. Second, the absence of discrete events is used to provide incentives continuously and deterministically. Finally, dynamic and continuous incentives can be provided through gradual information arrival. The use of a continuous-time model allows us to avoid the artifact of a discrete time period, which limits incentives to the first kind.

Discontinuous information is modeled through Poisson processes in this paper. However, the methodology does not rely on any specifics of a Poisson process other than it being a purely discontinuous process with totally inaccessible jump times. In future research, it would be interesting to characterize the boundary of the PPE payoff set for information that arrives according to general Lévy processes. Poisson processes arise naturally in many settings, in which timing of an event is the only crucial information. For example, consider stochastic games with a finite state space, in which players can observe the state but not the chosen actions. The state process in such a game is precisely described by a set of Poisson processes, whose intensities govern the rate at which the underlying state changes. It would thus be interesting to generalize the methods to stochastic games.

The optimality equations that describe the boundary of $\mathcal{E}(r)$ describe the trade-off between incentives provided through the different types of information. Due the quantitative nature of the result, the characterization may provide useful for future applied 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.

8 Data availability statement

There are no new data associated with this article. The code underlying this research is available on Zenodo at `https://dx.doi.org/10.5281/zenodo.7992464`.

References

Abreu, Dilip, and Yuliy Sannikov. 2014. "An algorithm for two-player repeated games with perfect monitoring." *Theoretical Economics*, **9(2)**: 313–338.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1986. "Optimal cartel equilibria with imperfect monitoring." *Journal of Economic Theory*, **39(1)**: 251–269.

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica*, **58(5)**: 1041–1063.

Abreu, Dilip, Paul Milgrom, and David Pearce. 1991. "Information and Timing in Repeated Partnerships." *Econometrica*, **59(6)**: 1713–1733.

Anderson, Robert M. 1984. "Quick-response equilibrium." *Working Papers in Economic Theory and Econometrics #IP-323, Center for Research in Management, University of California, Berkeley.*

Bernard, Benjamin, and Christoph Frei. 2016. "The folk theorem with imperfect public information in continuous time." *Theoretical Economics*, **11(2)**: 411–453.

Fudenberg, Drew, and David K. Levine. 1994. "Efficiency and Observability with Long-Run and Short-Run Players." *Journal of Economic Theory*, **62(1)**: 103–135.

Fudenberg, Drew, and David K. Levine. 2007. "Continuous time limits of repeated games with imperfect public monitoring." *Review of Economic Dynamics*, **10(2)**: 173–192.

Fudenberg, Drew, David K. Levine, and Eric Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica*, **62(5)**: 997–1039.

Green, Edward J., and Robert H. Porter. 1984. "Noncooperative Collusion under Imperfect Price Information." *Econometrica*, **52(1)**: 87–100.

Hashimoto, Tadashi. 2010. "Corrigendum to "Games With Imperfectly Observable Actions in Continuous Time"." *Econometrica*, **78(3)**: 1155–1159.

Hauser, Christine, and Hugo Hopenhayn. 2008. "Trading Favors: Optimal Exchange and Forgiveness." Collegio Carlo Alberto Carlo Alberto Notebooks **88**.

Jacod, J., and A. Shiryaev. 2002. *Limit Theorems for Stochastic Processes. Grundlehren der mathematischen Wissenschaften*. 2 ed., Springer Berlin Heidelberg.

Judd, Kenneth L., Şevin Yeltekin, and James Conklin. 2003. "Computing Supergame Equilibria." *Econometrica*, **71(4)**: 1239–1254.

Kalesnik, Vitali. 2005. "Continuous Time Partnerships with Discrete Events." Job Market Paper.

Kallenberg, Olav. 2002. *Foundations of modern probability. Probability and its Applications.* Second ed., Springer-Verlag, New York.

Karatzas, I., and S. Shreve. 2017. *Methods of Mathematical Finance. Probability Theory and Stochastic Modelling*, Springer New York.

Kazamaki, N. 2006. *Continuous Exponential Martingales and BMO. Lecture Notes in Mathematics*, Springer Berlin Heidelberg.

Munkres, J.R. 2000. *Topology. Featured Titles for Topology.* 2 ed., Prentice Hall, Incorporated.

Protter, P.E. 2005. *Stochastic Integration and Differential Equations. Stochastic Modelling and Applied Probability.* 2 ed., Springer Berlin Heidelberg.

Rockafellar, R.Tyrrell. 1985. "Lipschitzian properties of multifunctions." *Nonlinear Analysis: Theory, Methods & Applications*, **9(8)**: 867–885.

Sannikov, Yuliy. 2007. "Games with Imperfectly Observable Actions in Continuous Time." *Econometrica*, **75(5)**: 1285–1329.

Sannikov, Yuliy, and Andrzej Skrzypacz. 2007. "Impossibility of Collusion under Imperfect Monitoring with Flexible Production." *American Economic Review*, **97(5)**: 1794–1823.

Sannikov, Yuliy, and Andrzej Skrzypacz. 2010. "The Role of Information in Repeated Games With Frequent Actions." *Econometrica*, **78(3)**: 847–882.

Simon, Leo K., and Maxwell B. Stinchcombe. 1989. "Extensive Form Games in Continuous Time: Pure Strategies." *Econometrica*, **57(5)**: 1171–1214.

Sun, Yeneng. 2006. "The exact law of large numbers via Fubini extension and characterization of insurable risks." *Journal of Economic Theory*, **126(1)**: 31–69.

Yen, N. D. 1995. "Lipschitz Continuity of Solutions of Variational Inequalities with a Parametric Polyhedral Constraint." *Mathematics of Operations Research*, **20(3)**: 695–708.

A Proofs of auxiliary results in the main text

A.1 Dynamics of the continuation value and continuation promises

The proofs of the results in Section 3 follow the proofs of the analogous results in Bernard and Frei (2016) rather closely, extending them to information structures that also contain abrupt information.

We begin with the stochastic differential representation of the continuation value in Lemma 3.1. Intuitively, we wish to apply a martingale representation result to the continuation value at time ∞ under the family Q^A of probability measures induced by play of A. However, since the limiting measure Q^A_∞ may not be absolutely continuous with respect to P on $\mathcal{F}_\infty$, we instead use a martingale representation at finite times and show that the representation is time consistent.[23]

Proof of Lemma 3.1. Fix a strategy profile A and let $W := W(A)$. We will show that W^i satisfies (2) for $i = 1, 2$. Observe first that W is bounded as it remains in $\mathcal{V}$ at all times. Fix a player i and a time $T > 0$, and define the bounded random variable $w^i_T := W^i_T - r\int_0^T (W^i_t - g^i(A_t))\,\mathrm{d}t$. Because $(J^y)_{y\in Y}$ are pairwise orthogonal and orthogonal to Z, the stable subspace generated by Z and $(J^y)_{y\in Y}$ is the space of all stochastic integrals with respect to these processes by Theorem IV.36 in Protter (2005). Therefore, we obtain the unique martingale representation property for square-integrable martingales by Corollary 1 to Theorem IV.37 in Protter (2005). That is, there exist an $\mathcal{F}_0$-measurable c^i_T, predictable and square-integrable processes $(\beta^i_{t,T})_{0\le t\le T}$ and $\left(\delta^i_{t,T}(y)\right)_{0\le t\le T}$ for $y \in Y$ with $\mathbb{E}_{Q^A_T}\left[\int_0^T \left|\beta^i_{t,T}\right|^2 \mathrm{d}t\right] < \infty$ and $\mathbb{E}_{Q^A_T}\left[\int_0^T \left|\delta^i_{t,T}(y)\right|^2 \lambda(y\,|A_t)\,\mathrm{d}t\right] < \infty$ 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

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

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

$$w^i_T = W^i_T + r\int_0^T g^i(A_t)\,\mathrm{d}t - r\int_0^\infty \int_0^{s\wedge T} r\mathrm{e}^{-r(s-t)}\mathbb{E}_{Q^A_s}\left[g^i(A_s)\,\middle|\,\mathcal{F}_t\right]\mathrm{d}t\,\mathrm{d}s. \tag{17}$$

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

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

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

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

[23]The existence of such a limiting measure Q^A_∞, which coincides with Q^A_t on $\mathcal{F}_t$ for every $t \ge 0$, is given by Proposition 1.7.4 in Karatzas and Shreve (2017).

which yields $\beta^i_{\cdot,T} = \beta^i_{\cdot,\tilde{T}}$ and $\delta^i_{\cdot,T}(y) = \delta^i_{\cdot,\tilde{T}}(y)$ for every $y \in Y$ a.e. on $[0,\tilde{T}]$ and $M^i_{\tilde{T},T} = M^i_{\tilde{T},\tilde{T}}$ a.s. by the uniqueness of the orthogonal decomposition. Taking $\mathcal{F}_t$-conditional expectations, we deduce $M^i_{t,\tilde{T}} = M^i_{t,T}$ a.s. for $t \in [0,\tilde{T}]$, proving that the integral representation is independent of T. We thus omit the subscript T and $\tilde{T}$ of the constructed processes β^i, $(\delta^i(y))_{y\in Y}$, and M^i, hence W^i satisfies (2). To show the converse, we derive from Itō's formula that

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

Since M^i is strongly orthogonal to Z and $(J^y)_{y\in Y}$, it is also strongly orthogonal to the density process given in Footnote 7. Therefore, M^i is a martingale also under the family of probability measures Q^A. Integrating (18) from t to T and taking Q^A_T-conditional expectations on $\mathcal{F}_t$ thus yields

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

Since W is bounded, the second term converges to zero a.s. as $T\to\infty$, which completes the proof. □

Proof of Lemma 3.3. Fix a strategy profile A and let $\tilde{A}$ be a strategy profile involving a unilateral deviation of some player i, that is, $\tilde{A}^{-i} = A^{-i}$ a.e. For (β,δ) related to $W(A)$ by (2), integrating (18) from t to $u > t$ yields

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

Taking $\mathcal{F}_t$-conditional expectations under $Q^{\tilde{A}}_u$ causes the the $\mathrm{d}M^i$ term to disappear since M is a martingale up to time u also under $Q^{\tilde{A}}_u$. Moreover, $\mathrm{e}^{-r(u-t)}\mathbb{E}_{Q^A_u}\left[W^i_u(A)\,|\,\mathcal{F}_t\right]$ vanishes as $u\to\infty$ because $W(A)$ remains in the bounded set $\mathcal{V}$, hence we obtain

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

Because β is constructed using a martingale representation result for the bounded random variable w^i_T in (17), $\int_t^{\cdot} r\mathrm{e}^{-r(s-t)}\beta^i_s\left(\mathrm{d}Z_s - \mu(A_s)\,\mathrm{d}s\right)$ is a bounded mean oscillation (BMO) martingale under the probability measure Q^A_u up to any time $u > t$. Therefore, Theorem 3.6 in Kazamaki (2006) implies that $\int_t^{\cdot} r\mathrm{e}^{-r(s-t)}\beta^i_s\left(\mathrm{d}Z_s - \mu(\tilde{A}_s)\,\mathrm{d}s\right)$ is a martingale under $Q^{\tilde{A}}_u$. Since $W(A)$ lies in $\mathcal{V}$ $P\otimes$ *Lebsgue*-a.e., each $\delta(y)$ is uniformly bounded P-a.s. and hence also $Q^{\tilde{A}}_u$-a.s. for any $u > t$. The

lemma after Theorem IV.29 in Protter (2005) thus implies that $\int_t^\cdot re^{-r(s-t)}\delta_s^i(y)\left(\mathrm{d}J_s^y - \lambda(y \mid A_s)\,\mathrm{d}s\right)$ is a $Q_u^{\tilde{A}}$-martingale up to any time $u > t$. Together with Fubini's theorem, this implies

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

If (β, δ) enforces A, the above conditional expectation is non-positive, hence A is a PPE. To show the converse, assume towards a contradiction that there exist a player i, a set $\Xi \subseteq \Omega \times [0,\infty)$ with $P \otimes Lebesgue(\Xi) > 0$, and a strategy $\hat{A}^i$ such that

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

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

Proof of Lemma 3.5. The proof of Lemma 2 in Bernard and Frei (2016) works for signals given by any Lévy process. It thus directly applies here. □

A.2 Concatenation of $\mathcal{W}$-enforceable solutions

The main idea behind the proof of Lemma 3.7 is the following. Any payoff pair in $\mathcal{B}(\mathcal{W})$ can be attained by an en enforceable solution to (2) that remains in $\mathcal{B}(\mathcal{W})$ until the time σ of the first event and $W_\sigma \in \mathcal{W}$. Since $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$, we would like to attain W_σ by another enforceable solution W to (2) until the arrival of the next event and then concatenate the two solutions. Such a concatenation, however, is subject to some subtle measurability issues that we will address here.

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

$$\mathcal{B}(\mathcal{W}) = \left\{ w \in \mathcal{V} \;\middle|\; \begin{array}{l} \text{There exists a stochastic framework } (\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y}) \\ \text{containing a solution } (W, A, \beta, \delta, M) \text{ to (2) on } [\![0, \sigma]\!] \text{ with} \\ W_0 = w \text{ and } W_\sigma \in \mathcal{W} \text{ } P\text{-a.s., and on } [\![0,\sigma)\!), W \in \mathcal{B}(\mathcal{W}) \text{ and} \\ (\beta, \delta) \text{ enforces } A, \text{where } \sigma \text{ is the first jump time of } (J^y)_{y\in Y}. \end{array} \right\}.^{24,25}$$

When we aim to concatenate solutions at time σ, the continuation solutions may live on separate probability spaces for each realization of W_σ. It is thus not a priori clear that there exists a probability space that contains the entire concatenation. The following result establishes that this

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

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

is possible: because the probability spaces share the same path space, we can construct a regular conditional probability that allows us to aggregate the different probability spaces.

Lemma A.1. *Let X be a $\mathcal{V}$-valued random variable with distribution ν. The following are equivalent:*

(a) $X \in \mathcal{B}(\mathcal{W})$ *ν-a.s.,*

(b) There exists a solution (W, A, β, δ, M) to (2) in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$ such that X is $\mathcal{F}_0$-measurable with $\nu = P \circ X^{-1}$, $W_0 = X$ P-a.s., $W_\sigma \in \mathcal{W}$ P-a.s., and on $[\![0, \sigma)\!)$, we have $W \in \mathcal{B}(\mathcal{W})$ and (β, δ) enforces A, where σ is the first jump time of $(J^y)_{y\in Y}$.

Proof. We first show that (a) implies (b). Observe that we can parametrize a weak solution to (2) completely through $(A, W, Z, (J^y)_{y\in Y})$ since the processes β, δ, and M can be recovered by the unique martingale representation theorem. For each realization w of W^*, the path space of a solution is thus $\mathcal{A}^{[0,\infty)} \times \mathcal{D}^{2+d+|Y|}$, where $\mathcal{D}^{2+d+|Y|}$ is the space of càdlàg functions from $[0, \infty)$ into $\mathbb{R}^{2+d+|Y|}$. By Theorem A.2.2 in Kallenberg (2002), there exists a metric on $\mathcal{D}^{2+d+|Y|}$ that induces the Skorohod topology, under which $\mathcal{D}^{2+d+|Y|}$ is complete and separable. Since $\mathcal{A}^{[0,\infty)}$ is compact by Tychonov's theorem (see Theorem 37.3 in Munkres (2000)), it follows that $\Omega := \mathcal{V} \times \mathcal{A}^{[0,\infty)} \times \mathcal{D}^{2+d+|Y|}$ is complete and separable. Thus, by Theorem V.3.19 in Karatzas and Shreve (2017), there exists a regular conditional probability $(w, F) \mapsto P_w(F)$ for $(w, F) \in \mathcal{V} \times \mathcal{F}$, that is,

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

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

(iii) $P_w(F) = P(F \,|\, X = w)$ for each $F \in \mathcal{F}$ and ν-a.e. $w \in \mathcal{V}$.

For each $w \in \mathcal{B}(\mathcal{W})$, by definition there exists a weak solution $(W^w, A^w, \beta^w, \delta^w, M^w, Z^w, (J^{w,y})_{y\in Y})$ to (2) such that $W_0^w = w$ and on $[\![0, \sigma^w)\!)$, we have $W^w \in \mathcal{B}(\mathcal{W})$, $w + r\delta^w(y) \in \mathcal{W}$, and (β^w, δ^w) enforces A^w, where σ^w is the first jump time of any of the processes $J^{w,y}$. Define W pointwise as W^w on $\{X = w\}$ for each $w \in \mathcal{V}$. Define A, β, δ, M, Z, and $(J^y)_{y\in Y}$ similarly. Let Ξ denote the set on which $(W, A, \beta, \delta, M, Z, (J^y)_{y\in Y})$ satisfies (2) with $W_0 = X$ such that before the first jump time of any of the processes $(J^y)_{y\in Y}$, (β, δ) enforces A, $W + r\delta(y) \in \mathcal{W}$ for every $y \in Y$, and $W \in \mathcal{B}(\mathcal{W})$. It follows from the properties of a regular conditional probability that

$$P(\Xi) = \int_{\mathcal{B}(\mathcal{W})} P(\Xi \,|\, W^* = w)\, \mathrm{d}\nu(w) = 1. \tag{20}$$

For the converse, let $(W, A, \beta, \delta, M, Z, (J^y)_{y\in Y})$ be a solution to (2) as indicated in (b). Let Ξ be defined as above and set $\Xi(w) := \Xi \cap \{X = w\}$. Observe that $P_w(\Xi(w)) = 1$ if and only if $w \in \mathcal{B}(\mathcal{W})$. If X lies outside of $\mathcal{B}(\mathcal{W})$ with positive ν-measure, then $P(\Xi) < 1$ by (20), a contradiction. □

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

Lemma A.2. *Fix two sets $\mathcal{W}, \mathcal{X} \subseteq \mathbb{R}^2$ such that for every $w \in \mathcal{X}$ there exists a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$ containing a solution (W, A, β, δ, M) to (2) on $[\![0, \tau \wedge \sigma]\!]$ for an $\mathbb{F}$-stopping time τ with $W_0 = w$ P-a.s. such that*

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

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

where σ is the first jump time of $(J^y)_{y\in Y}$. Then for every $w \in \mathcal{X}$, there exists a stochastic framework $(\hat{\Omega}, \hat{\mathcal{F}}, \hat{\mathbb{F}}, \hat{P}, \hat{Z}, (\hat{J}^y)_{y\in Y})$ containing a solution $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ to (2) that coincides with (W, A, β, δ, M) on $[\![0, \tau \wedge \sigma))$, such that $(\hat{\beta}, \hat{\delta})$ enforces $\hat{A}$ on $[\![0, \hat{\sigma})\!)$, $\hat{W} \in \mathcal{B}(\mathcal{W})$ on $[\![\tau, \hat{\sigma})\!)$, and $\hat{W}_{\hat{\sigma}} \in \mathcal{W}$ P-a.s., where $\hat{\sigma}$ is the first jump time of $(\hat{J}^y)_{y\in Y}$. In particular, $\mathcal{X} \subseteq \mathcal{B}(\mathcal{W})$.

Proof. Fix $w \in \mathcal{X}$, a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$, an $\mathbb{F}$-stopping time τ, and a solution (W, A, β, δ, M) to (2) on $[\![0, \tau]\!]$ with all the stated properties. By Lemma A.1, on $\{\tau < \sigma\}$ there exists a stochastic framework $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{\mathbb{F}}, \tilde{P}, \tilde{Z}, (\tilde{J}^y)_{y\in Y})$ containing a $\mathcal{W}$-enforceable solution $(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}, \tilde{M})$ to (2) such that $\tilde{W}_0$ is distributed under $\tilde{P}$ as W_τ under P. As in the proof of Lemma A.1, without loss of generality Ω and $\tilde{\Omega}$ are Polish so that there exists a regular conditional probability $\hat{P}$ on $\hat{\Omega} = \Omega \times \tilde{\Omega}$ with marginal P on Ω. Let $\hat{\mathbb{F}} := (\hat{\mathcal{F}}_t)_{t\geq 0}$ denote the $\hat{P}$-augmented filtration such that $\hat{\mathcal{F}}_t$ contains $\mathcal{F}_{t\wedge\tau} \vee \tilde{\mathcal{F}}_{t-\tau\vee 0}$. Define the processes $\hat{Z}$ and $\hat{J}^y$ by setting

$$\hat{Z}_t = Z_{t\wedge\tau} + \tilde{Z}_{t-\tau} 1_{\{t>\tau\}}, \qquad\qquad \hat{J}^y_t = J^y_{t\wedge\tau} + \tilde{J}^y_{t-\tau} 1_{\{t>\tau\}}.$$

Because Brownian motion and Poisson processes have independent and identically distributed increments, $\hat{Z}$ is an $\hat{\mathbb{F}}$-Brownian motion and $\hat{J}^y$ for any $y \in Y$ is an $\hat{\mathbb{F}}$-Poisson process. Define the concatenated processes $\hat{W}$, $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, and $\hat{M}$ by setting

$$\hat{W} := (W 1_{[\![0,\tau]\!]} + \tilde{W}_{\cdot-\tau} 1_{(\!(\tau,\infty)\!)}) 1_{\{\tau<\sigma\}} + W 1_{\{\tau\geq\sigma\}}$$

and similarly for $\hat{A}$, $\hat{\beta}$, $\hat{\delta}$, and $\hat{M}$. By construction, $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ coincides with (W, A, β, δ, M) on $[\![0, \tau \wedge \sigma))$. By the properties of a regular conditional probability, $\hat{W}_0 = w$ and $\hat{W}_{\hat{\sigma}} \in \mathcal{W}$ $\hat{P}$-a.s. Moreover, $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M})$ has all the desired properties $\hat{P} \otimes$ *Lebesgue*-a.e. on $[\![0, \hat{\sigma})\!)$. □

Remark A.1. Relevant in the proof of Lemma A.2 is that $\hat{\sigma}$ is the time of the first event strictly after the time τ of the concatenation—if a concatenation occurs. For general sets $\mathcal{W}$, the event that a concatenation occurs is $\{\tau < \sigma\}$. The qualifier $\{\tau < \sigma\}$ in (ii) is needed because W_σ may attain values in $\mathcal{W} \setminus \mathcal{B}(\mathcal{W})$, in which case we cannot concatenate the solutions. If $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$, however, then the proof carries through for $\sigma = \tau$ because $W_\sigma \in \mathcal{B}(\mathcal{W})$. A concatenation occurs with certainty in this case (hence the qualifier $\{\tau < \sigma\}$ in (ii) is not needed) and yields a $\mathcal{W}$-enforceable solution on $[\![0, \sigma_2)\!)$, where σ_2 is the time of the second event. This is essentially the proof of Lemma 3.7.

Proof of Lemma 3.7. Suppose first that $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$. Because Poisson processes have only countably many jumps, an iteration of the procedure in Remark A.1 yields an enforceable solution to (2) that remains in $\mathcal{B}(\mathcal{W})$ forever, showing that $\mathcal{B}(\mathcal{W})$ is self-generating. If $\mathcal{W}$ is self-generating, then any

$w \in \mathcal{W}$, is attainable by an enforceable solution W to (2) with $W \in \mathcal{W}$ almost everywhere. In particular, $W_\sigma \in \mathcal{W}$ a.s. It follows that $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$ by maximality of $\mathcal{B}(\mathcal{W})$. □

A.3 Convergence of the algorithm

Before we prove the convergence of the algorithm in Proposition 3.8 to $\mathcal{E}(r)$, we need one more auxiliary result, stating that $\mathcal{B}(\mathcal{W})$ is monotone in $\mathcal{W}$.

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

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

The proof of Proposition 3.8 relies on the characterization of $\mathcal{B}(\mathcal{W})$ in Theorem 5.11, as well as properties of $\mathcal{W}$-enforceable solutions that attain extremal payoffs of $\mathcal{B}(\mathcal{W})$ that are derived in the following two appendices. Specifically, we need that any $w \in \mathcal{B}(\mathcal{W})$ can be attained by an enforceable solution to (2) with bounded quadratic variation; see Lemmas B.12 and C.5.

Proof of Proposition 3.8. We first show that $\mathcal{B}(\mathcal{V}^*) \subseteq \mathcal{V}^*$, i.e., each player i attains at least their minmax payoff in $\mathcal{B}(\mathcal{V}^*)$. Suppose towards a contradiction that this is not the case, i.e., there exists $w \in \mathcal{B}(\mathcal{V}^*)$ with $w^i < \underline{v}^i$ for some i. Since $w \in \mathcal{B}(\mathcal{V}^*)$, there exists a solution (W, A, β, δ, M) to (2) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$ with $W_0 = w$ a.s., $W_\sigma \in \mathcal{V}^*$ a.s., such that on $[\![0, \sigma)\!)$, we have $W \in \mathcal{B}(\mathcal{V}^*)$ and (β, δ) enforces A. By Lemma 3.1, we can write

$$w^i = \int_0^\sigma r \mathrm{e}^{-rs} \mathbb{E}_{Q_s^A}\Big[g^i(A_s)\Big]\,\mathrm{d}s + \mathrm{e}^{-r\sigma}\mathbb{E}_{Q_\sigma^A}\Big[W_\sigma^i\Big]. \tag{21}$$

Since $w^i < \underline{v}^i$ but $W_\sigma^i \geq \underline{v}^i$ a.s., it follows that the first term in (21) is smaller than $(1-\mathrm{e}^{-r\sigma})\underline{v}^i$. Define now the strategy $\tilde{A}^i = \arg\max_{a^i \in \mathcal{A}^i} g^i(\,\cdot\,, A^{-i}) 1_{[\![0,\sigma)\!)} + A^i 1_{[\![\sigma,\infty)\!)}$. Because $\tilde{A}^i$ is i's myopic best reply to A^{-i} on $[\![0, \sigma)\!)$, it follows that $g^i(\tilde{A}^i, A^{-i}) \geq \underline{v}^i$ on that same stochastic interval. We deduce that

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

This contradicts the fact that (β, δ) enforces A on $[\![0, \sigma)\!)$.[26] We have thus shown that $\mathcal{W}_1 = \mathcal{B}(\mathcal{V}^*)$ is contained in $\mathcal{W}_0 = \mathcal{V}^*$. Since $\mathcal{E}(r) \subseteq \mathcal{W}_0$, monotonicity of $\mathcal{B}$ implies that $\mathcal{B}(\mathcal{E}(r)) \subseteq \mathcal{W}_1$. It follows from Lemma 3.7 that $\mathcal{E}(r) \subseteq \mathcal{B}(\mathcal{E}(r))$, hence $\mathcal{E}(r) \subseteq \mathcal{B}(\mathcal{E}(r)) \subseteq \mathcal{W}_1 \subseteq \mathcal{W}_0$. An iterated application of Lemma A.3 thus shows that $(\mathcal{W}^n)_{n \geq 0}$ is decreasing in the set-inclusion sense and that it is bounded from below by $\mathcal{E}(r)$. It must converge to a limit $\mathcal{W}_\infty \supseteq \mathcal{E}(r)$ that satisfies $\mathcal{B}(\mathcal{W}_\infty) \subseteq \mathcal{W}_\infty$ by monotonicity. The last step of the proof is to show that $\mathcal{W}_\infty \subseteq \mathcal{B}(\mathcal{W}_\infty)$ so that $\mathcal{W}_\infty = \mathcal{B}(\mathcal{W}_\infty)$ is self-generating by Lemma 3.5 and, hence, contained in $\mathcal{E}(r)$.

[26] This follows along the same lines as the proof of Lemma 3.3 by restricting attention to $[\![0, \sigma)\!)$.

To that end, fix $w \in \mathcal{W}_\infty$. Since $w \in \mathcal{W}_n$ for each n, there exist enforceable solutions $(W^n, A^n, \beta^n, \delta^n, M^n)$ to (2) on some stochastic basis $(\Omega^n, \mathcal{F}^n, \mathbb{F}^n, Z^n, J^n, P^n)$ such that $W^n \in \mathcal{W}_n$ a.e. on $[\![0, \sigma)\!)$ and $W^n_\sigma \in \mathcal{W}_{n-1}$ a.s. Through the use of public randomization, we may choose W^n that remains on the boundary of ∂W^n on $[\![0, \sigma)\!)$. Lemmas 4.3 and 4.5 and Theorem 5.11 imply that public randomization is needed only at time 0 and at boundary payoffs that are either strictly decomposable or that are attainable only through limit strategies. The latter two are both attained by $\mathcal{W}$-enforceable solutions to (2) constructed in the proof of Lemma C.5, hence we may choose M^n with finite variation. Because β^n is given by the maximizer of the optimality equation (9) where $\beta^n \neq 0$, Lemma B.12 shows that there exists a constant $K > 0$ such that $\|\beta^n\| \leq K$, i.e., the quadratic variation of the continuation value is uniformly bounded. We will show that the solutions converge in law along a subsequence to a limit (W, A, β, δ, M) that solves (2). Convergence in law then implies that the limit has the remaining desired properties. Indeed, if there exists a set $\Xi \subseteq [\![0, \sigma)\!)$ of positive $P \otimes Lebesgue$-measure on which $W \notin \mathcal{W}_\infty$, then closedness of $\mathcal{W}_\infty$ implies that there exists n sufficiently large, for which $W^n \notin \mathcal{W}_n$ with positive measure, a contradiction. Similarly, convergence in law implies $W_\sigma \in \mathcal{W}_\infty$ a.s., and that (β, δ) enforce A a.e. on $[\![0, \sigma)\!)$.

Convergence of $(W^n)_{n\geq 0}$ in law along a subsequence is shown in two steps. First, we show that the sequence is tight by verifying that the associated semimartingale characteristics satisfy the sufficient conditions of Theorem VI.5.17 in Jacod and Shiryaev (2002). Second, Prokhorov's theorem states that any tight sequence converges in law along a subsequence.[27] We show convergence in law of $\tilde{W}^n := (e^{-rt} W^n_t)_{t\geq 0}$ because its semimartingale characteristics or more easily computed. Convergence of $(\tilde{W}^n)_{n\geq 0}$ in law then implies convergence in law of $(W^n)_{n\geq 0}$. Since each $\tilde{W}^n$ is uniformly bounded under Q_{A^n}, it is a so-called special semimartingale, admitting a unique decomposition $\tilde{W}^n = W^n_0 + B^n + \tilde{M}^n$ into a predictable process B^n of finite variation and a Q_{A^n}-local martingale $\tilde{M}^n$. The semimartingale characteristics of $\tilde{W}^n$ under Q_{A^n} is the triplet (B^n, C^n, ν^n), where C^n is the predictable quadratic variation of $\tilde{M}^n$ and ν^n is the compensated jump measure of $\tilde{W}^n$; see Chapter II of Jacod and Shiryaev (2002) for a detailed introduction.[28] It follows from (18) that the first two characteristics of $\tilde{W}^n$ are given by

$$B^n = -\int_0^\cdot r\mathrm{e}^{-rt} g(A^n_t)\,\mathrm{d}t, \qquad\qquad C^{n,ij} = \int_0^\cdot r^2 \mathrm{e}^{-2rt} \beta^{n,i\top}_t \beta^{n,j}_t\,\mathrm{d}t.$$

Define the auxiliary processes $G^n := \mathrm{Var}(B^{n,1}) + \mathrm{Var}(B^{n,2}) + C^{n,11} + C^{n,22}$ for each n. Since we chose β^n uniformly bounded by K, it follows that each G^n is majorized by

$$F := \int_0^\cdot r\mathrm{e}^{-rs} \max_{a\in\mathcal{A}} \left| g^1(a) + g^2(a) \right| \mathrm{d}s + 2\int_0^\cdot r^2 \mathrm{e}^{-2rs} K\,\mathrm{d}s.^{29}$$

Aldous' criterion implies that $(G^n)_{n\geq 0}$ is tight; see Theorem VI.4.5 in Jacod and Shiryaev (2002). Thus, the sequence $(G^n)_{n\geq 0}$ converges in law along a subsequence to some limit process G by

[27] See, for example, Theorem VI.3.5 in Jacod and Shiryaev (2002).

[28] Because each $\tilde{W}^n$ is bounded, we do not need to truncate its jumps to compute the semimartingale characteristics. In the notation of Jacod and Shiryaev (2002), we can choose "truncation function" $h(x) = x$, hence the second modified characteristic coincides with the second characteristic.

[29] A process F majorizes G if $F - G$ is increasing.

Prokhorov's theorem. Since F is deterministic, F majorizes G. Together with the fact that $(\tilde{W}_0^n)_{n\geq 0} = (M_0^n)_{n\geq 0}$ is tight because each M_0^n takes values in the compact set $\mathcal{V}$, this shows that $(\tilde{W}^n)_{n\geq 0}$ is tight by virtue of Theorem VI.5.17 in Jacod and Shiryaev (2002).

With the same argument and majorizing process F, it follows that $(B^n)_{n\geq 0}$ and $(\tilde{M}^n)_{n\geq 0}$ are tight. Since the path space $\mathcal{A}^{[0,\infty)}$ of any A^n is compact by Tychonov's theorem, the sequence $(A^n)_{n\geq 0}$ is uniformly tight. By Prokhorov's theorem, there exists a subsequence $(n_k)_{k\geq 0}$, along which $(\tilde{W}^n)_{n\geq 0}$, $(A^n)_{n\geq 0}$, $(B^n)_{n\geq 0}$, and $(\tilde{M}^n)_{n\geq 0}$ all converge to limits $\tilde{W}$, A, B, and $\tilde{M}$, respectively. Convergence in law implies that $-\int_0^{\cdot} re^{-rt} g(A_t)\,\mathrm{d}t = \tilde{B}$. After a suitable application of Girsanov's theorem, the laws of $(Z^n)_{n\geq 0}$ and $(J^{y,n})_{n\geq 0}$ are constant and, hence converge trivially to Brownian motion Z and Poisson processes J^y along $(n_k)_{k\geq 0}$. By Proposition IX.1.1 in Jacod and Shiryaev (2002), the limit process $\tilde{M}$ of $(\tilde{M}^n)_{n\geq 0}$ is an $(\mathbb{F}, Q^A)$-martingale. Let β and δ be defined by the martingale representation of $\tilde{M}$. By Theorem VI.6.26 in Jacod and Shiryaev (2002), $[\tilde{M}^{n_k}, \tilde{M}^{n_k}] \to [\tilde{M}, \tilde{M}]$ in law, showing that $\beta^{n_k} \to \beta$ and $\delta^{n_k} \to \delta$ in law as well. This concludes the proof that any payoff pair $w \in \mathcal{W}_\infty$ can be attained by an enforceable solution to (2) that remains in $\mathcal{W}_\infty$ up until and including time σ. Thus, $\mathcal{W}_\infty$ is $\mathcal{W}_\infty$-relaxed self-generating, hence it is contained in $\mathcal{B}(\mathcal{W}_\infty)$ by maximality. □

Proof of Proposition 6.4. Observe first that the operators $\mathcal{S}$ and $\mathcal{K}_a$ are monotone, that is, $\mathcal{S}(\mathcal{W}) \subseteq \mathcal{S}(\mathcal{W}')$ and $\mathcal{K}_a(\mathcal{W}) \subseteq \mathcal{K}_a(\mathcal{W}')$ for two payoff sets $\mathcal{W} \subseteq \mathcal{W}'$. Because (9) is solved over a larger set of controls in $\tilde{\mathcal{B}}(\mathcal{W}')$ than in $\tilde{\mathcal{B}}(\mathcal{W})$, it follows that $\tilde{\mathcal{B}}(\mathcal{W}') \supseteq \tilde{\mathcal{B}}(\mathcal{W})$, i.e., $\tilde{\mathcal{B}}$ is monotone as well. The same reasoning shows that $\tilde{\mathcal{B}}(\mathcal{W}) \subseteq \mathcal{B}(\mathcal{W})$ for a fixed payoff set $\mathcal{W}$, hence $\mathcal{W}_1 = \tilde{\mathcal{B}}(\mathcal{W}_0) \subseteq \mathcal{B}(\mathcal{W}_0) \subseteq \mathcal{W}_0$. Together with monotonicity of $\tilde{\mathcal{B}}$, 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 6.2 asserts that $\tilde{\mathcal{B}}(\mathcal{E}(r)) = \mathcal{B}(\mathcal{E}(r)) = \mathcal{E}(r)$. Monotonicity of $\tilde{\mathcal{B}}$ 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 $\mathcal{B}$ to $\mathcal{V}^*$. Since $\tilde{\mathcal{B}}(\mathcal{W}) \subseteq \mathcal{B}(\mathcal{W})$ for any set $\mathcal{W}$, it follows that $\mathcal{W}_1 \subseteq \mathcal{W}'_1$ and hence $\mathcal{W}_{n+1} \subseteq \tilde{\mathcal{B}}(\mathcal{W}'_n) \subseteq \mathcal{B}(\mathcal{W}'_n) = \mathcal{W}'_{n+1}$ 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 equations

The purpose of this appendix is to prove that the optimality equations (6) and (9) are locally Lipschitz continuous so that they admit unique solutions that are continuous in initial conditions. The main idea is to show that both ODEs are a maximization of a locally Lipschitz continuous function over a locally Lipschitz continuous set of available incentives.

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

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

Proof. For any x, let U be a neighborhood 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

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

□

While, in general, the intersection of two locally Lipschitz continuous maps need not be locally Lipschitz, or even lower semicontinuous, the following lemma establishes that local Lipschitz continuity is preserved in the interior of the effective domain for closed and convex-valued maps.

Lemma B.3. *Let F and G be closed- and convex-valued maps that are locally Lipschitz continuous at $x_0 \in \operatorname{dom} F \cap \operatorname{dom} G$, at least one of them locally bounded at x_0. If $F(x_0) \cap \operatorname{int} G(x_0) \neq \emptyset$, then $F \cap G$ is locally Lipschitz continuous at x_0.*

Proof. Since $F(x_0)$ intersects the interior of $G(x_0)$, the two sets cannot be separated by a hyperplane through any $y \in \partial(F(x_0) \cap G(x_0))$. Theorem 4.12 of Rockafellar (1985) implies that $F \cap G$ is pseudo-Lipschitzian at (x_0, y). Since y was arbitrary, Theorem 2.2 of Rockafellar (1985) implies that $F \cap G$ is sub-Lipschitzian at x_0. Because $F \cap G$ is locally bounded, Theorem 2.1 of Rockafellar (1985) implies that $F \cap G$ is locally Lipschitz at x_0. □

B.1 Abrupt-information optimality equation

For any action profile a, let Ψ_a^i denote the set of all δ^i that solve the enforceability constraint (3) for player $i = 1, 2$. We view $\Psi_a = \Psi_a^1 \times \Psi_a^2$ as a subset of $\mathbb{R}^{2|Y|}$ with typical element $\delta = (\delta^1, \delta^2)$. Note that Ψ_a is the intersection of closed half-spaces with normal vectors $\Delta\lambda_a^1(\tilde{a}^1) := (\lambda(\tilde{a}^1, a^2) - \lambda(a), 0)$ for $\tilde{a}^1 \in \mathcal{A}^1 \setminus \{a^1\}$ and $\Delta\lambda_a^2(\tilde{a}^2) := (0, \lambda(a^1, \tilde{a}^2) - \lambda(a))$ for $\tilde{a}^2 \in \mathcal{A}^2 \setminus \{a^2\}$. Let $(\frac{\mathcal{W}-w}{r})^{|Y|}$ denote the set of all $\delta \in \mathbb{R}^{2|Y|}$ with $w + r\delta(y) \in \mathcal{W}$ for each $y \in Y$. Since $\mathcal{W}$ is typically fixed, let us abbreviate $\Psi_a(w) := \Psi_a(w, \mathcal{W}) = \Psi_a \cap (\frac{\mathcal{W}-w}{r})^{|Y|}$. Recall that $\mathcal{K}_a(\mathcal{W})$ is the set of all payoff pairs w, for which $\Psi_a(w)$ is non-empty and that $\mathcal{K}(\mathcal{W}) = \bigcup_a \partial\mathcal{K}_a(\mathcal{W})$. If $\mathcal{W}$ is closed or convex, respectively, then so is $\Psi_a(w)$ and $\mathcal{K}_a(\mathcal{W})$. We begin with the following two observations.

Lemma B.4.

(i) If there exists $w \in \mathcal{K}_a(\mathcal{W})$, for which Ψ_a intersect the interior of $(\frac{\mathcal{W}-w}{r})^{|Y|}$, then Ψ_a intersects the interior of $(\frac{\mathcal{W}-v}{r})^{|Y|}$ for every $v \in \operatorname{int} \mathcal{K}_a(\mathcal{W})$.

(ii) Let $H_a^i(\tilde{a}^i)$ denote the hyperface of Ψ_a with normal vector $\Delta\lambda_a^i(\tilde{a}^i)$. If $\Psi_a(w) \subseteq H_a^i(\tilde{a}^i)$ for some $w \in \operatorname{int} \mathcal{K}_a(\mathcal{W})$, then $\Psi_a(v) \subseteq H_a^i(\tilde{a}^i)$ for all $v \in \mathcal{K}_a(\mathcal{W})$.

Proof. For the first statement, fix $v \in \mathcal{K}_a(\mathcal{W})$ and $\delta_v \in \Psi_a \cap \operatorname{int}(\frac{\mathcal{W}-v}{r})^{|Y|}$. Because Ψ_a is constant and $(\frac{\mathcal{W}-w}{r})^{|Y|}$ is Lipschitz continuous, there exists an open neighborhood U of v such that Ψ_a

intersects the interior of $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$ for any $w \in U$. In particular, $v \in \operatorname{int} \mathcal{K}_a(\mathcal{W})$. By convexity of $\mathcal{K}_a(\mathcal{W})$, we can write any other $w \in \operatorname{int} \mathcal{K}_a(\mathcal{W})$ as a strict convex combination $w = xv + (1-x)u$ for some $u \in \mathcal{K}_a(\mathcal{W})$. Convexity of Ψ_a and $\mathcal{W}$ imply that $\delta_* := x\delta_v + (1-x)\delta_u \in \Psi_a \cap \operatorname{int}\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$.

For the second statement, fix $w \in \operatorname{int} \mathcal{K}_a(\mathcal{W})$, a player i, and a deviation $\tilde{a}^i$ with $\Psi_a(w) \subseteq H_a^i(\tilde{a}^i)$. Suppose towards a contradiction that there exists $v \in \mathcal{K}_a(\mathcal{W})$ and $\delta_v \in \Psi_a(v)$ for which deviations to $\tilde{a}^i$ by player i are deterred strictly. Write w as a strict convex combination $w = xv + (1-x)u$ of v and $u \in \mathcal{K}_a(\mathcal{W})$. Convexity of Ψ_a and $\mathcal{W}$ imply that $\delta_* := x\delta_v + (1-x)\delta_u \in \Psi_a(w)$ for arbitrary $\delta_u \in \Psi_a(u)$. Since deviations to $\tilde{a}^i$ are deterred strictly by δ_*, this is a contradiction. □

Lemma B.5. *$w \mapsto \Psi_a(w)$ is continuous on $\mathcal{K}_a(\mathcal{W})$ and locally Lipschitz continuous in* $\operatorname{int} \mathcal{K}_a(\mathcal{W})$.

Proof. We first show local Lipschitz continuity. If $\mathcal{K}_a(\mathcal{W})$ contains a single element v, for which Ψ_a intersects the interior of $\left(\frac{\mathcal{W}-v}{r}\right)^{|Y|}$, then $w \mapsto \Psi_a(w)$ is locally Lipschitz continuous in the interior of $\mathcal{K}_a(\mathcal{W})$ by Lemmas B.3 and B.4. Suppose, therefore, that Ψ_a does not intersect the interior of $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$ for any $w \in \mathcal{K}_a(\mathcal{W})$. Then $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$ moves parallel to the hyperfaces of Ψ_a that contain $\Psi_a(w)$. Since the set of binding hyperfaces does not change in the interior of $\mathcal{K}_a(\mathcal{W})$ by Lemma B.4, the maximal change in $\Psi_a(w)$ is bounded by the maximal change in $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$.

It remains to show that $w \mapsto \Psi_a(w)$ is continuous at the boundary of $\mathcal{K}_a(\mathcal{W})$. Because it is the intersection of two upper semicontinuous maps, it is again upper semicontinuous. For lower semicontinuity, fix $w \in \partial\mathcal{K}_a(\mathcal{W})$ and let $\tilde{\mathcal{A}}^i$ be the set of those deviations to $\tilde{a}^i$ for which incentives bind at w. Thus, the hyperfaces of Ψ_a that intersect $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$ are precisely those with normal vectors $\Delta\lambda_a^i(\tilde{a}^i)$ for $\tilde{a}^i \in \tilde{\mathcal{A}}^1 \cup \tilde{\mathcal{A}}^2$. Let $(w_n)_{n\geq 0} \subseteq \mathcal{K}_a(\mathcal{W})$ be any sequence approximating w. Since $w_n \in \mathcal{K}_a(\mathcal{W})$, there must exist at least one face of Ψ_a that intersects $\left(\frac{\mathcal{W}-w_n}{r}\right)^{|Y|}$. Since Ψ_a has finitely many hyperfaces, by passing to a subsequence we may assume that the same set of faces intersect $\left(\frac{\mathcal{W}-w_n}{r}\right)^{|Y|}$ for any $n \geq 0$. Moreover, those hyperfaces are a subset of $\tilde{\mathcal{A}}^1 \cup \tilde{\mathcal{A}}^2$. Let $\hat{\Psi}_a(v)$ denote the intersection of those faces with $\left(\frac{\mathcal{W}-v}{r}\right)^{|Y|}$. Then $\Psi_a(w) \subseteq \hat{\Psi}_a(w)$ and $\hat{\Psi}_a(w_n) \subseteq \Psi_a(w_n)$ for any n.

Observe that any δ in the relative interior of $\hat{\Psi}_a(w)$ is in the relative interior of $\hat{\Psi}_a(w_n)$ for n sufficiently large since Ψ_a is constant and $\left(\frac{\mathcal{W}-w_n}{r}\right)^{|Y|}$ converges to $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$. In particular, any $\delta_* \in \Psi_a(w)$ that lies in the relative interior of $\hat{\Psi}_a(w)$ can be approximated by a constant sequence $\delta_n = \delta_*$ for n sufficiently large. If $\hat{\Psi}_a(w) = \{\delta_*\}$ is a singleton (hence so is $\Psi_a(w)$), then an arbitrary sequence in $\hat{\Psi}_a(w_n)$ approaches δ_*. If $\hat{\Psi}_a(w)$ is not a singleton, fix arbitrary δ_0 in the relative interior of $\hat{\Psi}_a(w)$ (hence in the relative interior of $\hat{\Psi}_a(w_n)$ for n sufficiently large) and parametrize δ in the relative boundary of $\hat{\Psi}_a(v)$ through the direction e_δ of $\delta - \delta_0$. Any $\delta_* \in \Psi_a(w)$ on the relative boundary of $\hat{\Psi}_a(w)$ can then be approximated by δ_n on the relative boundary of $\hat{\Psi}_a(w_n)$ parametrized by e_{δ_*}. □

B.2 Optimality equation

The main results of this appendix are Lemmas B.9 and B.10. They establish regularity of (9) in a fashion suitable for the proofs in Appendix C.1 that characterize the boundary of $\mathcal{B}(\mathcal{W})$ outside of

$\Gamma(\mathcal{W})$. Consider the optimality equation first for a fixed action profile a:

$$\kappa_a(w,N) = \max_{(\beta,\delta)\in\Xi_a(w,N,\mathcal{W})} \frac{2N^\top(g(a)+\delta\lambda(a)-w)}{r\|\beta\|^2}. \tag{22}$$

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

Lemma B.6. *For any $a\in\mathcal{A}$, the map $(N,\delta)\mapsto\phi(a,N,\delta)$ is locally Lipschitz continuous for $N\notin\{\pm e_1,\pm e_2\}$ and $\delta\in\Upsilon_a$.*

Proof. Let $G^i(a)$ denote the row vector with entries $g^i(\tilde{a}^i,a^{-i})-g^i(a)$ and let $\Lambda^i(a)$ denote the matrix with column vectors $\lambda(\tilde{a}^i,a^{-i})-\lambda(a)$. For non-coordinate N, $\Phi_a(N,\delta)$ is the solution to $\phi M^i(a)\le-\frac{1}{T^i}(G^i(a)+\delta^i\Lambda^i(a))$. Since $M^i(a)$ does not depend on N or δ and the right-hand side is locally Lipschitz continuous in N,δ, the set $\Phi_a(N,\delta)$ is a constant-rank polyhedron with locally Lipschitz continuous right-hand side. It follows from the main result of Yen (1995) that the projection of 0 onto $\Phi_a(N,\delta)$ is locally Lipschitz continuous in (N,δ). □

For any $h\ge0$, let $\mathcal{W}_h:=\{v\in\mathbb{R}^2\mid\min_{w\in\mathcal{W}}\|v-w\|\le h\}$ denote the set of all payoff pairs within distance h from $\mathcal{W}$. Call a set-valued map $\mathcal{L}:\mathbb{R}^2\to\mathbb{R}^{2\times|Y|}$ a *Lipschitz expansion* of $\mathcal{W}$ if

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

for a non-negative Lipschitz-continuous function $h:\mathbb{R}^2\to\mathbb{R}$ with Lipschitz constant K_h. Observe that a Lipschitz expansion is Lipschitz continuous in the sense of set-valued maps with Lipschitz constant $\sqrt{|Y|}(K_h+1)$. We denote by $\mathcal{L}_0$ the trivial Lipschitz expansion with $h\equiv0$. Set $\Upsilon_a(w,N,\mathcal{L}):=\{\delta\in\Upsilon_a\cap\mathcal{L}(w)\mid\Phi_a(N,\delta)\neq\emptyset\}$ and consider the optimality equation in the form

$$\kappa_{a,\mathcal{L}}(w,N)=\max_{\delta\in\Upsilon_a(w,N,\mathcal{L})}\frac{2N^\top(g(a)+\delta\lambda(a)-w)}{r\|\phi(a,N,\delta)\|^2}\vee0. \tag{23}$$

Observe that ODE (23) for $\mathcal{L}_0$ reduces to (22) if the expression in the right-hand side of (22) is understood to be 0 when $\Upsilon_a(w,N,\mathcal{L})=\emptyset$. Denote by

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

the effective domain of of $(w,N)\mapsto\Upsilon_a(w,N,\mathcal{L})$.

Lemma B.7. *Suppose that Assumptions 1 and 3 are satisfied. Then there exists $\varepsilon>0$ such that $\mathcal{W}_\varepsilon\times(S^1\setminus\{\pm e_1,\pm e_2\})\subseteq E_a(\mathcal{L})$ for any action profile a with $E_a(\mathcal{L})\neq\emptyset$ and any Lipschitz expansion $\mathcal{L}$. Moreover, $(w,N)\mapsto\Upsilon_a(w,N,\mathcal{L})$ is locally Lipschitz continuous on $\mathcal{W}_\varepsilon\times(S^1\setminus\{\pm e_1,\pm e_2\})$.*

Proof. Any enforceable action profile a is pairwise identifiable by Assumption 1. Fix any $\delta \in \Upsilon_a$ and any non-coordinate N. It follows from Lemma 2 in Sannikov (2007) applied to payoff function $\tilde{g}(a) = g(a) + \delta\lambda(a)$ that there exists β with $N^\top\beta = 0$ such that (β, δ) enforces a. In particular, $\Phi_a(N, \delta)$ is non-empty, hence $\Upsilon_a(w, N, \mathcal{L}) = \Upsilon_a \cap \mathcal{L}(w)$ for any non-coordinate N. This shows that $E_a(\mathcal{L})$ contains $\{w \mid \Upsilon_a \cap \mathcal{L}(w) \neq \emptyset\} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$ and that $N \mapsto \Upsilon_a(w, N, \mathcal{L})$ is locally Lipschitz continuous because it is locally constant.

By Assumption 3, there exists ε such that $\delta \in \Upsilon_a$ for any δ with $\|\delta(y)\| \leq \varepsilon$ and any enforceable action profile a. In particular, for any $w \in \mathcal{W}_{r\varepsilon/2}$ there exists $\delta \in \Upsilon_a$ with $w + r\delta(y) \in \operatorname{int} \mathcal{W}$ for every event y. This shows that $E_a(\mathcal{L})$ contains $\mathcal{W}_{r\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. Moreover, Υ_a intersects the interior of $\mathcal{L}(w)$, hence local Lipschitz continuity for $w \in \mathcal{W}_{r\varepsilon/2}$ follows from Lemma B.3. □

The last ingredient to establish local Lipschitz continuity of $\kappa_{a,\mathcal{L}}$ for non-coordinate directions is to characterize the set of payoff-direction pairs, for which the right-hand side of (23) is unbounded. This is the case on the set $\Gamma_a(\mathcal{L})$ of all (w, N), for which there exists a $\delta \in \Upsilon_a(w, N, \mathcal{L})$, for which $\phi(a, N, \delta) = 0$ and $N^\top(g(a) + \delta\lambda(a) - w) \geq 0$. Set $\Gamma(\mathcal{L}) := \bigcup_{a \in \mathcal{A}} \Gamma_a(\mathcal{L})$.

Remark B.1. Note that $\Gamma(\mathcal{L}_0) = \Gamma(\mathcal{W})$, where $\Gamma(\mathcal{W})$ is defined in Section 5.1.

Lemma B.8. *Let $\mathcal{L}$ be a Lipschitz expansion. Then $\Gamma_a(\mathcal{L})$ is closed for any $a \in \mathcal{A}$. In particular, $\Gamma(\mathcal{L})$ is closed as the finite union of closed sets.*

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

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

Lemma B.9. *Suppose that Assumptions 1 and 3 hold. For any convex $\mathcal{W}$ with $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$ and any Lipschitz expansion $\mathcal{L}$, there exists $\varepsilon > 0$ such that*

$$\kappa_{\mathcal{L}}(w, N) = \max_{a \in \mathcal{A}} \kappa_{a,\mathcal{L}}(w, N) \tag{24}$$

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

Proof. Let $\varepsilon > 0$ be given by Lemma B.7 so that $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \subseteq E_a(\mathcal{L})$ for any action profile a with $E_a(\mathcal{L}) \neq \emptyset$ and that $\Upsilon_a(w, N, \mathcal{L})$ is locally Lipschitz continuous on $\mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. It follows from Lemma B.2 that $\kappa_{a,\mathcal{L}}$ is locally Lipschitz continuous at any $(w, N) \in \mathcal{W}_\varepsilon \times (S^1 \setminus \{\pm e_1, \pm e_2\}) \setminus \Gamma(\mathcal{L})$ for any action profile a. Therefore, $\kappa_{\mathcal{L}}$ is locally Lipschitz continuous except at $\bigcup_{a \in \mathcal{A}} \partial E_a(\mathcal{L})$. Since the intersection of $\bigcup_{a \in \mathcal{A}} \partial E_a(\mathcal{L})$ with $\mathcal{W}_{\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$ is empty by Lemma B.7, $\kappa_{\mathcal{L}}$ is locally Lipschitz continuous on $\mathcal{W}_{\varepsilon/2} \times (S^1 \setminus \{\pm e_1, \pm e_2\})$. □

Let us turn our attention to coordinate directions. The following lemma implies that coordinate payoffs of $\partial\mathcal{B}(\mathcal{W})$ are isolated points, hence $\partial\mathcal{B}(\mathcal{W})$ is well-behaved at those payoffs.

Lemma B.10. *Suppose that Assumptions 2 and 3 are satisfied. Then there exists $\varepsilon > 0$ such that for any $w \in \mathcal{W}_\varepsilon \cap \operatorname{int} \mathcal{V}^*$ and any $N \in \{\pm e_1, \pm e_2\}$, there exists $a \in \mathcal{A}$ such that $\kappa_a(w, N) > 0$ in a neighborhood of $\mathcal{U}$ of (w, N) with $\mathcal{U} \subseteq E_a(\mathcal{L}_0)$.*

Proof. By Assumptions 2 and 3, the minmax profile $\underline{a}_i$ against player i and the global maximizer $\bar{a}_i$ of g^i are enforceable and there exists $\varepsilon > 0$ such that $\Delta_\varepsilon := \{\delta \mid \|\delta(y)\| \leq \varepsilon \text{ for every } y \in Y\}$ is contained in $\Upsilon_{a(N)}$ for every $N \in \{\pm e_1, \pm e_2\}$, where we denote $a(e_i) = \bar{a}_i$, $a(-e_i) = \underline{a}_i$. Let $G^i(a)$ and $\Lambda^i(a)$ be defined as in the proof of Lemma B.6. If Assumption 2.(i) is satisfied, we additionally choose ε small enough such that $G^i(a(N)) + \delta^i \Lambda^i(a(N)) < 0$ for player i with static best reply in $a(N)$ and $\delta \in \Delta_\varepsilon$. Let $k \leq 1$ be such that for any $w \in \mathcal{W}_{rk\varepsilon} \cap \partial\mathcal{V}^*$ with outward normal vector N, there exists $\delta \in \Delta_\varepsilon$ with $\delta^i = 0$ such that $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$. For any $w \in \mathcal{W}_{rk\varepsilon} \cap \operatorname{int} \mathcal{V}^*$, choose $\delta \in \Delta_\varepsilon$ with $w + r\delta(y) \in \mathcal{W}$ for every event $y \in Y$ such that $N^\top(g(a(N)) + \delta\lambda(a(N)) - w) > 0$. Since $\delta \in \Upsilon_{a(N)}$, there exists β such that (β, δ) enforces $a(N)$. If Assumption 2.(i) holds, the choice of ε implies that $(\frac{\tilde{T}}{\tilde{T}^{-i}}\beta^{-i}, \delta)$ enforces $a(N)$ for $\tilde{T}$ orthogonal to $\tilde{N}$ sufficiently close to N. In particular, $(w, N) \in \operatorname{int} E_{a(N)}(\mathcal{L}_0)$ and $\kappa_a(w, N) > 0$ since the numerator is strictly positive. If Assumption 2.(ii) holds, then player i may have more than one static best reply to $a^{-i}(N)$. Let $a_+(N)$ and $a_-(N)$ denote action profiles a with $a^{-i} = a^{-i}(N)$ that maximize and minimize $\sum_y \lambda(y \mid a)$, respectively, among all static best replies to $a^{-i}(N)$. For any $w \in \mathcal{W}_{rk\varepsilon}$, choose now $\delta \in \Delta_\varepsilon$ with $w + r\delta(y) = v$ for some $v \in \mathcal{W}$ with $N^\top(g(a(N)) + \delta\lambda(a(N)) - w) > 0$. If $v^i \geq w^i$, then $G^i(a_+(N)) + \delta^i \Lambda^i(a_+(N)) \leq 0$ and if $v^i \leq w^i$, then $G^i(a_-(N)) + \delta^i \Lambda^i(a_-(N)) \leq 0$. Since $\delta \in \Upsilon_{a(N)}$, there exists β such that (β, δ) enforces $a(N)$. Let β' denote the projection of β^{-i} onto $\operatorname{span} M^{-i}(a)$. By Assumption 2.(ii), $\frac{\tilde{T}}{\tilde{T}^{-i}}\beta' M^i(a) = 0$, hence $(\frac{\tilde{T}}{\tilde{T}^{-i}}\beta', \delta)$ enforces $a(N)$ for $\tilde{T}$ orthogonal to $\tilde{N}$ sufficiently close to N. □

Corollary B.11. *$\bigcup_{a \in \mathcal{A}} \Xi_a(w, N, \mathcal{W})$ is non-empty for any (w, N) with $w \in \mathcal{W}_\varepsilon \cap \operatorname{int} \mathcal{V}^*$.*

Finally, the following two lemmas establish bounds on incentives related to the Brownian information. Those bounds are used to show convergence in the algorithm of Proposition 3.8 and to derive the escaping lemma in Lemma C.2.

Lemma B.12. *There exists $K > 0$ such that $F(w, N) := \min_{a \in \mathcal{A} \setminus \mathcal{A}^N} \max_{\delta \in \Upsilon_a(w,N)} \|\phi(a, N, \delta)\| \leq K$ for any (w, N) with $w \in \mathcal{W}_\varepsilon \cap \operatorname{int} \mathcal{V}^*$.*

Proof. Let $\mathcal{A}_* = \{\bar{a}_1, \bar{a}_2, \underline{a}_1, \underline{a}_2\}$ and denote by $\tilde{\kappa}_a$ the optimality equation restricted to $\delta = 0$. Assumption 2 implies via Lemma 4 in Sannikov that $\max_{a \in \tilde{A}} \tilde{\kappa}_a(w, N)$ is locally Lipschitz continuous for N in a neighborhood $\mathcal{N}$ of all coordinate directions. Choose $\mathcal{N}$ such that $\max_{a \in \mathcal{A}} \kappa_a(w, N) \geq \max_{a \in \tilde{A}} \tilde{\kappa}_a(w, N) \geq \varepsilon > 0$ for all $N \in \mathcal{N}$, showing that $F(w, N)$ is uniformly bounded for all $N \in \mathcal{N}$. By Lemma B.6, $\phi(a, N, \delta)$ is Lipschitz continuous, hence bounded, for $N \in \operatorname{cl} \mathcal{N}^c$. □

Lemma B.13. *Let $\mathcal{C} \subseteq \mathcal{W}$ be a C^1-solution to (24) for Lipschitz expansion $\mathcal{L}$ oriented by $w \mapsto N_w$ such that $\mathcal{N}_\mathcal{C} \cap (\Gamma(\mathcal{L}) \cup \mathcal{P}) = \emptyset$. For every $a \in \mathcal{A}$, there exists $K_a > 0$ such that for any $(w, N_w) \in \mathcal{N}_\mathcal{C} \cap$*

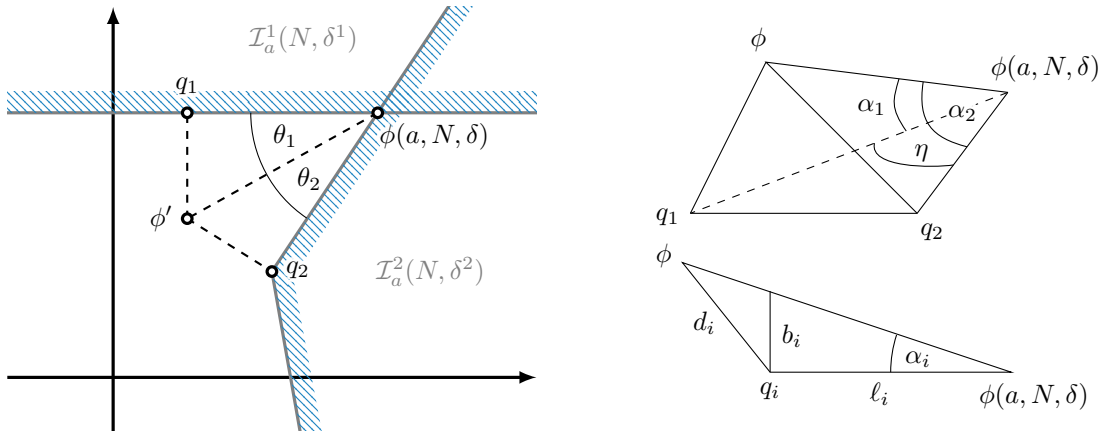

Figure 22: The left panel illustrates the positions of $\phi(a,N,\delta)$, ϕ', q_1, and q_2 relative to $\mathcal{I}_a^1(N,\delta^1)$ and $\mathcal{I}_a^2(N,\delta^2)$. The top right panel illustrates that $\alpha_j \geq \theta_j \geq \eta/2$. The bottom right panel illustrates the identity $\|\phi(a,N,\delta)-\phi\| \leq d_i(1+1/\tan(\alpha_i))$: because q_i is the projection of ϕ onto $\mathcal{I}_a^i(N,\delta^i)$, the angle at q_i is at least 90°. Therefore, $d_i \geq b_i = \ell_i \tan(\alpha_i)$ and the desired identity follows from the triangle inequality.

$E_a(\mathcal{L})$, and any $(T_w\phi + N_w\chi, \delta)$ enforcing a with $\delta \in \mathcal{L}(w)$ and ${N_w}^\top(g(a)+\delta\lambda(a)-w) \geq 0$, we have

$$1 - \frac{\big(\|\phi\| - \gamma\|\chi\|\big)^2}{\|\phi(a,N_w,\delta)\|^2} \leq \frac{2K_a + 2\gamma}{\bar{\Psi}_a}\|\chi\|, \tag{25}$$

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

Proof. Fix an action profile $a \in \mathcal{A}$. Since $\Gamma(\mathcal{L})$ is closed by Lemma B.8, $\mathcal{N}_\mathcal{C}$ is bounded away from $\Gamma(\mathcal{L})$, hence $\bar{\Psi}_a$ is indeed strictly positive. Observe that it is sufficient to show the existence of a constant $K_a > 0$ such that

$$\|\phi(a,N_w,\delta)\| - \|\phi\| \leq K_a\|\chi\| \tag{26}$$

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

$$1 - \frac{\|\phi\| - \gamma\|\chi\|}{\|\phi(a,N_w,\delta)\|} \leq \frac{\|\phi(a,N_w,\delta)\| - \|\phi\| + \gamma\|\chi\|}{\bar{\Psi}_a} \leq \frac{K_{a,2}(K_a+1)+\gamma}{\bar{\Psi}_a}\|\chi\|, \tag{27}$$

hence (25) follows from (27) in conjunction with $1 - x \geq \frac{1}{2}(1-x^2)$.

It remains to show (26). Clearly, the inequality is satisfied if $\|\phi\| \geq \|\phi(a,N,\delta)\|$, hence suppose that $\|\phi\| < \|\phi(a,N,\delta)\|$. For $i = 1,2$, define

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

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

For $i = 1, 2$, let q_i denote the point in $\mathcal{I}_a^i(N, \delta^i)$ closest to ϕ and let $d_i = \|\phi - q_i\|$. Note that $\phi + N^i/T^i\chi \in \mathcal{I}_a^i(N, \delta^i)$ for $i = 1, 2$ and hence $\left|N^i/T^i\right|\|\chi\| \geq d_i$. Let ϕ' be the projection of ϕ onto the plane through $\phi(a, N, \delta)$, q_1, and q_2. Let $j \in \{1, 2\}$ be the index i for which the angle θ_i between $\phi(a, N, \delta) - \phi'$ and $\phi(a, N, \delta) - q_i$ is maximal. Then $\theta_j \geq \eta/2$, where η is the angle between $\phi(a, N, \delta) - q_1$ and $\phi(a, N, \delta) - q_2$. Let α_i denote the angle between $\phi(a, N, \delta) - \phi$ and $\phi(a, N, \delta) - q_i$; See Figure 22 for an illustration. Observe that $\alpha_i \geq \theta_i$ and hence

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

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

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

Equation (26) now follows from the reverse triangle inequality. □

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

C.1 Characterization outside of $\Gamma(\mathcal{W})$

In this appendix, we show that the boundary of $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$ outside of $\Gamma(\mathcal{W})$ satisfies the optimality equation. Because the continuous part of the signal is what creates the curvature, these steps are similar in ideas to Sannikov (2007). We begin a generalization of Lemmas 4.3 and 4.5, showing how to construct Markovian strategy profiles whose continuation values remain on a given curve.

Lemma C.1. *Let $\mathcal{C}$ be a continuously differentiable curve oriented by $w \mapsto N_w$ such that there exist measurable selectors a_*, β_*, δ_* on $\mathcal{C}$ that for any $w \in \mathcal{C}$,*

(i) $(\beta_*(w), \delta_*(w)) \in \Xi_{a_*(w)}(w, N_w, \mathcal{W})$,

(ii) The curvature $\kappa(w)$ of $\mathcal{C}$ at w satisfies

$$\kappa(w)\|\beta_*(w)\|^2 = \frac{2}{r}N_w{}^\top(g(a_*(w)) + \delta_*(w)\lambda(a_*(w)) - w).$$

(iii) If $\beta_*(w) = 0$ for a segment of positive length, then $T_w{}^\top(g(a_*(w)) + \delta_*(w)\lambda(a_*(w)) - w)$ is either non-positive or non-negative throughout the segment.

Then the solution (W, A, β, δ, M) to (2) with $A = a_*(W_-)$, $\beta = \beta_*(W_-)$, $\delta = \delta_*(W_-)$, and $M \equiv 0$ is $\mathcal{W}$-enforceable and it remains on $\mathcal{C}$ until an end point of $\mathcal{C}$ is reached or a discrete event occurs.

Proof. Fix w in the relative interior of $\mathcal{C}$ and choose $\eta > 0$ small enough such that $N_w{}^\top N_v > 0$ for all $v \in \mathcal{C} \cap B_\eta(w)$, where $B_\eta(w)$ denotes the closed ball around w with radius η. On $B_\eta(w)$,

$\mathcal{C}$ admits a local parametrization f in the direction N_w. For any $v \in B_\varepsilon(w)$, define the orthogonal projection $\hat{v} = {T_w}^\top v$ onto the tangent and denote by $\pi(v) = (\hat{v}, f(\hat{v}))$ the projection of $v \in B_\eta(w)$ onto $\mathcal{C}$ in the direction N_w. Let $(W, A, \beta, \delta, Z, (J^y)_{y\in Y}, M)$ be a weak solution to (2) starting at $W_0 = w$ with $A = a_*(\pi(W_-))$, $\beta = \beta_*(\pi(W_-))$, $\delta = \delta_*(\pi(W_-))$, and $M \equiv 0$ on $[\![0, \tau[\![$, where we set $\tau := \sigma \wedge \inf\{t \geq 0 \mid W_t \notin B_\eta(w)\}$ for the first jump time σ of any of the processes $(J^y)_{y\in Y}$. By (iii), such a solution exists. Since π is measurable, the processes A, β and δ are all predictable.

We measure the distance of W to $\mathcal{C}$ by $D_t = {N_w}^\top W_t - f(\hat{W}_t)$. Note that f is differentiable by assumption and $(-f'(\hat{W}_t), 1) = \ell_t N_t$, where $\ell_t := \|(-f'(\hat{W}_t), 1)\|$. 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 Itō's formula that

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

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

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

where we used ${N_t}^\top (W_t - \pi(W_t)) = {N_t}^\top N_w D_t = \ell_t^{-1} D_t$ in the second equality. It follows that $D_t = D_0 \mathrm{e}^{rt} = 0$ since $D_0 = 0$. On $\{\tau < \sigma\}$ we can repeat this procedure and concatenate the solutions as in the proof of Lemma A.2 to obtain a solution (W, A, β, δ, M) to (2) with $W \in \mathcal{C}$ until time $\rho \wedge \sigma$, where ρ is the first time W reaches an end point of $\mathcal{C}$. Condition (i) implies $(\beta, \delta) \in \Xi_A(\pi(W), N_W, \mathcal{W})$ on $[\![0, \rho \wedge \sigma[\![$, which implies that the solution is $\mathcal{W}$-enforceable since $\pi(W) = W$ on $[\![0, \rho \wedge \sigma[\![$. □

Proofs of Lemmas 4.3 and 4.5. These are immediate corollaries to Lemma C.1 since $\Psi_a(w, \mathcal{W})$ is non-empty in $\mathcal{K}_a(\mathcal{W})$ and $\bigcup_{a\in\mathcal{A}} \Xi_a(w, N, \mathcal{W})$ is non-empty for any $(w, N) \in \mathcal{W}_\varepsilon \times S^1$ by Corollary B.11 and the selection of a maximizer is measurable. □

The following result is known as the "escaping lemma," stating that it is impossible for a solution $\mathcal{C}$ to the optimality equation to cut through $\mathcal{B}(\mathcal{W})$ as indicated in Figure 23. The main idea is that the optimality equation maximizes the curvature over all enforceable action profiles and all restricted-enforceable incentives. Thus, the continuation value of no enforceable strategy profile that attains a payoff pair above $\mathcal{C}$—in the shaded area of Figure 23—can cross $\mathcal{C}$ unless the continuation value jumps across $\mathcal{C}$ due to the occurrence of a discrete event. Because arrival times of Poisson jumps are arbitrarily large with positive probability and Brownian motion has infinite variation, the continuation value escapes $\partial\mathcal{B}(\mathcal{W})$ to the sides with positive probability before a discrete event occurs.

The presence of abrupt information complicates the argument in two ways. First, simply because Brownian incentives are needed to provide incentives on the curve does not imply that they are

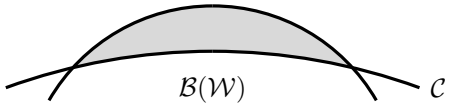

Figure 23: A solution $\mathcal{C}$ to (9) that cuts through $\mathcal{B}(\mathcal{W})$.

needed in the shaded area as well. Second, the curvature in the optimality equation is the maximum over all restricted-enforceable incentives at that point on the curve. Incentives provided in an enforceable strategy profile with continuation value in the shaded area may not come from the same set of incentives because the set of admissible jump incentives depends on the location of the payoff pair. Luckily, both problems can be solved by considering a Lipschitz expansion over jump incentives that guarantees that the set of admissible jump incentives on the curve contains the set of admissible jumps incentives in the shaded area.

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

$$\mathcal{L}_k(v) := \left\{ \delta \in \mathbb{R}^{2\times m} \;\middle|\; v + r\delta(y) \in \mathcal{W}_{k\|\pi(v)-v\|} \text{ for all } y \in Y \right\} \tag{28}$$

Let $\mathcal{C}$ be a C^1-solution to (24) for $\mathcal{L}_k$ with $k \geq 1$, oriented by $v \mapsto N_v$ with end points $v_L, v_R \in U$. It is impossible that the following properties hold simultaneously:

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

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

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

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

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

Proof. Suppose that there exists such a curve $\mathcal{C}$ for some $k \geq 1$. Since $\mathcal{L}_k$ is a Lipschitz expansion, it follows from Conditions (iv), (v) and Lemma B.9 that $\mathcal{C}$ is twice differentiable. By Condition (iii), there exists a local parametrization f of $\mathcal{C}$ in the direction N'. Let T' denote the counterclockwise rotation of N' by 90° and denote by $\hat{v} = {T'}^\top v$ the projection onto the tangent for any $v \in U$. Let $\hat{\pi} : U \to \mathcal{C}$ denote the projection of U onto $\mathcal{C}$ in direction N', i.e., $\hat{\pi}(v) = \big(\hat{v}, f(\hat{v})\big)$. By Condition (ii), there exists a solution $\big(W, A, \beta, \delta, M, Z, (J^y)_{y\in Y}\big)$ to (2) starting at $W_0 = v_0 + \eta N'$ such that on $[\![0, \sigma)\!)$, we have $W \in \mathcal{B}(\mathcal{W})$ and (β, δ) enforces A with $\delta \in \mathcal{L}_0(W)$. Let $\tau_1 := \inf\{t \geq 0 \mid W_t \notin U\}$ denote the time when W leaves the neighborhood U.

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

by Condition (iii). Then, similarly as in Footnote 3 of Hashimoto (2010), Itō's formula implies

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

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

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

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

Specifically, we show that there exists a probability measure R, absolutely continuous with respect to Q^A, under which the drift rate of D is bounded from below by rD. Then, D becomes arbitrarily large with positive R-probability, and hence positive Q^A-probability. Because it may take arbitrarily long until an infrequent event occurs, it follows that W escapes the bounded set $\mathcal{V}$ with positive probability. Let $\Xi_1 \subseteq [\![0, \tau_2)\!)$ denote the set of all (ω, t), for which ${N_t(\omega)}^\top(\hat\pi(W_t(\omega)) - g(A_t(\omega)) - \delta_t(\omega)\lambda(A_t(\omega))) \geq 0$. Since $\zeta_t(\omega) \geq rD_t(\omega)$ for $(\omega, t) \in \Xi_1$, there is no need to change the probability measure on Ξ_1. Let $\Xi_1^c = [\![0, \tau_2)\!) \setminus \Xi_1$ denote the complement of Ξ_1 relative to $[\![0, \tau_2)\!)$. Because W has not crossed $\mathcal{C}$ yet, it follows that $\|\hat\pi(W_t(\omega)) - W_t(\omega)\| \leq \|\hat\pi(W_t(\omega)) - \pi(W_t(\omega))\|$ for $(\omega, t) \in [\![0, \tau_2)\!)$ and hence $\delta_t(\omega) \in \Psi_{A_t(\omega)} \cap \mathcal{L}_k(\hat\pi(W_t(\omega)))$. This implies that $\beta_t(\omega) \neq 0$ for $(\omega, t) \in \Xi_1^c$ as otherwise $(\hat\pi(W_t(\omega)), N_t(\omega)) \in \Gamma(\mathcal{L}_k)$, contradicting Condition (v). It follows from Condition (iv) and Lemma B.7 that $(\hat\pi(W), N) \in E_A(\mathcal{L}_k)$ on $[\![0, \tau_2)\!)$. Since $\kappa(\hat\pi(W))$ solves (24), it follows that

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

Denote $\Lambda := \max_{a\in\mathcal{A}}\sum_{y\in Y}\lambda(y \mid a)$ and observe that ${N}^\top(g(A) + \delta\lambda(A) - \hat\pi(W))$ is uniformly bounded above by the constant $K_1 := \operatorname{diam}\mathcal{V} + \sup(\mathcal{W} - \mathcal{V})\Lambda < \infty$. Due to Lemma B.13, there exist constants $K_2, \bar\Psi$ such that

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

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

$$\frac{\mathrm{d}L_t}{L_t} = (\psi_t - \mu(A_t))(\mathrm{d}Z_t - \mu(A_t)\,\mathrm{d}t) - \sum_{y\in Y}(\mathrm{d}J_t^y - \lambda(y \mid A_{t-})\,\mathrm{d}t),$$

where

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

By the Lenglart-Girsanov theorem (Theorem III.38 in Protter (2005)), L is the density process of a probability measure R that is absolutely continuous with respect to Q_T^A on $\mathcal{F}_T$. Under R, each J_t^y is identically 0, $Z_t' = Z_t - \int_0^t \psi_s \, \mathrm{d}s$ is an R-Brownian motion on $[0,T]$ because $\int_0^T \|\psi_t\|^2 \, \mathrm{d}t < \infty$ Q_T^A-a.s., and $\tilde{M}_t$ is an R-martingale because it is orthogonal to L. Thus, under R, we have

$$D_t \geq D_0 + \int_0^t rD_s \, \mathrm{d}s + \int_0^t \xi_s^\top \, \mathrm{d}Z_s' + \tilde{M}_t. \tag{29}$$

Since W is bounded, $\int_0^t \xi_s^\top \, \mathrm{d}Z_s$ is a $BMO(Q_T^A)$-martingale. Therefore, $\int_0^t \xi_s^\top \, \mathrm{d}Z_s'$ is an R-martingale by Theorem 3.6 in Kazamaki (2006). Define the stopping time $\tau_3 := \inf\{t \geq 0 \mid D_t \leq D_0(1+rt)/2\}$ and observe that $\tau_3 \leq \tau_2 \wedge T$. Abbreviate by $F_t = \int_0^t \xi_s \, \mathrm{d}Z_s' + \tilde{M}_t$ the martingale part in (29). It follows from (29) that

$$D_{\tau_3} - \frac{D_0}{2}(1 + r\tau_3) \geq \frac{D_0}{2} + F_{\tau_3}.$$

Since $F_t = 0$, the optional stopping theorem implies

$$0 \geq \mathbb{E}_R\left[D_{\tau_3} - \frac{D_0}{2}(1 + r\tau_3)\right] \geq \mathbb{E}_R\left[\frac{D_0}{2} + F_{\tau_3}\right] = \frac{D_0}{2} > 0,$$

a contradiction. □

A first application of Lemma C.2 shows that the boundary of $\mathcal{B}(\mathcal{W})$ is smooth outside of $\Gamma(\mathcal{W})$. If it did have a corner, we could construct a solution to (24) that cuts through that corner, which is an impossibility by Lemma C.2. Additionally, Lemma C.2 implies that if $(w,N) \in \Gamma(\mathcal{W})$ for some normal vector $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$, then $(w,N) \in \Gamma(\mathcal{W})$ for all $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$.

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

Proof. Fix a corner w of $\mathcal{B}(\mathcal{W})$ and set $\mathcal{N}_0 := \operatorname{int} \mathcal{N}_w(\mathcal{B}(\mathcal{W})) \setminus \{\pm e_1, \pm e_2\}$. Fix $N' \in \mathcal{N}_0$ and let π denote the projection of a neighborhood U of w onto $\partial\mathcal{B}(\mathcal{W})$ in the direction of N'. For $\eta > 0$ and $k \geq 1$, let $\mathcal{C}_{\eta,k}$ be a C^1 solution to (22) for $\mathcal{L}_k$ defined in (28) starting at (w_η, N'), where $w_\eta := w - \eta N'$. Since N' is in the interior of $\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$, the curve $\mathcal{C}_{\eta,k}$ satisfies Conditions (i)–(iv) of Lemma C.2 for any $\eta > 0$ sufficiently small. Suppose towards a contradiction that $(w,N) \notin \Gamma(\mathcal{W})$ and hence $(w,N) \notin \Gamma(\mathcal{L}_k)$. Since $\Gamma(\mathcal{L}_k)$ is closed by Lemma B.8, there exists a neighborhood $\mathcal{U}$ of (w,N) with $\mathcal{U} \cap \Gamma(\mathcal{L}_k) = \emptyset$. Therefore, for η sufficiently small such that $\mathcal{N}_{\mathcal{C}_{\eta,k}} \subseteq \mathcal{U}$, Condition (v) of Lemma C.2 is satisfied. Such a curve cannot exist by Lemma C.2, contradicting $(w,N') \in \Gamma(\mathcal{W})$. Finally, $(w,N) \in \Gamma(\mathcal{W})$ for any $N \in \operatorname{cl} \mathcal{N}_0 = \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ since $\Gamma(\mathcal{W}) = \Gamma(\mathcal{L}_0)$ is closed by Lemma B.8. □

Finally, the following lemma is an integral part of the proof of Theorem 5.11, characterizing all boundary payoffs in $\mathcal{B}(\mathcal{W})$, at which Brownian information is needed.

Lemma C.4. *Suppose that $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$. The boundary $\partial\mathcal{B}(\mathcal{W})$ coincides with a solution to (9) in a neighborhood of any $(w,N) \in \mathcal{N}_{\mathcal{B}(\mathcal{W})} \setminus (\Gamma(\mathcal{W}) \cup \mathcal{P})$. Moreover, $\mathcal{N}_{\mathcal{B}(\mathcal{W})} \cap \mathcal{P} \setminus \Gamma(\mathcal{W})$ consists of at most 4 isolated points.*

Proof. Fix $(w, N) \in \mathcal{N}_{\mathcal{B}(\mathcal{W})} \setminus (\Gamma(\mathcal{W}) \cup \mathcal{P})$. Due to Lemma B.9, the optimality equation (9) is locally Lipschitz continuous at (w, N). It follows that (9) admits a unique C^2 solution that is continuous in initial values that lie in a sufficiently small neighborhood of (w, N). Let $\mathcal{C}$ denote the solution to (9) starting at (w, N) and suppose towards a contradiction that $\mathcal{C}$ escapes $\operatorname{cl}\mathcal{B}(\mathcal{W})$ in a neighborhood U of w. Without loss of generality, we may assume that $\mathcal{C}$ escapes $\operatorname{cl}\mathcal{B}(\mathcal{W})$ to the right of w. Let π denote the projection of U onto $\partial\mathcal{B}(\mathcal{W})$ in direction N. Define

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

and $\mathcal{L}(v) := \{\delta \in \mathbb{R}^{2\times m} \mid v + r\delta(y) \in \mathcal{W}(v) \text{ for all } y \in Y\}$. Then, $\mathcal{C}$ is also the unique solution to (24) starting in (w, N) for Lipschitz expansion $\mathcal{L}$. Since $(w, N) \notin \Gamma(\mathcal{W})$, the boundary $\partial\mathcal{B}(\mathcal{W})$ is C^1 at w by Lemma C.3. Therefore, a solution $\mathcal{C}'$ to (24) with initial conditions (w, N') for arbitrarily slight rotations N' of N cuts through $\mathcal{B}(\mathcal{W})$ as illustrated in Figure 23. Since $(w, N) \notin \Gamma(\mathcal{W}) \subseteq \Gamma(\mathcal{L})$ and $\Gamma(\mathcal{L})$ is closed, such a curve $\mathcal{C}'$ satisfies the conditions of Lemma C.2 for sufficiently slight rotations N'. Such a curve $\mathcal{C}'$ cannot exist by Lemma C.2, hence $\mathcal{C}$ cannot escape $\operatorname{cl}\mathcal{B}(\mathcal{W})$.

Suppose next that $\mathcal{C}$ falls into the interior of $\mathcal{B}(\mathcal{W})$ in a neighborhood of (w, N), that is, there exists $v \in \mathcal{C} \cap \operatorname{int}\mathcal{B}(\mathcal{W})$ arbitrarily close to w. By convexity of $\mathcal{B}(\mathcal{W})$, this is not possible if $\mathcal{C}$ is a trivial solution to (9), hence $\mathcal{C}$ is a solution with positive curvature. We may assume without loss of generality that this happens to the right of w as illustrated in the right panel of Figure 10. Let $\varepsilon > 0$ be sufficiently small such that $B_\varepsilon(v) \subseteq \operatorname{int}\mathcal{B}(\mathcal{W})$. By continuity in initial conditions, the solution $\mathcal{C}'$ to (9) with initial condition (w, N') for sufficiently slight counterclockwise rotations N' must contain $v_R \in B_\varepsilon(v)$. Moreover, $\mathcal{C}'$ enters $\mathcal{B}(\mathcal{W})$ to the left of w so that it contains a point $v_L \in \operatorname{int}\mathcal{B}(\mathcal{W})$. Finally, $\mathcal{C}'$ contains a point $w' \notin \operatorname{cl}\mathcal{B}(\mathcal{W})$ between w and v_R as illustrated in the right panel of Figure 10. Since $\mathcal{C}'$ satisfies the conditions of Lemma 4.5, it follows that $w' \in \mathcal{B}(\mathcal{W})$, a contradiction.

For the second statement, suppose towards a contradiction that $\partial\mathcal{B}(\mathcal{W})$ contains a straight line segment L of positive length with coordinate outward normal N with $L \times \{N\} \cap \Gamma(\mathcal{W}) = \emptyset$. Let w be in the relative interior of L. Due to Lemma B.10, there exists a with $(w, N) \in \operatorname{int} E_a(\mathcal{L}_0)$ and $\kappa_a > 0$ in a neighborhood of (w, N). Let $\mathcal{C}_\eta$ be a strictly curved solution to (9) starting at $(w + \eta N, N)$. For η sufficiently small, $\mathcal{N}_{\mathcal{C}_\eta} \subseteq E_a(\mathcal{L}_0)$ and $\mathcal{C}_\eta$ contains end points in $\operatorname{int}\mathcal{B}(\mathcal{W})$. Lemma 4.5 thus implies that $w + \eta N \in \mathcal{B}(\mathcal{W})$, which is a contradiction. □

C.2 Characterization within $\Gamma(\mathcal{W})$

In this appendix, we characterize all boundary points of $\mathcal{B}(\mathcal{W})$ within $\Gamma(\mathcal{W})$, where Brownian information is not needed to provide incentives. The first result, whose proof relies on an exact law of large numbers, characterizes payoffs that can be attained by limit strategies.

Lemma C.5. *Consider a payoff pair w, a set of action profiles $\mathcal{A}_w$, a set of weights $(x_a)_{a\in\mathcal{A}_w}$, and*

$(\delta_a)_{a\in\mathcal{A}_w}$ *with* $\delta_a \in \Psi_a(w,\mathcal{W})$ *for each* $a \in \mathcal{A}_w$. *If* w *coincides with*

$$d := \sum_{a\in\mathcal{A}_w} x_a(g(a) + \delta_a\lambda(a)), \tag{30}$$

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

Proof. Suppose that such $w, v, \mathcal{A}_w, (\delta_a)_{a\in\mathcal{A}_w}$, and $(x_a)_{a\in\mathcal{A}_w}$ exist. Consider a solution (W, A, δ, β, M) to (2) on a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y\in Y})$ satisfying with $W_0 = w$ such that:

(i) $(\Omega, \mathcal{F}, \mathbb{F}, P)$ is rich enough to admit partitions $\Xi_t = (\Xi_t(a))_{a\in\mathcal{A}_w}$ of Ω with $P(\Xi_t(a)) = x_a$ for each $a \in \mathcal{A}_w$ and each $t > 0$ such that $\Xi = (\Xi_t)_{t\geq 0}$ is independent of Z, $(J^y)_{y\in Y}$ and Ξ satisfies an exact law of large numbers; see Sun (2006),

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

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

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

It follows from an exact law of large numbers (see Proposition 2.5 in Sun (2006)) that

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

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

Lemma C.5 implies that the set $\mathcal{S}(\mathcal{W})$ of all stationary profiles and, in fact, all of $\bar{\mathcal{S}}(\mathcal{W})$ is contained in $\mathcal{B}(\mathcal{W})$. Next, we characterize the set $\mathcal{K}(\mathcal{W})$ of all payoffs, at which incentives are binding and extremal. Before we give the proof of Proposition 5.5, we state two preliminary lemmas.

Lemma C.6. *Let* $\mathcal{A}_w$ *strictly and minimally decomposes* $w \notin \mathcal{S}(\mathcal{W})$. *Then for each* $a \in \mathcal{A}_w$, *there exists* $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ *and* $\delta \in \Psi_a(w,\mathcal{W})$ *with* $N^\top(g(a) + \delta\lambda(a) - w) > 0$.

Proof. Fix such w and $\mathcal{A}_w$ and suppose that towards a contradiction that there exists $a \in \mathcal{A}_w$ such that $\mathcal{D}_a := g(a) + \Psi_a(w,\mathcal{W})\lambda(a)$ lies in lies in every lower half-space $H(w, N) = \{v \mid N^\top(v - w) \leq 0\}$ for $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$. Since w is not stationary, $w \notin \mathcal{D}_a(w)$, hence $N^\top(g(a) + \delta\lambda(a) - w) < 0$ for all non-extremal $N \in \mathcal{N}_w(\mathcal{B}(\mathcal{W}))$. By definition of strict decomposability, for any extremal normal vector N in $\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$, there is $\hat{a} \in \mathcal{A}_w$ and $\hat{\delta} \in \Psi_{\hat{a}}(w,\mathcal{W})$ with $N^\top(g(\hat{a}) + \hat{\delta}\lambda(\hat{a}) - w) > 0$, hence $\hat{a} \neq a$. It follows that $\mathcal{A}_w \setminus \{a\}$ decomposes w, contradicting minimality of $\mathcal{A}_w$. □

Lemma C.7. *Let* $w \in \partial\mathcal{B}(\mathcal{W}) \setminus \mathcal{S}(\mathcal{W})$ *be minimally and strictly decomposed by* $\mathcal{A}_w$. *Then* $w \in \mathcal{B}(\mathcal{W})$, $w \in \partial\mathcal{K}_a(\mathcal{W})$ *for each* $a \in \mathcal{A}_w$ *and* (12) *is satisfied.*

Figure 24: Since ${N'}^\top\big(g(a)+\delta_*\lambda(a)-w\big)\geq 0$ for slight clockwise rotations of N, $g(a)+\delta_*\lambda(a)$ must lie to the right of w and, hence, decompose w.

Proof. Fix such a payoff pair $w\in\partial\mathcal{B}(\mathcal{W})\setminus\mathcal{S}(\mathcal{W})$ and a set of action profiles $\mathcal{A}_w$. By Lemma C.6, for each $a\in\mathcal{A}_w$, there exist $N\in\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ and $\delta\in\Psi_a(w,\mathcal{W})$ with $N^\top\big(g(a)+\delta\lambda(a)-w\big)>0$. Thus, Lemma 4.2 implies that $w\in\partial\mathcal{K}_a(\mathcal{W})$ for each $a\in\mathcal{A}_w$. Strict decomposition implies that there exists a convex combination d of $g(a)+\delta_a\lambda(a)$ for $a\in\mathcal{A}_w$ that satisfies the conditions of Lemma C.5, hence $w\in\mathcal{B}(\mathcal{W})$. For the last statement, suppose towards a contradiction that $\partial\mathcal{B}(\mathcal{W})$ enters the interior of $\mathcal{K}_{\mathcal{A}_w}(\mathcal{W})$ and, hence, (12) is not satisfied. Since $w\mapsto\Psi_a(w,\mathcal{W})$ is continuous in $\mathcal{K}_a(\mathcal{W})$ by Lemma B.5, strict decomposition of w implies that there exists $a\in\mathcal{A}_w$ and $\delta\in\Psi_a(v,\mathcal{W})$ with ${N_v}^\top\big(g(a)+\delta\lambda(a)-v\big)>0$ for $v\in\partial\mathcal{B}(\mathcal{W})$ sufficiently close to w, contradicting Lemma C.6. □

Proof of Proposition 5.5. Fix $(w,N)\in\mathcal{N}_{\mathcal{B}(\mathcal{W})}\cap\Gamma(\mathcal{W})$ such that $w\in\mathcal{K}(\mathcal{W})\setminus\mathcal{S}(\mathcal{W})$. Lemma C.3 implies that $(w,N')\in\Gamma(\mathcal{W})$ for every $N'\in\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$. In particular, the payoff pair w is decomposable. If w is strictly decomposable, the result follows from Lemma C.7. Suppose, therefore, that w is not strictly decomposable, i.e., there exists an extremal normal vector N, for which $N^\top\big(g(a)+\delta\lambda(a)-w\big)\leq 0$ for all $\delta\in\Psi_a(w,\mathcal{W})$ and every $a\in\mathcal{A}$. Because $(w,N)\in\Gamma(\mathcal{W})$, there must exist at least one action profile, for which the expression is non-negative and hence (6) holds. If $\mathcal{N}_w(\mathcal{B}(\mathcal{W}))=\{N\}$, then the maximizer a_* of (6) minimally decomposes w. If w is a corner, then, because w is decomposable, there exists $a\in\mathcal{A}$ and $\delta\in\Psi_a(w,\mathcal{W})$ with ${N'}^\top\big(g(a)+\delta\lambda(a)-w\big)\geq 0$ for arbitrarily slight rotations $N'\in\operatorname{int}\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ of the extremal normal vector N. Without loss of generality, suppose that N' are clockwise rotations of N. Since $\Psi_a(w,\mathcal{W})$ is closed, there must exist $\delta_*\in\Psi_a(w,\mathcal{W})$ with ${N'}^\top\big(g(a)+\delta_*\lambda(a)-w\big)\geq 0$ for $N'=N$ as well as for slight clockwise rotations N' of N. Therefore, a maximizes (6) and $g(a)+\delta_*\lambda(a)$ lies on the tangent to $\mathcal{B}(\mathcal{W})$ to the right of w as illustrated in Figure 24. It follows that a decomposes w. □

Finally, the following lemma shows that the boundary of $\mathcal{B}(\mathcal{W})$ is a smooth solution to the abrupt-information optimality equation (6) outside of $\mathcal{K}(\mathcal{W})\cup\mathcal{S}(\mathcal{W})$.

Lemma C.8. *Let $\mathcal{C}\subseteq\partial\mathcal{B}(\mathcal{W})$ be a segment of positive length outside of $\mathcal{K}(\mathcal{W})\cup\mathcal{S}(\mathcal{W})$ such that for each $w\in\mathcal{C}$, there exists $N\in\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ with $(w,N)\in\Gamma(\mathcal{W})$. Then $\mathcal{C}$ is a continuously differentiable solution to* (6)*.*

Proof. Lemma 4.2 implies that the boundary satisfies (6). Suppose that the boundary is not continuously differentiable outside of $\mathcal{K}(\mathcal{W})\cup\mathcal{S}(\mathcal{W})$, i.e., $\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$ is not a singleton for some w in the relative interior of $\mathcal{C}$. Lemma C.3 implies that $(w,N)\in\Gamma(\mathcal{W})$ for each $N\in\mathcal{N}_w(\mathcal{B}(\mathcal{W}))$. Since $\mathcal{A}$ is finite, there is an action profile a such that $(w,N_k)\in\Gamma_a(\mathcal{W})$ for three distinct normal vectors N_1,N_2,N_3. Lemma 4.2 implies that

$$N_k{}^\top w=\max_{\delta\in\Psi_a(w,r,\mathcal{W})}N_k{}^\top\big(g(a)+\delta\lambda(a)\big)$$

for $k=1,2,3$. This implies that w is stationary as in Figure 8; a contradiction. □

Having characterized the individual parts of the boundary in Proposition 5.5 and Lemmas C.4 and C.8, it remains to show that $\mathcal{B}(\mathcal{W})$ is the largest closed set that characterized by those individual parts. The following lemma is the essential component of the proof of Theorem 5.11.

Lemma C.9. *Fix a closed, convex set $\mathcal{W}$ with $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$. Let $\mathcal{X} \subseteq \mathcal{V}^*$ be a convex set with $\partial\mathcal{X}$ satisfying (i) and (ii) of Theorem 5.11. Then* $\operatorname{cl}\mathcal{X} \subseteq \mathcal{B}(\mathcal{W})$.

Proof. Fix a closed, convex set $\mathcal{W}$ with $\mathcal{B}(\mathcal{W}) \subseteq \mathcal{W}$ and a set $\mathcal{X}$ whose boundary satisfies (i) and (ii) of Theorem 5.11. Let $\mathcal{X}_1 = \partial\mathcal{X} \cap \mathcal{S}(\mathcal{W})$, let $\mathcal{X}_2 \subseteq \partial\mathcal{X} \setminus \mathcal{X}_1$ denote those boundary payoffs in $\mathcal{K}(\mathcal{W})$ that are strictly decomposable, let $\mathcal{X}_3 \subseteq \partial\mathcal{X} \setminus (\mathcal{X}_1 \cup \mathcal{X}_2)$ denote the remaining corners of $\mathcal{X}$ in $\mathcal{K}(\mathcal{W})$. Let $\mathcal{X}_4$ denote those boundary payoffs $w \in \partial\mathcal{X}$, where $\partial\mathcal{X}$ solves the abrupt-information optimality equation (6) in a neighborhood of w such that the drift rate on either side of w points towards w.

Lemma C.5 implies that $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_4$ is contained in $\mathcal{B}(\mathcal{W})$. At any boundary payoff $w \notin \mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3 \cup \mathcal{X}_4$, the boundary is a continuously differentiable solution to (6) or (9) satisfying the conditions of Lemma C.1. Thus, Lemma C.1 implies that there exists a $\mathcal{W}$-enforceable solution that remains on $\partial\mathcal{X}$ until a discrete event occurs or a point in $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_3 \cup \mathcal{X}_4$ is reached. If a payoff pair in $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_4 \subseteq \mathcal{B}(\mathcal{W})$ is reached, the solution can be concatenated with another $\mathcal{W}$-enforceable solution by Lemma A.2. If a payoff pair $\hat{w}$ in $\mathcal{X}_3$ is reached, decomposition of $\hat{w}$ by a singleton $\{a\}$ implies that a $\mathcal{W}$-enforceable solution $\hat{W}$ to (2) attaining $\hat{w}$ as in Lemma 4.3 enters the interior of $\mathcal{K}_a(\mathcal{W})$. Thus, the concatenation of W with $\hat{W}$ remains on $\partial\mathcal{X} \cup \mathcal{B}(\mathcal{W})$ until a discrete event occurs or another payoff pair in $\mathcal{X}_1 \cup \mathcal{X}_2 \cup \mathcal{X}_4$ is reached, etc. A countable concatenation thus yields a $\mathcal{W}$-enforceable solution that remains in $\partial\mathcal{X} \cup \mathcal{B}(\mathcal{W})$ until a discrete event occurs. This shows that $\partial\mathcal{X} \subseteq \mathcal{B}(\mathcal{W})$. Any payoff pair in the interior of $\mathcal{X}$ can be attained by using public randomization supported on $\partial\mathcal{X}$ at time 0, showing that $\operatorname{cl}\mathcal{X} \subseteq \mathcal{B}(\mathcal{W})$. □

Proof of Lemma 5.9 and Theorem 5.11. The proof of Theorem 5.11 proceeds in four steps. First, stationary payoffs are contained $\mathcal{B}(\mathcal{W})$ by Lemma C.5. Second, $\operatorname{cl}\mathcal{X} \subseteq \mathcal{B}(\mathcal{W})$ for any convex set $\mathcal{X} \subseteq \mathcal{V}^*$, whose boundary satisfies conditions (i) and (ii) by Lemma C.9. Third, Proposition 5.5 and Lemmas C.4 and C.8 show that $\partial\mathcal{B}(\mathcal{W})$ satisfies (i) and (ii) of Theorem 5.11, hence $\operatorname{cl}\mathcal{B}(\mathcal{W}) \subseteq \mathcal{B}(\mathcal{W})$ by Lemma C.9, showing that $\mathcal{B}(\mathcal{W})$ is closed. Fourth, we need to show that straight line segments on the boundary are contained in $\mathcal{V}^* \cap \operatorname{conv}\mathcal{S}(\mathcal{W})$ if Assumptions 1–3 are satisfied.

To that end, let $\mathcal{C}$ be a straight line segment in $\partial\mathcal{B}(\mathcal{W})$. Lemmas B.7 and B.10 show that for any $(w, N) \in \operatorname{int}\mathcal{V}^* \times S^1$ with w sufficiently close to $\mathcal{B}(\mathcal{W})$, there exists some action profile, for which $\Xi_a(w, N, \mathcal{W})$ is non-empty with $\kappa_a > 0$ in a neighborhood of (w, N). Consequently, if $\mathcal{C} \subseteq \operatorname{int}\mathcal{V}^*$, then for any w in the relative interior of $\mathcal{C}$, a solution to (9) with initial conditions $(w + \varepsilon N_w, N_w)$, intersects $\mathcal{B}(\mathcal{W})$ on both sides of w, hence $w + \varepsilon N_w \in \mathcal{B}(\mathcal{W})$ by Lemma 4.5, a contradiction.

Suppose next that $\mathcal{C}$ is contained in a face F of $\partial\mathcal{V}^*$ with outward normal vector N. Without loss of generality, suppose $\mathcal{C} = F \cap \mathcal{B}(\mathcal{W})$. Let w be an extremal payoff of $\mathcal{C}$. Since w is an extremal payoff of $\mathcal{B}(\mathcal{W})$, there exists a with $\Xi_a(w, N, \mathcal{W}) \neq \emptyset$. Fix $(\beta, \delta) \in \Xi_a(w, N, \mathcal{W})$ and observe that $\mathcal{W} \subseteq \mathcal{V}^*$ implies $N^\top\delta(y) \leq 0$ for any event y. Since F is either a face of $\mathcal{V}$ or $N = -e_i$ for some player i

and the straight extension of F goes through player i's minmax payoff, $N^\top(g(a)+\delta\lambda(a)-w)\geq 0$ can only hold with equality. This shows that neither is w strictly decomposable, nor can the boundary $\partial\mathcal{B}(\mathcal{W})$ be a solution to (6) or (9) in a neighborhood of w since the solution to both is a straight line segment, violating that w is extremal in $\mathcal{C}$. Consequently, w is stationary. □

Proof of Lemma 5.12. Note that $\phi(a,N,\delta)$ scales proportionally to $1/\gamma$, hence the curvature of $\partial\mathcal{B}(\mathcal{W})$ outside of $\Gamma(\mathcal{E}(r))$ goes to 0 as $\gamma\to 0$. Thus, for γ sufficiently small, no continuously differentiable solution to (8) can connect two non-adjacent vertices. □

Finally, we prove the results related to public randomization. The following lemma allows us to relate solutions to the abrupt-information optimality equation for different initial conditions.

Lemma C.10. *Let $w\in\operatorname{int}\mathcal{K}_a(\mathcal{W})$ with $(w,N)\in\Gamma_a(\mathcal{W})$. Then for any $\varepsilon>0$ with $w-\varepsilon N\in\operatorname{int}\mathcal{K}_a(\mathcal{W})$, there exists a neighborhood $\mathcal{N}$ of N such that $(w-\varepsilon N,N')\in\Gamma_a(\mathcal{W})$.*

Proof. Fix such (w,N) and observe that any $\delta\in\Psi_a(w)$ that maximizes $N^\top\delta\lambda(a)$ lies on the boundary of $\Psi_a(w)$ with outward normal $(N^1\lambda(a),N^2\lambda(a))$. The shift $w\mapsto w-\varepsilon N$ corresponds to a shift of $\left(\frac{\mathcal{W}-w}{r}\right)^{|Y|}$ by $\varepsilon(N^1,\dots,N^1,N^2,\dots,N^2)$. Since $(N^1\lambda(a),N^2\lambda(a))^\top(N^1,\dots,N^1,N^2,\dots,N^2)\geq 0$, the maximum over $N^\top\delta\lambda(a)$ weakly increases with ε. In particular,

$$\max_{\delta\in\Psi_a(w-\varepsilon N)} N^\top(g(a)+\delta\lambda(a)-w+\varepsilon) > \max_{\delta\in\Psi_a(w)} N^\top(g(a)+\delta\lambda(a)-w)\geq 0,$$

hence the inequality holds also for N' close to N. □

Proof of Proposition 5.10. Let $\mathcal{X}_1$, $\mathcal{X}_2$, X_3, and $\mathcal{X}_4$ be defined as in the proof of Lemma C.9. Suppose first that conditions (i) and (ii) are satisfied, so that for every $w_k\in\mathcal{X}_2\cup\mathcal{X}_4$, there exists a solution $\mathcal{C}_k$ to (6) with end point v_k in the interior of $\mathcal{B}(\mathcal{W})$. By using public randomization at v_k, the arguments from the proof in Lemma C.9 carry through without the use of limit strategies. For necessity, observe that infinitesimally quick public randomization (which violates predictability) is necessary to attain boundary payoffs in $\mathcal{X}_2$ where (i) fails to hold and boundary payoffs in $\mathcal{X}_4$ where (ii) fails to hold.

For the first statement, observe that $(w,N_w)\in\partial\Gamma_a(\mathcal{W})$ along any solution $\mathcal{C}$ to (10) for some action profile a. If (w,N') were in $\Gamma_a(\mathcal{W})$ for any N' in a neighborhood of N_w, then $w+\Psi_a(w,\mathcal{W})\lambda(a)$ would contain payoff pairs strictly above the tangent to $\mathcal{C}$ at w, contradicting (6). Thus, Lemma C.10 implies that solutions to (10) for maximizers a_1 and a_2 to the left and right of $w\in\mathcal{X}_4$ starting at $w-\varepsilon N_w$ must intersect as in the right panel of Figure 18. In particular, $w-\varepsilon N_w$ can be attained by a $\mathcal{W}$-enforceable solution without limit strategies for any $\varepsilon>0$ via the construction in Section 5.3. Similarly, the construction in the left panel of Figure 18 shows that any payoff in $\mathcal{K}_{\mathcal{A}_w}\setminus\{w\}$ in a neighborhood of $w\in\mathcal{X}_2$ can be attained without limit strategies.

To conclude the proof, we show that any payoff w in the interior of $\mathcal{W}$, there exists a solution to (6) or (9) that locally remains in the interior. Note that a solution to (6) starting at w cannot intersect $\partial\mathcal{B}(\mathcal{W})$ where the boundary is a solution to (6) because solutions to (6) are unique, and it cannot intersect $\partial\mathcal{B}(\mathcal{W})$ where the boundary is a solution to (9) because that would imply the

intersection point is in $\Gamma(\mathcal{W})$. A solution to (9) starting at (w, N) cannot intersect $\partial\mathcal{B}(\mathcal{W})$ for every N. Continuity of initial conditions would imply that for some N, the solution would either merge differentiably with $\partial\mathcal{B}(\mathcal{W})$, violating uniqueness in initial conditions, or it would cut through $\mathcal{B}(\mathcal{W})$, contradicting Lemma C.2. Thus any payoff pair in the interior of $\mathcal{B}(\mathcal{W})$ can be attained by a $\mathcal{W}$-enforceable solution W to (2) that remains close to the boundary until a payoff pair close to a corner of $\mathcal{B}(\mathcal{W})$ is reached. If the corner satisfies (i) of Proposition 5.5, we can concatenate W with the construction in the left panel of Figure 18. If the corner satisfies (ii) of Proposition 5.5, we can concatenate W with a $\mathcal{W}$-enforceable solution that remains on a solution to (6) that remains in the interior of $\mathcal{B}(\mathcal{W})$. Finally, if a stationary corner is reached, we concatenate W with a $\mathcal{W}$-stationary strategy profile. An iteration of this construction yields a $\mathcal{W}$-enforceable solution to (2) without the use of limit strategies that remains in $\operatorname{int}\mathcal{B}(\mathcal{W}) \cup \mathcal{S}(\mathcal{W})$ before the arrival of the first event. □

Proof of Lemma 5.13. If the conditions of Proposition 5.10 are met, any extremal payoff of $\mathcal{E}(r)$ can be attained without public randomization. If condition (i) is met, any non-extremal boundary payoff is stationary by Theorem 5.11. Moreover, any enforceable action profile can be enforced with $\delta = 0$, hence any payoff in the interior can be attained by a continuous enforceable solution to (2) until the boundary is reached. If condition (ii) is met, any non-extremal payoffs can be attained by a solution W to (2) with $A \in \mathcal{A}^N$ until an extremal payoff is reached. □