

**Continuous-Time Repeated Games with Imperfect Information:
Folk Theorems and Explicit Results**

by

Benjamin Bernard

A thesis submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

in

Mathematical Finance

Department of Mathematical and Statistical Sciences
University of Alberta

© Benjamin Bernard, 2016

Abstract

This thesis treats continuous-time models of repeated interactions with imperfect public monitoring. In such models, players do not directly observe each other's actions and instead see only the impacts of the chosen actions on the distribution of a random signal. Often, there are two reasons why this signal imperfectly reflects the chosen actions: (a) information is continuously available but it is noisy, or (b) events are observable but occur only at intermittent occasions. In a continuous-time setting, these two different types of information can be cleanly distinguished, where Brownian motion is used to model noise in the continuous information and Poisson processes indicate the arrival of informative events with an intensity that depends on players' actions.

The first major result of this thesis is a folk theorem for continuous-time repeated games even when players receive only noisy information about past play. The folk theorem gives sufficient conditions such that players achieve asymptotic efficiency as they get arbitrarily patient. Because more outcomes are sustainable in equilibrium when more information is observed, this result also applies when players receive both aforementioned types of imperfect information. In the proof, we restrict ourselves to strategies that are adjusted only at identical copies of certain stopping times. This has two important implications: (1) despite the possibility of switching actions infinitesimally fast, players do not need to do so to attain asymptotic efficiency, and (2) continuous-time equilibria can be attained as limits of equilibria in discrete-time repeated games where the length of the time period is random, rather than fixed.

The other main result of this thesis is a characterization of all payoffs that are attainable in equilibrium in such games with two finitely patient players. Relating optimal actions and incentives to the boundary of the equilibrium payoff set, we obtain a differential equation describing the curvature of the set at almost every point. The equilibrium payoff set is obtained from an iterative procedure, which is similar to that known for discrete-time repeated games but leads to an explicit characterization in our setting. Our result shows that the two types of information have drastically different impacts on the equilibrium payoff set. This is due to the fundamental difference in which the two types of information are used to provide incentives: while the continuous information can be used only to transfer value between players, the discontinuous information may be used to transfer or destroy value upon the arrival of an infrequent event. The quantitative nature of the result makes it possible to precisely measure the impact of abrupt information on the efficiency of players' payoffs in equilibrium. Thus, one can compare the value of additional information to the cost of procuring or providing it, which may lead to interesting applications in mechanism design and information disclosure.

Preface

Chapters 2 and 3 of this thesis are an extension of the article [6], which is accepted for publication in *Theoretical Economics* and written jointly by Benjamin Bernard and Christoph Frei. The proof of the main results were done by Benjamin Bernard and the idea to prove a folk theorem in this setting was Benjamin Bernard's, while the model formulation and additional results on the monotonicity of the equilibrium payoff set came from Christoph Frei. Chapter 4 constitutes original work done solely by Benjamin Bernard and not yet published.

Acknowledgements

First and foremost, I would like to thank my supervisor Prof. Christoph Frei for his continued support and encouragement. He has been an invaluable source of advice, supporting me on myriad occasions with his helpful critiques and our inspirational discussions. I am deeply grateful for his trust and to have had the opportunity to work together on such an interesting topic.

I would also like to express my gratitude to my entire dissertation committee for valuable comments and their guidance throughout my degree, especially to Prof. Tilman Klumpp for helping me find my way in the world of economics with our many insightful discussions. I would like to thank Prof. Drew Fudenberg for readily accepting to act as external examiner. I am also grateful to Profs. George J. Mailath, Semyon Malamud and Yuliy Sannikov for valuable comments. I am indebted to Prof. Byron Schmuland for supporting me as my teaching mentor and assisting me in various other endeavours.

I am grateful to the University of Alberta and especially the Department of Mathematical and Statistical Sciences for providing a great learning environment. Funding through teaching appointments at the University of Alberta and research support by the Natural Sciences and Engineering Research Council of Canada through grant RGPIN/402585-2011 are gratefully acknowledged.

Last but not least, I would like to thank all my friends throughout the years in Edmonton. I have been lucky to meet such exceptional individuals that turned my stay into a new place I call home.

Table of Contents

1	Introduction	1
1.1	Preliminaries on game theory	4
1.2	Scope of this thesis	10
2	Continuous-Time Model	14
2.1	Pure strategies	16
2.2	Example: climate agreement	23
2.3	Mixing in continuous time	27
2.4	Incentive compatibility and self-generation	35
2.5	Execution speed and the use of public randomization	47
3	Folk Theorems in Games with Continuous Information	57
3.1	Pairwise identifiability and enforceability on hyperplanes	61
3.2	Folk theorems in pure strategies	68
3.3	Folk theorems in behaviour strategies	72
3.4	Finite variation of equilibrium profiles	80
3.5	Comments on abrupt information	85
3.6	Proof of Proposition 3.1.4	86

4	Explicit Characterization of $\mathcal{E}(r)$ in Two-Player Games	96
4.1	Motivation	101
4.2	Main result	105
4.3	Interpretation of results	108
4.3.1	Climate agreement	108
4.3.2	Cournot duopoly	110
4.4	Algorithm	111
4.5	Discussion of Assumption 4.2.1	115
4.5.1	Assumption 4.2.1 and Public Randomization	117
4.5.2	Verification of Candidate	120
4.6	Computation	123
4.6.1	Climate agreement	124
4.6.2	Partnership game	125
4.7	Proof of Proposition 4.4.2	131
4.7.1	Strategic definition of $\mathcal{B}(r, \mathcal{W})$	131
4.7.2	Regularity of the optimality equation	134
4.7.3	Characterization of $\partial\mathcal{B}(r, \mathcal{W})$	142
4.7.4	Closedness of $\mathcal{B}(r, \mathcal{W})$	157
4.8	Proof of Lemma 4.5.4	163
4.A	Appendix: Auxiliary results	165
4.A.1	Lipschitz continuity of set-valued maps	165
4.A.2	Bounds on incentives and solutions to (4.9)	167
5	Conclusion	175
	Bibliography	182

List of Tables

2.1	Drift rate and instantaneous intensities in the climate agreement example.	25
2.2	Expected flow payoff in the climate agreement example.	25

List of Figures

2.1	Decomposition of information in a continuous-time game. . . .	15
2.2	Sample paths of public signal in climate agreement example under two different strategy profiles.	24
2.3	Important payoff sets in the climate agreement example. . . .	26
2.4	Restriction of incentives on the boundary of a self-generating payoff set.	43
3.1	Decomposition of payoffs in a discrete-time game.	66
3.2	For a payoff set to be self-generating, the outward drift from the tangential diffusion has to be compensated by the inward drift from the expected flow payoff.	67
3.3	The payoff set \mathcal{V}^Δ of mixed action payoffs is strictly smaller than \mathcal{V}^* in some games.	75
3.4	Explicit construction of finite-variation strategy profiles in a partnership game.	81
3.5	Comparison of our construction of strategies to the construction in Sannikov [37].	82

3.6	Abrupt information cannot be used to attain approximate efficiency.	86
3.7	Illustration of how we measure the distance of a solution to (2.8) to the boundary of a self-generating payoff set \mathcal{W}	90
4.1	Comparison of incentives in our setting to the setting in Sannikov and Skrzypacy [39]	98
4.2	Illustration that a solution to (2.8) remains on the curve \mathcal{C} defined by (4.2).	103
4.3	Procedure to enlarge $\mathcal{E}(r)$ to a larger self-generating payoff set, contradicting Proposition 2.4.3.	104
4.4	Comparison of equilibrium payoff sets in the climate agreement example for two different monitoring techniques.	109
4.5	$\mathcal{E}(r)$ in a Cournot duopoly in a homogeneous good.	110
4.6	Illustration how Assumption 4.2.1 is used in the proof of Proposition 4.4.2.	119
4.7	Computational procedure to calculate $\mathcal{E}(r)$ in the climate agreement example using the algorithm in Lemma 4.4.1.	125
4.8	Computation of $\mathcal{E}_c(r)$ when there are straight line segments on its boundary in a partnership game.	126
4.9	Computation of $\tilde{\mathcal{B}}(r, \mathcal{V}^*)$ when there are straight line segments and corners outside of \mathcal{V}^N and verification of $\tilde{\mathcal{B}}(r, \mathcal{V}^*)$	127
4.10	Comparison of $\mathcal{E}_c(r)$ to $\mathcal{E}(r)$ in partnership game.	129
4.11	The convex hull of the union of locally $(\varepsilon, r, \mathcal{W})$ -admissible payoff sets is $(\varepsilon, r, \mathcal{W})$ -admissible.	133

4.12	Construction of a curve \mathcal{C}' that cuts through $\mathcal{B}(r, \mathcal{W})$, which is impossible by Lemma 4.7.12.	150
4.13	A solution to (4.4) cannot fall into the interior of $\mathcal{B}(r, \mathcal{W})$ as otherwise one could construct a larger (r, \mathcal{W}) -admissible payoff set by changing initial conditions slightly.	152
4.14	Under suitable conditions, the intersection of Lipschitz continuous maps is Lipschitz again.	166
4.15	Construction of bound on optimal incentives, part I.	169
4.16	Construction of bound on optimal incentives, part II.	171

List of Symbols

Symbol	Description
\mathcal{A}	Set of pure action profiles
\mathcal{A}^N	Set of static Nash profiles
A	Strategy profile
a	Pure action profile
α	Mixed action profile
a^i, α^i	Action of player i in action profile a, α , respectively
a^{-i}, α^{-i}	Action profile of player i 's opponents
a_e, α_e	Static Nash profiles
$\underline{\alpha}_i$	Minmax profile against player i
$\mathcal{B}(r, \mathcal{W})$	Largest (r, \mathcal{W}) -admissible payoff set
$B_\varepsilon(\mathcal{W})$	Set of payoffs whose distance to \mathcal{W} is smaller than ε
β	Incentives through continuous component of public signal

\mathcal{C}	Curve
C^1	Continuously differentiable
\mathcal{D}^d	Space of càdlàg (right-continuous with left limits) functions from $[0, \infty)$ to \mathbb{R}^d endowed with the Skorohod topology
d	Dimension of the public signal
d_c	Number of dimensions of public signal with a non-trivial continuous component
d_z	Dimension of driving Brownian motion
δ	Incentives provided through infrequent events
$\partial\mathcal{W}$	Boundary of the set \mathcal{W}
$\mathbb{E}_Q[\cdot]$	Expectation under probability measure Q
$\mathcal{E}(r)$	Equilibrium payoff set for discount rate r
$\mathcal{E}_c(r)$	Equilibrium payoff set for discount rate r when information is restricted to continuous component of public signal
$\mathcal{E}_t(\cdot)$	Stochastic exponential
\mathbb{F}	Filtration of public information
\mathcal{F}_t	σ -algebra of public information at time t
$\mathcal{G}(r)$	Set of payoffs in $\partial\mathcal{E}(r)$, where incentives can be provided through the abrupt information exclusively

$\mathcal{G}(r, \mathcal{W})$	Set of payoffs in $\partial\mathcal{B}(r, \mathcal{W})$, where incentives can be provided through the abrupt information exclusively
$g(a)$	Expected flow payoff under action profile a
$H(w, N)$	Lower half-space defined by $\{v \in \mathbb{R}^I \mid N^\top(v - w) \leq 0\}$
$h(y)$	Impact of event of type y on the public signal
$\mathcal{I}_a^i(N, \delta^i)$	Set of all $\phi \in \mathbb{R}^{d_c}$ such that $(T^i\phi, \delta^i)$ satisfies the enforceability conditions for action profile a and player i
I	Number of players
i, j	Specific players
J^y	Poisson process indicating arrivals of events of type y
$\kappa(w)$	Curvature at point w
$\Lambda^i(\alpha)$	Matrix with column vectors $\lambda(a^i, \alpha^{-i}) - \lambda(\alpha)$ for all deviations $a^i \in \mathcal{A}^i \setminus \{\alpha^i\}$
$\lambda(y a)$	Intensity of event y under action profile a
\mathcal{M}	Payoff bound established in Sannikov and Skrzypacz [39]
M	Martingale used for public randomization
$M^i(\alpha)$	Matrix with column vectors $\mu(a^i, \alpha^{-i}) - \mu(\alpha)$ for all deviations $a^i \in \mathcal{A}^i \setminus \{\alpha^i\}$
m	Number of different types of events

$\mu(a)$	Drift rate of continuous component of public signal under action profile a
N	Unit vector, also called direction
N_w	Outward normal vector at w
$\mathcal{N}_{\mathcal{W}}(w)$	Set of outward normal vectors to $\partial\mathcal{W}$ at $w \in \partial\mathcal{W}$
$\mathcal{N}_{\mathcal{W}}$	Normal bundle to $\partial\mathcal{W}$
(Ω, \mathcal{F}, P)	Probability space
$\mathcal{O}(\cdot)$	Optional projection onto \mathbb{F}
P	Reference probability measure
$\Phi_a(N, \delta)$	Set of all $\phi \in \mathbb{R}^{d_c}$ such that $(T\phi, \delta)$ enforces a
ϕ	Row vector such that $\beta = T\phi$ are incentives for players
$\Psi_a^\varepsilon(w, N, r, \mathcal{D})$	Set of all $\delta \in \mathbb{R}^m$ with $\tilde{N}^\top \delta(y) \leq 0$ for every $\tilde{N} \in B_\varepsilon(N) \cap S^1$ such that a is (w, N, r, \mathcal{D}) -restricted-enforced by (ϕ, δ) for some $\phi \in \mathbb{R}^{d_c}$
$(Q_t^A)_{t \geq 0}$	Family of probability measures induced by strategy profile A
r	Discount rate common to all players
S^1	Unit circle
σ	Volatility of continuous component of public signal

σ_k	Arrival time of k^{th} infrequent event
T_w	Tangent hyperplane at w
τ	Stopping time
\mathcal{V}	Set of feasible payoffs
\mathcal{V}^*	Set of feasible and individually rational payoffs
\mathcal{V}^N	Set of static Nash payoffs
v, w	Payoffs
\underline{v}^i	Minmax payoff of player i
\mathcal{W}	Generic payoff set
W	Continuation value of a strategy profile
X	Public signal
$\Xi_a(w, N, r, \mathcal{W})$	Set of all pairs (ϕ, δ) that (w, N, r, \mathcal{W}) -restricted-enforce a
Y	Set of different types of infrequent (sometimes rare) events
y	A specific type of infrequent event
Z	Driving Brownian motion of continuous component of the public signal
Z^A	Brownian motion under family $(Q_t^A)_{t \geq 0}$ of probability measures

Chapter 1

Introduction

A key question in game theory is whether or not cooperation between different parties is sustainable. In many situations, a socially desirable outcome is imperfectly aligned with the parties' individual interests and hence parties need extra incentives to enforce cooperative behaviour. Consider, for example, a climate agreement where each participating country agrees to reduce its greenhouse gas emissions below a certain threshold. While the reduction of the atmospheric greenhouse gas concentration is of global benefit, the cost of reducing the greenhouse gas output is borne by the countries individually. This creates an incentive for the countries to violate the agreement and “free-ride” on the efforts of the other signatories. To deter such behaviour in the absence of a supranational court, it is necessary that the agreement is self-enforcing: If a country violates the agreement at some point in time, other countries can impose appropriate penalties in the future (e.g., they can levy import tariffs, withhold international aid, or simply respond by violating the agreement themselves). If these penalties are

sufficiently severe, they will deter countries from violating the agreement and thereby enforce it. The construction of such a self-enforcing agreement is easiest when countries can perfectly observe each other's actions. In many situations, however, observation is only *imperfect*: When a country violates its obligations and emits more than its agreed-upon share of greenhouse gases, other countries will not know immediately that this has happened. They may perhaps see an increase in industrial production in the country, or an increase in the atmospheric greenhouse gas concentration — information that is suggestive, but not conclusive proof, that a country has violated the agreement. The noisiness of the available information makes it easier for every party to get away with cheating, which substantially complicates the enforcement of the agreement. Ultimately, the question is how much coordination can be achieved when observation is imperfect.

In the language of game theorists, the above situation is a repeated game: a strategic interaction between several parties (*players*) facing the same decision repeatedly. The fact that interactions are repeated is crucial to enforce cooperative behaviour: if, in the above example, participating countries did not have to fear punishments by their opponents after violating the agreement, cooperation could not be sustained. Indeed, without repetition, the only Nash equilibrium is that of a mutual violation of the agreement by all countries. When interactions are repeated, players need to maintain a good reputation with their opponents to prevent adversary moves in the future. This intuitive way of providing intertemporal incentives by future interaction is one of the prime appeals of repeated games, together with the fact that a larger set of outcomes is consistent with equilibrium behaviour than in one-shot games.

While the theory of repeated games with imperfect information is well developed in discrete time, it has been studied only recently in a continuous-time setting. Continuous-time games exhibit the realistic feature that information may arrive both continuously and intermittently as occurrences of random events. In addition to the continuously observable information on industrial production and the atmospheric greenhouse gas concentration, countries may observe infrequent political events that are pertinent to a country's stance on climate issues. Since these events may also affect the industrial production, their impacts should be subtracted from the industrial production before using it to estimate a country's greenhouse gas emissions. In a continuous-time setting a clean separation is possible as jumps in an otherwise continuous signal are easily detected. Thus, one obtains two separate signals that are both relevant: the increase in industrial production without the impact of these political events and the frequency at which these political events occur. In a discrete-time setting, such a distinction is not unambiguously possible.

In addition to the informational advantages, continuous-time models are valuable because they often give rise to explicit results that are not available in a discrete-time setting. For example, in his introduction of continuous-time repeated games with imperfect information, Sannikov [37] shows that the continuous-time techniques make it possible to explicitly characterize the set of all payoffs that are attainable in equilibrium. Its boundary is described by a differential equation, relating optimal actions and incentives to the curvature of the set. One drawback of continuous-time games is that one often needs to make model assumptions to capitalize on mathematical tractability. In Sannikov [37], information arrives continuously only, which excludes infor-

mative but infrequent events. One of the main goals of this thesis is to extend the existing model of continuous-time repeated games with imperfect information to include the observation of discrete events in addition to the continuous information. We characterize the set of all equilibrium payoffs by a differential equation similarly to Sannikov [37] and discuss asymptotic efficiency.

1.1 Preliminaries on game theory

It was John Nash's famous one-page paper [33] that ignited the study of non-cooperative, non-zero sum games. Nash introduced a solution concept for these games that should shape game theory until the present day: in a *Nash equilibrium*, no player has an incentive to change his/her action, given that no other player does so. While this does not require that players know each others' actions in advance, it assumes that every player acts rationally and that every player's rationality is common knowledge. Thus, in any other outcome a perfectly rational player could have foreseen that his/her choice is suboptimal against perfectly rational opponents.

In non-cooperative games without repetition (one-shot games), the Nash solution concept often predicts outcomes that are not socially optimal. Because players cannot react to each other, they will not sacrifice personal utility for the social benefit as their opponents cannot retaliate to such an action. In other words, players have no incentive to maintain a reputation for being 'good'. One way to model reputation effects is by repeating the one-shot game: because players face each other again, they have to consider each other's re-

actions when making their decisions.¹ There are several ways how one can model repeated interactions. In this thesis, we will focus exclusively on discounted repeated games with an infinite time horizon. Even though this does not exactly apply when players are humans due to their finite lifespans, it is a good approximation when players cannot foresee the end of their interaction.

In a Nash equilibrium of a repeated game, it is in no player's interest to deviate from a prescribed outcome if deviations are punished sufficiently severely in the future. As is standard in the literature, we require equilibria in repeated games to be *subgame perfect* (see Selten [40]), that is, the continuation strategy after any possible history is a Nash equilibrium of the continuation game. This precludes punishment threats that are not credible, meaning that players would not execute the punishment if required. The presence of such intertemporal incentives has a drastic impact on the set of equilibrium payoffs. The celebrated folk theorem shows that any feasible and individually rational payoff can be attained in a subgame perfect equilibrium of a repeated game if players are sufficiently patient and the game satisfies a full-dimensionality condition (see Fudenberg and Maskin [17]).

In many situations, players cannot perfectly observe each other's actions and instead see only the impact of the chosen actions on a random signal. Games of this kind are called games with *imperfect monitoring*. In discrete time, games of imperfect monitoring are a strict subset of games with imperfect information, but in continuous time the two terms are used interchangeably.²

¹We omit here a whole branch of literature on reputation effects in games of incomplete information, where a long-run player of a hidden type faces a sequence of short-run players that uses Bayesian updating to deduce the type of the long-run player (see, for example, Harsanyi [20], Fudenberg, Kreps and Maskin [13], or Faingold and Sannikov [12]).

²A discrete-time game has imperfect information if there exists one information set in

If the same signal is observed by every player, it is a game of (imperfect) *public* monitoring, whereas in a game of (imperfect) *private* monitoring, players may receive classified information about past play. It is in this subfield of infinitely repeated games with imperfect public monitoring that this thesis is situated. When monitoring is only imperfect, coordination between players becomes more difficult. Nevertheless, Abreu, Pearce and Stacchetti [1, 2] provide a tractable framework for the study of these games with an iterative construction of equilibrium strategies using sequential rationality: They define a monotone operator \mathcal{B} that assigns to every payoff set \mathcal{W} the set $\mathcal{B}(\mathcal{W})$ of payoffs that can be attained with an incentive-compatible action profile in the current period and a continuation payoff in \mathcal{W} . If $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$, an iteration of this decomposition into current-period action profile and continuation payoff leads to a strategy profile that is incentive compatible after every history, and hence is a subgame perfect equilibrium. A set \mathcal{W} with the property $\mathcal{W} \subseteq \mathcal{B}(\mathcal{W})$ is called *self-generating*, and one can show that the equilibrium payoff set $\mathcal{E}(r)$ is the largest self-generating payoff set. Note that the equilibrium payoff set depends on the rate r at which players discount their future payoffs. Abreu, Pearce and Stacchetti [2] show that $\mathcal{E}(r)$ is the largest fixed-point of the operator \mathcal{B} and show that a successive application of \mathcal{B} to the set of all feasible and individually rational payoffs converges to $\mathcal{E}(r)$.

Incentive compatible continuation payoffs can be constructed in two main ways: The first way is to attach punishments to undesired outcomes similarly

its extensive form that contains multiple decision nodes (see, for example, Fudenberg and Tirole [18, p. 80]). Games with simultaneous interaction are games of imperfect information even when monitoring is perfect because players are uncertain about other players' moves in the current turn. In a continuous-time setting, this distinction vanishes since each chosen action profile is played only for an instant, which has a negligible impact.

as it is done in the proof of the folk theorem with perfect monitoring. This destruction of value (also called burning of value) upon the arrival of bad news has first been introduced by Green and Porter [19], who show that equilibrium behaviour can be sustained above the static Nash outcomes. Since these punishments are necessarily tied to the outcomes of a random signal in games of imperfect monitoring, it is inevitable that players are sometimes punished without having deviated from the desired equilibrium profile. To obtain a folk theorem, the destruction of value is too costly when monitoring is imperfect. The second way of providing incentives is by transferring future value between players so that every player is content with the current action profile after transfers. With this method, Fudenberg, Levine and Maskin [16] show that a folk theorem also holds in games with imperfect public monitoring if deviations of any two players are statistically distinguishable.

These results show that an impressive amount of cooperation can be sustained in repeated games. However, they do not give rise to an explicit description of the equilibrium payoff set $\mathcal{E}(r)$ for finitely patient players with positive discount rate. Although there exist implementations of the algorithm in Abreu, Pearce and Stacchetti [2] for games with perfect monitoring, an efficient computation for imperfect monitoring is currently unknown.³ In his introduction of continuous-time repeated games with imperfect public monitoring, Sannikov [37] shows that these problems can be remedied for a class of continuous-time two-player games, where players actions affect the drift rate of an arithmetic Brownian motion. Developing the concept of a self-generating

³In perfect monitoring, Judd, Yeltekin and Conklin [23], implement an algorithm based on Abreu, Pearce and Stacchetti [2], and Cronshaw and Luenberger [9]. An improved algorithm for two-player games in Abreu and Sannikov [3] significantly reduces the running time.

payoff set in continuous time, he relates the motion of an equilibrium profile's continuation payoff to the boundary of the equilibrium payoff set. As a result, he obtains an explicit characterization of $\mathcal{E}(r)$ by describing its boundary with an ordinary differential equation. Not only is his result of great theoretical interest, relating optimal incentives, inefficiency and noisiness of the public signal to the curvature of the equilibrium payoff set, but it also leads to extremely efficient computation times for specific examples.

The use of Brownian information in continuous-time games with imperfect monitoring has been used in several works since (see, for example, Faingold and Sannikov [12], Daley and Green [10] as well as Bernard and Frei [6]). However, Brownian information is suitable only to model information that is imperfect because observation is noisy. In addition to noisy information that is continuously observable, players may also receive sudden new information when infrequent but informative events occur such as demand shocks, accidents, equipment failure or political events in the aforementioned climate agreement example. Players have usually only limited influence on the damage caused by such an event, but their actions affect the frequency at which these events occur. The arrival of these events are thus suitably modelled by Poisson processes, whose intensities depend on players' actions.

In many situations, there is a natural decomposition of the available information into one of these two main types of information. Consider a partnership between two firms, where each firm chooses hidden effort levels and observes only the total revenue of the partnership. The total revenue is subject to both day-to-day fluctuations in supply and demand conditions, and to demand shocks when one of the firms receives bad press. In a continuous-time

model, these two different sources of information can be cleanly separated: The public signal moves continuously under normal market conditions and suddenly jumps when a demand shock occurs. This reflects the fact that bad press itself is publicly observable, whereas in discrete time, one cannot clearly distinguish between large swings of the market in regular conditions and the impact of an infrequent event. Moreover, continuous-time models rule out the possibility that two rare events occur in the same time period before players have a chance to react to it. Thus, even aside from the advantage of obtaining explicit results, continuous-time models are suitable here because they do not create an artificial information processing problem.

This is precisely why Sannikov and Skrzypacz [39] introduced this continuous monitoring structure in their analysis of discrete-time games with frequent actions. They show that the two types of information are used very differently to provide incentives because of their fundamental difference in structure. At the boundary of $\mathcal{E}(r)$, the continuous information cannot be used to burn value as it would be too costly due to infinite variation of Brownian motion. Thus, the only way to provide incentive through the continuous information is by transferring value tangentially between players as introduced in Fudenberg, Levine and Maskin [16]. The arrival of infrequent events, however, may be used to both transfer and burn value. Because actions are taken at discrete time points in Sannikov and Skrzypacz [39], they are able to apply the methods from Abreu, Pearce and Stacchetti [2] together with the aforementioned informational restrictions to obtain a bound for the equilibrium payoff set when the length of the time between two periods becomes sufficiently short. This result is similar in spirit to Fudenberg and Levine [14].

1.2 Scope of this thesis

This thesis introduces continuous-time repeated games with both continuous and abrupt information and thereby complements the existing literature with the important addition of infrequent events. In this setting with a more general information structure, our Theorem 4.2.1 proves a two-player characterization of the equilibrium payoff set via a differential equation of its boundary. Similarly to Sannikov [37], this differential equation depends on the optimal incentives at the boundary $\partial\mathcal{E}(r)$. However, the presence of infrequent events significantly complicates the analysis: Because an infrequent event causes a jump in the continuation value, self-generation of $\mathcal{E}(r)$ restricts the possible incentives through infrequent events such that the continuation payoff remains in $\mathcal{E}(r)$ even after the jump. The amount of value that can be transferred or destroyed upon the arrival of an event thus depends on $\mathcal{E}(r)$, thereby creating a fixed-point problem. We solve this problem by introducing an operator $\mathcal{B}(r, \cdot)$ similarly to the discrete-time operator \mathcal{B} from Abreu, Pearce and Stacchetti [2]: for any set \mathcal{W} , the set $\mathcal{B}(r, \mathcal{W})$ is the largest set that is self-generating up to the arrival of the first event, at which point the continuation payoff jumps to \mathcal{W} . It is easy to check that $\mathcal{W} \subseteq \mathcal{B}(r, \mathcal{W})$ implies self-generation, and we show that the continuous-time analogue to the algorithm in Abreu, Pearce and Stacchetti [2] converges to $\mathcal{E}(r)$. In contrast to its discrete-time counterpart, however, the payoff set at every step of the iteration can be characterized explicitly through a differential equation describing the curvature of its boundary similarly to Sannikov [37]. The curvature depends on the abrupt information only through the amount of value that is burnt upon the arrival of an infre-

quent event. Even though the amount of value burnt has to be nonnegative, the possible tradeoffs with incentives provided through the continuous component of the public signal increases overall efficiency: Even small amounts of value burning may lead to a significant reduction of tangential volatility that is necessary to enforce a strategy profile. Players are thus more certain when adapting to new situations, leading to an increase of efficiency.

Contrary to Sannikov [37], we do not require that action profiles are pairwise identifiable, that is, deviations of two players are not necessarily statistically distinguishable. This has several implications. Most notably, our result is also applicable when the signal is continuous but only one-dimensional. This is an important extension in itself as it contains the frequently studied partnership games, where only the total revenue is observed, and Cournot duopolies in a single homogeneous good. Because action profiles are not pairwise identifiable, players may not be willing to transfer future payoffs at any rate and instead have a cap on the exchange rate. At payoffs where these limiting rates are attained, the equilibrium payoff set may be flat, i.e., the boundary may have straight line segments. Sannikov shows that equilibrium profiles attaining extremal payoffs in $\mathcal{E}(r)$ are unique up to a set of measure zero with a continuation value that remains in $\partial\mathcal{E}(r)$. This is not the case anymore if $\partial\mathcal{E}(r)$ has straight line segments even when the public signal is entirely continuous. We present an example in Section 4.6.2, where the continuation value of equilibrium profiles must enter the interior of $\mathcal{E}(r)$, showing that in continuous-time games, a bang-bang result may fail to hold. Because of the complexity in proving an explicit characterization of the equilibrium payoff set, we present this part only towards the end of the thesis in Chapter 4.

This characterization of the equilibrium payoff set for a fixed discount rate r is complemented by an analysis as r goes to 0: We show in Chapter 3 that a folk theorem holds for continuously repeated n -player games even if information is Brownian only as in Sannikov [37]. The conditions on the folk theorem are similar to the conditions of its discrete-time counterpart in Fudenberg, Levine and Maskin [16], essentially requiring that deviations of any two players can be statistically distinguished. As such, the result itself may not be surprising, but its proof shows the intricacies that need to be dealt with to provide sufficient incentives in a continuous-time setting.⁴ Conditions need to be strengthened slightly to ensure that small changes in payoffs lead to only small changes in the volatility of the induced continuation value. As a consequence, it is possible to prove the folk theorem by using locally constant strategy profiles exclusively. This is very desirable from an implementation standpoint as these strategies switch actions only finitely many times on finite time intervals. In comparison, taking limits of strategies obtained with discrete-time techniques typically lead to strategies of unbounded oscillation.

The remainder of this thesis is structured as follows: Chapter 2 introduces the continuous-time model including both continuous and abrupt information and proves first results on incentive compatibility, self-generation, execution speed of strategy profiles and the use of public randomization. Chapter 3 establishes several versions of the folk theorem and Chapter 4 derives an exact description of $\mathcal{E}(r)$ in the two-player case.

⁴This stands in contrast to equilibrium behaviour between one long-run player with a hidden type and a sequence of short-run players. Faingold and Sannikov [12] and Fudenberg and Levine [15] show that reputation effects vanish in the limit as the time period goes to zero in games where players' actions affect the drift rate of a Brownian signal. In these games, only the static Nash payoff is attainable in equilibrium.

Throughout the thesis, I will compare features of continuous-time games to their discrete-time counterparts without properly introducing discrete-time repeated games as this would inflate the thesis. For a detailed overview of discrete-time games and reputations, see Mailath and Samuelson [31]. This thesis is intended for an audience of both mathematicians and economists, and hence will inevitably contain information in more detail than needed for either one of these groups of researchers. Nevertheless, I think this is the best approach to present a topic that contains advanced methods of both fields and to make it accessible to a broader audience.

Chapter 2

Continuous-Time Model

Players in continuous-time games often have the same actions available at any point in time and instantaneous payoffs depend on the past only through the choice of the current action profile. Such games, which are stationary in time, are called *continuous-time repeated games*. Frequently, they have a public signal that can be split into a continuous and a discontinuous component. Consider, for example, a joint venture of two partners that secretly choose effort levels. Instead of observing each other's effort levels, the partners observe only the total revenue of the joint venture, which represents the public signal of the game. Demand is stochastic so that the revenue only imperfectly reflects players' strategies. The total revenue has a continuous component that is due to the instantaneous profits of production and sale, and it has a discontinuous component that reflects demand shocks because of bad press or large expenses for equipment upkeep, settlements of lost law suits, etc.; see Figure 2.1. In such an application it is natural to distinguish between continuous and discontinuous information.

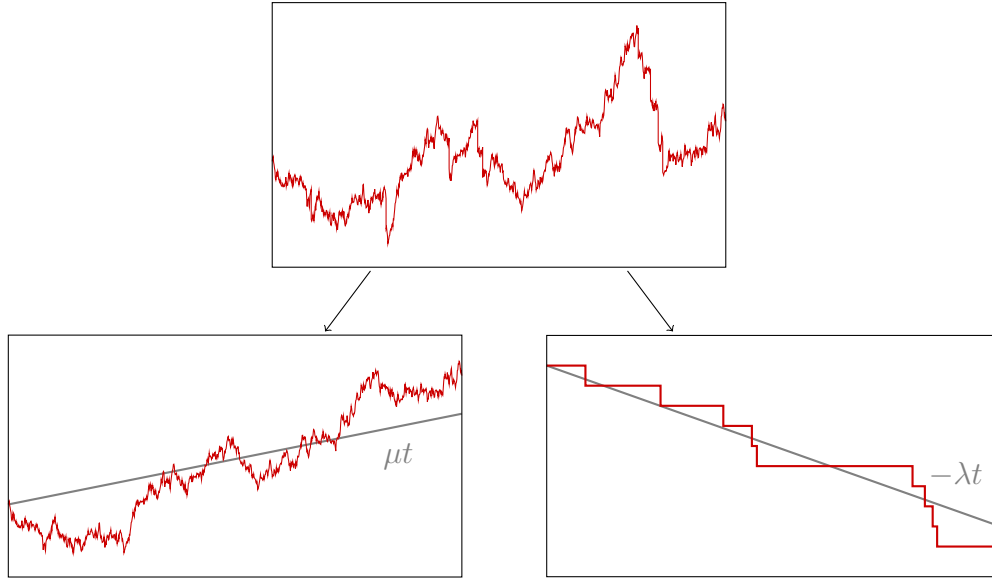


Figure 2.1: The total revenue of the partnership (top panel) consists of normal market fluctuations (left panel) and shocks due to infrequent events (right panel). Both components carry important information about past play. The expected rise in continuous revenue μ and the intensity of the events λ depend on the chosen effort levels.

Mathematically, the public signal in a continuous-time repeated game is a stochastic process. Because of stationarity in time, conditional on players' strategies, increments of the public signal are independent of the past and identically distributed. That is, the public signal is a Lévy process. Instead of considering all Lévy processes, we restrict ourselves to the sum of arithmetic Brownian motion and Poisson processes in this thesis. According to the Lévy decomposition theorem (see, for example, Theorem 15.4 in Kallenberg [24]), this includes all continuous Lévy processes but not all discontinuous Lévy processes. The reason why we restrict ourselves to this class of public signals is illustrated in the aforementioned example: the players' strategies have an impact mainly on the frequency of the randomly occurring events so that we model each type of event with a Poisson process, whose intensity depends on the chosen

strategies. Information thus arrives in one of two ways: it may be continuously observable but carry only noisy information about past play, or it may arrive infrequently with an intensity that depends on players' strategies.

We begin this chapter by describing pure strategies and how they affect the distribution of the public signal in Section 2.1. We also introduce the main solution concept, that of a perfect public equilibrium (PPE). To get a flair for these games, we present in Section 2.2 an example, for which the framework of repeated games is very natural. In Section 2.3, we consider extensions to mixed and behaviour strategies and show that the two notions are realization equivalent. This is the analogue to Kuhn's theorem in continuous time. Section 2.4 contains the important concepts of incentive compatibility and self-generation, which are essential to the construction of equilibria. Finally, we show in Section 2.5 how game primitives can be exchanged for each other without affecting equilibrium outcomes. This includes a time-change result, showing that the slower execution of a PPE corresponds to a PPE at a smaller discount rate if public randomization is used suitably.

This chapter extends the models in Bernard and Frei [6] to include discontinuous information, which is reflected in the results in Sections 2.4 and 2.5. In its general outline, the chapter follows [6] fairly closely.

2.1 Pure strategies

In a continuous-time repeated game, I players $i = 1, \dots, I$ continuously take actions from the finite sets \mathcal{A}^i at each moment of time $t \in [0, \infty)$. An action in \mathcal{A}^i is also called a *pure action* to indicate that players abstain from any form of

randomization by choosing an action in \mathcal{A}^i . The set of all *pure action profiles* $a = (a^1, \dots, a^I)$ is denoted by $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^I$. In a game with imperfect monitoring, chosen actions affect the distribution of a public signal X , which in turn affects the actions that players choose. To avoid a recursion in the definition of players' strategies, the public signal is defined as a stochastic process with a fixed distribution under a preliminary probability measure P . This stochastic process generates the filtration of public information, based on which players choose their actions. Finally, the impact of a strategy profile on the distribution of the public signal is modelled by changing the probability measure, under which players observe the game.

To formalize this construction, let (Ω, \mathcal{F}, P) be a probability space containing a d_z -dimensional Brownian motion Z and Poisson processes $(J^y)_{y \in Y}$, where y is in some finite index set $Y = \{y_1, \dots, y_m\}$. We assume that $(J^y)_{y \in Y}$ are independent of each other and independent of Z . Under this preliminary probability measure P , the Poisson processes all have intensity 1. The Brownian motion represents the noise in the continuous component of the public signal $X^c = \sigma Z$, where the volatility matrix $\sigma \in \mathbb{R}^{d_c \times d_z}$ is of rank d_c . An informative event of type y arrives according to the jumps of Poisson process J^y and has an impact of $h : Y \rightarrow \mathbb{R}^d$ on the public signal. The public signal is thus d -dimensional, of which the first $d_c \leq d$ dimensions have non-trivial continuous components. The arrival of public information is captured by the *public filtration* $\mathbb{F} = (\mathcal{F}_t)_{t \geq 0}$, where \mathcal{F}_t contains the history of the processes Z and $(J^y)_{y \in Y}$ up to time t . The filtration may be strictly larger than the filtration generated by Z and $(J^y)_{y \in Y}$, containing additional information that players may use as a public randomization device.

Definition 2.1.1. A *(public) pure strategy* of player i is an \mathbb{F} -predictable process $A^i : \Omega \times [0, \infty) \rightarrow \mathcal{A}^i$ taking values in \mathcal{A}^i . The tuple $A = (A^1, \dots, A^I)$ is called a *pure strategy profile*.

Remark 2.1.1.

1. We will omit the term public when referring to strategies since this entire dissertation treats games of public information exclusively. In a game of public information, it is possible to identify the probability space with the path space of the public signal. A strategy A^i is thus a decision rule that maps a realized path of the public signal $\omega \in \Omega$ to an action $A_t^i(\omega)$ at any point in time t . Observe that the outcome $A(\omega)$ is unique, which is in contrast to continuous-time games with perfect information as shown by Anderson [4] as well as Simon and Stinchcombe [41].
2. A stochastic process A is *predictable* with respect to a filtration \mathbb{F} if it is measurable with respect to the σ -algebra generated by all left-continuous and \mathbb{F} -adapted processes; see, for example, Billingsley [7]. This is the suitable concept of measurability because it means that player i 's decision $A_t^i(\omega)$ at time t is based only on the public information before time t , and it has the right integrability properties for discontinuous integrators. If the public signal is entirely Brownian as in Sannikov [37] or Bernard and Frei [6], then the slightly weaker notion of progressive measurability may be used instead.

The game primitives $\mu : \mathcal{A} \rightarrow \mathbb{R}^{d_c}$ and $\lambda(y|\cdot) : \mathcal{A} \rightarrow (0, \infty)$ for $y \in Y$ determine the impact of a chosen pure action profile on the drift rate of the public signal and the intensity of events of type y , respectively. Denote

by $\lambda(a) = (\lambda(y_1|a), \dots, \lambda(y_m|a))^\top$ the vector containing the intensities of all events. We assume that events of any type y are possible after any history, that is, it is a game of full support public monitoring.

Assumption 2.1.1. $\lambda(y|a) > 0$ for every $a \in \mathcal{A}$ and every $y \in Y$.

The choice of a strategy profile affects the future distribution of the public signal through a change of probability measure. This is the natural analogue to taking expectations under conditional distributions in discrete time; see, for example, Section 7.1.1 of Mailath and Samuelson [31]. Strategy profile A induces a family $Q^A = (Q_t^A)_{t \geq 0}$ of probability measures, under which players observe the game when A is played.⁵ Under Q^A , the signal is decomposed into

$$X = \int \mu(A_s) ds + \sigma Z^A + \sum_{y \in Y} h(y) J^y,$$

where we use the convention that the d_c -dimensional continuous component is added to the first d_c components of the jump component. It follows from Girsanov's theorem that for any $t > 0$, under Q_t^A on the interval $[0, t]$, the process $Z^A := Z - \int \sigma^\top (\sigma \sigma^\top)^{-1} \mu(A_s) ds$ is a Brownian motion and that J^y is a counting process with instantaneous intensity $\lambda(y|A)$ for every $y \in Y$. Thus, from the players' perspective, the continuous component of the public signal X^c consists of a drift term determined by $\mu(A)$ and Brownian noise.

⁵The family $Q^A = (Q_t^A)_{t \geq 0}$ is defined via its density process

$$\frac{dQ_t^A}{dP} := \mathcal{E}_t \left(\int_0^\cdot \mu(A_s)^\top (\sigma \sigma^\top)^{-1} \sigma dZ_s + \sum_{y \in Y} \int_0^\cdot (\lambda(y|A_s) - 1) (dJ_s^y - ds) \right), \quad (2.1)$$

where $\mathcal{E}_t(X) := \exp(X_t - X_0 - \frac{1}{2} \langle X^c \rangle_t) \prod_{0 < s \leq t} (1 + \Delta X_s) e^{-\Delta X_s}$ is the stochastic exponential.

Remark 2.1.2. By observing the continuous component of the public signal X^c , players also observe its quadratic variation $\langle X^c \rangle_t = \sigma \sigma^\top t$. In a continuous-time setting, the instantaneous volatility can be perfectly estimated by players, which is why σ is constant in games of imperfect public monitoring. This is different from discrete time, where changes in volatility due to players' actions lead to very efficient ways of providing incentives to players; see Fudenberg and Levine [15]. Because the volatility is constant, one can normalize it to a diagonal matrix with ones on the first d_c components and zeroes thereafter. Indeed, σ has right-inverse $\sigma^\top (\sigma \sigma^\top)^{-1}$ and the game is equivalent to the game with drift rate $\tilde{\mu} = \sigma^\top (\sigma \sigma^\top)^{-1} \mu$ and jump sizes $\tilde{h} = \sigma^\top (\sigma \sigma^\top)^{-1} h$. We will thus consider normalized games in Chapters 3 and 4 for ease of notation, noting that it comes at no cost of generality.

Suppose that players have an affine payoff structure, that is, every player i earns a *discounted future payoff* according to

$$V_t^i(A) := \int_t^\infty r e^{-r(s-t)} (b^i(A_s^i) dX_s - c^i(A_s^i) ds),$$

where r is a discount rate common to all players, $b^i : \mathcal{A}^i \rightarrow \mathbb{R}^d$ is the sensitivity of player i 's payoff to the public signal and $c^i : \mathcal{A}^i \rightarrow \mathbb{R}$ is a cost-of-effort function. Observe that a player's payoff depends on the strategies of his/her opponents only through the distribution of the public signal. Players are assumed to be risk-neutral and aim to maximize the conditional expectation based on the information in \mathcal{F}_t ,

$$\lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^A} \left[\int_t^u r e^{-r(s-t)} (b^i(A_s^i) dX_s - c^i(A_s^i) ds) \middle| \mathcal{F}_t \right]. \quad (2.2)$$

Because A^i is valued in a finite set, the process $b^i(A^i)$ is uniformly bounded. It follows that for any $u > t$, the increments $\int_t^u e^{-r(s-t)} b^i(A_s^i) (\sigma dZ_s - \mu(A_s) ds)$ and $\int_t^u r e^{-r(s-t)} b^i(A_s^i) h(y) (dJ_s^y - \lambda(y|A_s) ds)$ for every $y \in Y$ are martingale increments under Q_u^A . Therefore, (2.2) equals

$$\int_t^\infty r e^{-r(s-t)} \mathbb{E}_{Q_s^A} \left[b^i(A_s^i) \left(\mu(A_s) + \sum_{y \in Y} h(y) \lambda(y|A_s) \right) - c^i(A_s^i) \middle| \mathcal{F}_t \right] ds.$$

Observe that $b^i(A_s^i) (\mu(A_s) + \sum_{y \in Y} h(y) \lambda(y|A_s)) - c^i(A_s^i)$ is the instantaneous payoff rate that player i expects to receive under strategy profile A . It is the equivalent to the ex-ante stage game payoff in discrete-time repeated games. Because players are risk-neutral, ex-ante payoffs are sufficient to analyze the behaviour of the players. We will thus allow payoffs to be defined directly through an *expected flow payoff* function $g^i : \mathcal{A} \rightarrow \mathbb{R}$ of any functional form that depends on the action profile a^{-i} of player i 's opponents only through the distribution of the public signal. That is, the expected flow payoff is of the form $g^i(a) = f^i(a^i, \mu(a), \lambda(a))$ for some function f^i .

Definition 2.1.2. Player i 's *discounted expected future payoff*, also known as player i 's *continuation value*, under a strategy profile A at time t equals

$$W_t^i(A; r) := \int_t^\infty r e^{-r(s-t)} \mathbb{E}_{Q_s^A} [g^i(A_s) | \mathcal{F}_t] ds, \quad (2.3)$$

We often omit the discount rate if there is no chance for confusion. Note that player i 's continuation value depends on players' strategies directly through the expected flow payoff $g^i(A)$ and indirectly through the change of measure. Because the weights $r e^{-r(s-t)}$ integrate up to one, the continuation value is a

convex combination of stage-game payoffs. The set of *feasible payoffs* is thus given by the convex hull of pure action payoffs $\mathcal{V} = \text{conv} \{g(a) \mid a \in \mathcal{A}\}$.

Definition 2.1.3. A strategy profile A is a *perfect public equilibrium (PPE)* for discount rate r if for every player $i = 1, \dots, I$ and every $t \geq 0$,

$$W_t^i(A; r) \geq W_t^i(\tilde{A}; r) \text{ a.s.}^6 \quad (2.4)$$

for all public strategy profiles \tilde{A} , for which the strategy profile of player i 's opponents is almost everywhere the same, i.e., $\tilde{A}^{-i} = A^{-i}$ a.e.⁷ This means that no player has a strictly profitable unilateral deviation at any point in time and hence, in the absence of cooperation, no player has an incentive to deviate from the equilibrium profile. Denote the set of payoffs achievable by perfect public equilibria by

$$\mathcal{E}(r) := \{x \in \mathcal{V} \mid \text{there exists a PPE } A \text{ with } W_0(A; r) = x \text{ a.s.}\}.$$

Observe that $\mathcal{E}(r)$ is convex because we allow for public randomization in \mathcal{F}_0 . Indeed, let $w = \kappa w' + (1 - \kappa)w''$ be a convex combination of two payoffs $w', w'' \in \mathcal{E}(r)$, achieved by PPE A' and A'' , respectively. Then w is attained by randomly selecting A' with probability κ and A'' with probability $1 - \kappa$, i.e., it is attained by $A = A'1_{\Xi} + A''1_{\Xi^c}$ for $\Xi \in \mathcal{F}_0$ with $P(\Xi) = \kappa$.

⁶We omit with respect to which probability measure the statement holds almost surely (a.s.) because Q_t^A is equivalent to P for any $t > 0$. Indeed, Assumption 2.1.1 ensures that the jumps $(\lambda(y|A_s) - 1)\Delta J_s^y$ are bounded away from -1 , hence the density process in (2.1) remains strictly positive throughout.

⁷Because players maximize their discounted expected future payoff, deviations with time measure zero or probability zero are irrelevant. Therefore, two strategy profiles lead to the same continuation value if they are $P \otimes \text{Lebesgue}$ -almost everywhere (a.e.) the same.

Because player i can always deviate to the strategy of *myopic best responses* $\arg \max g^i(\cdot, A^{-i})$, outcomes below his/her *pure action minmax payoff*

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

are precluded in equilibrium. Any payoff w with $w^i \geq \underline{v}^i$ for all $i = 1, \dots, I$ is called *individually rational* and the set of all feasible and individually rational payoffs is denoted by $\mathcal{V}^* := \{w \in \mathcal{V} \mid w^i \geq \underline{v}^i \text{ for all } i\}$. Particularly simple PPE of the repeated game are strategy profiles, where every player plays a myopic best response to their opponents simultaneously. That is, at every point in time t , A_t is a Nash profile of the stage game (g, \mathcal{A}) . Let $\mathcal{A}^N \subseteq \mathcal{A}$ denote the set of stage game Nash equilibria and denote by $\mathcal{V}^N := \text{conv} \{g(a) \mid a \in \mathcal{A}^N\}$ the set of static Nash payoffs. It follows that $\mathcal{V}^N \subseteq \mathcal{E}(r) \subseteq \mathcal{V}^* \subseteq \mathcal{V}$. The study of how precisely $\mathcal{E}(r)$ lies between \mathcal{V}^N and \mathcal{V}^* is an important part of game theory, and the topic of this thesis.

2.2 Example: climate agreement

Consider a climate agreement between two neighbouring countries that obligates each signatory to reduce greenhouse gas (GHG) emissions. In the absence of a supranational court, such an agreement must be self-enforcing, which requires repeated interaction: If a country violates the agreement at some point in time, the other country can impose appropriate penalties in the future. If these penalties are sufficiently severe they will deter countries from violating the agreement and thereby enforce it. Since countries cannot measure the




Figure 2.2: Depicted are sample paths of the industrial production index X^i in country i and the number of bad news J^{y^i} that have occurred under $A^i \equiv C$ and $A^i \equiv D$ in the left and right panel, respectively. Given only a short period of observation, it is difficult to correctly estimate the underlying action profile.

total GHG output of their counterparty, they will not know with certainty whether their counterparty has violated the terms of the agreement. They may, however, see an increase in industrial production in the country, or an increase in atmospheric GHG concentration — information that is suggestive, but not conclusive proof of such a violation. In addition to the observation of these continuous processes, countries also observe infrequent but informative political and economic events, such as the passing of an environmental bill, the commissioning of a coal power plant and similar events. While these events individually may not significantly affect a country’s total GHG emission, they may be a good indication of a country’s overall policy.

To formalize this setting, suppose that a country can either cooperate (C) or defect on the agreement (D). The public signal $X = (X^1, X^2, G)^\top$ has three continuous components. The first two are given by industrial production indices of the respective countries and the third component is the atmospheric greenhouse gas concentration. Suppose that the expected annual increase of country i ’s industrial production is 1.6% or 2% under policies C and D ,

$\mu(a)$	C	D	$\lambda(y a^i)$	C	D
C	(1.6, 1.6, 1.8)	(1.6, 2, 5.7)	\bar{y}^i	1/3	1/4
D	(2, 1.6, 1.9)	(2, 2, 2)	\underline{y}^i	1/10	1/6

Table 2.1: Drift rate of the continuous component of the public signal X^c in the left table and intensities of events in the right table.

$g(a)$	C	D
C	(1.3, 1.3)	(1, 1.4)
D	(1.4, 1)	(1.1, 1.1)

Table 2.2: Expected flow payoff of pure action profiles.

respectively, so that X^i satisfy

$$\frac{dX_t^i}{X_t^i} = (1.6 + 0.4 \cdot 1_{\{A_t^i=D\}}) dt + dZ_t^{A,i}, \quad i = 1, 2,$$

where $Z^{A,1}$, $Z^{A,2}$ are Brownian motions under Q^A with correlation coefficient 0.6. Let G denote the atmospheric GHG concentration with expected annual increase of 1.8% under compliance with the climate agreement. A violation by one country amplifies the increase by 0.1%. Suppose further that the atmospheric GHG concentration is less volatile than the industrial production, so that its law of motion under strategy profile $A = (A^1, A^2)$ is given by

$$\frac{dG_t}{G_t} = (1.8 + 0.1 \cdot 1_{\{A_t^1=D\}} + 0.1 \cdot 1_{\{A_t^2=D\}}) dt + \frac{1}{3} dZ_t^{A,3}$$

for a standard Brownian motion $Z^{A,3}$ under Q^A independent of $Z^{A,1}$, $Z^{A,2}$. In addition to the continuous processes, the countries observe infrequent but informative political and economics events about a country's policy. Suppose that for $i = 1, 2$, an indicator \bar{y}^i of a climate-friendly policy in country i occurs




Figure 2.3: The set of feasible payoffs (dashed lines) is the convex hull of stage-game payoffs. Any individually rational payoffs dominates the pure action minmax payoff pair $g(D, D) = (1.1, 1.1)$, which coincides with the static Nash payoff \mathcal{V}^N .

on average every three years under C and every four years under D . An event y^i that suggests an environmentally unfriendly policy in country i happens on average every six years under D and every ten years under C . See Table 2.1 for a summary of the drift function μ and the intensities λ in this game.

Suppose that a country's payoff is linear in its industrial production and that each country suffers an externality due to environmental damages. Specifically, the discounted future increase in country i 's welfare is given by

$$\int_t^\infty r e^{-r(s-t)} \left(\frac{dX_t^i}{X_t^i} - 3 \left(\frac{dG_t}{G_t} - 1.7 dt \right) \right).$$

This corresponds to an expected flow payoff $g^i(a) = \mu^i(a) - 3(\mu^G(a) - 1.7)$. Observe that the political events do not directly affect players' payoffs and are instead of purely informational nature. Figure 2.3 shows the payoff sets $\mathcal{V}^N \subseteq \mathcal{V}^* \subseteq \mathcal{V}$ of this game. The unique static Nash equilibrium corresponds to the situation where both countries disregard the terms of the agreement or,

equivalently, do not sign the agreement in the first place. This is a game of moral hazard, where countries have an incentive to free-ride on each other's efforts because both countries benefit equally from a low greenhouse gas concentration in the atmosphere. Because both parties prefer (C, C) to (D, D) , the continuation value of country i is closely related to the beliefs of its counterparty that country i is honouring the agreement.

2.3 Mixing in continuous time

A mixed action α^i of player i is a randomization over his/her available pure actions. It is an element of $\Delta(\mathcal{A}^i)$, the space of all distributions over \mathcal{A}^i . In a mixed action profile $\alpha = (\alpha^1, \dots, \alpha^I)$, outcomes of players' mixed actions are drawn independently from each other. Because players are risk-neutral, stage game payoffs are extended to mixed actions by multilinearity

$$g(\alpha) = \sum_{a \in \mathcal{A}} g(a) \alpha^1(a^1) \cdots \alpha^I(a^I), \quad (2.5)$$

where $\alpha^i(a^i)$ denotes the probability that player i assigns to pure action a^i . Denote by $\Delta(\mathcal{A}) := \Delta(\mathcal{A}^1) \times \cdots \times \Delta(\mathcal{A}^I)$ the set of all mixed action profiles.

Considering mixed actions has the major advantage that action spaces are convex. This leads to several nice properties, including the fact that any stage game has at least one Nash equilibrium as famously demonstrated by Nash [33] in 1950. In addition, we obtain approximation properties that lower the conditions on the stage game necessary for a folk theorem to hold; see also Fudenberg, Levine and Maskin [16]. When players have access to mixed actions, more payoffs are possible in equilibrium as the payoff bound is lowered

to the *mixed action minmax payoff*

$$\underline{v}^i = \min_{\alpha^{-i} \in \Delta(A^{-i})} \max_{a^i \in \mathcal{A}^i} g^i(a^i, \alpha^{-i}). \quad (2.6)$$

Observe that it is sufficient to consider player i 's deviations to pure actions since a mixed action's payoff is a convex combination of payoffs of the pure actions in its support. We denote by α_i^{-i} the *minmax profile* against player i , which is the action profile minimizing (2.6).

Definition 2.3.1. A (*public*) *behaviour strategy profile* is an \mathbb{F} -predictable process $A : \Omega \times [0, \infty) \rightarrow \Delta(\mathcal{A})$. Its discounted expected future payoff is defined as in (2.3), using multilinearity in (2.5).

The outcome of a behaviour strategy is unique up to realizations of the mixed actions. In a continuous-time setting, these realizations have to be drawn continuously: Suppose that player i plays a fixed mixed action over an interval $[s, t]$ and samples from his/her mixing distribution only at discrete intervals. An opponent who samples more frequently may realize this after a couple of his/her own samples, and henceforth play a best response to the already sampled action of player i . To avoid such a scenario, sampling has to be done continuously, where the realizations are drawn from a continuum of independent events. Because each realization is played only for an instant, different realizations should not affect the distribution of the public signal. This means that μ and λ are extended to mixed action profiles by multilinearity, which essentially assumes an exact law of large numbers (see Judd [22]).

The outcome of a mixed action profile α corresponds to the realization of an \mathcal{A} -valued random variable γ on some probability space, whose distribution

is given by α . We call such a random variable an *instantiation* of α . Sampling of a behaviour strategy profile could thus possibly involve a continuum of probability spaces, on which A_t are instantiated. We proceed to show that there exists a unified probability space Ω , rich enough to contain the public information and the outcomes of players' behaviour strategy profiles.

Definition 2.3.2. We call $(\gamma_t)_{t \geq 0}$ an *instantiation* of a behaviour strategy profile A if at each point in time t , the \mathcal{F}_t -conditional distribution of γ_t equals A_t . Then $\gamma(\omega)$ is the outcome of the strategy profile A .

An instantiation of a behaviour strategy profile A is constructed by setting

$$\gamma_t = \left(\sum_{a^1 \in \mathcal{A}^1} a^1 1_{\Xi_t^1(a^1)}, \dots, \sum_{a^I \in \mathcal{A}^I} a^I 1_{\Xi_t^I(a^I)} \right), \quad t > 0, \quad (2.7)$$

where $\Xi_t^i = (\Xi_t^i(a^i))_{a^i \in \mathcal{A}^i}$ is a partition of Ω independent of the public filtration \mathbb{F} , of all past partitions Ξ_s^i for $s < t$, and of all partitions Ξ_t^j of player i 's opponents $j \neq i$ such that $A_t^i(a^i) = P(\Xi_t^i(a^i) | \mathcal{F}_t)$. The following lemma justifies that such a construction is indeed possible.

Lemma 2.3.1. *There exist independent filtrations $\mathbb{M}^i = (\mathcal{M}_t^i)_{t \geq 0}$, $i = 1, \dots, I$ that are independent of the public filtration, such that at all $t > 0$, \mathcal{M}_t^i contains finite partitions of Ω of arbitrary size that are independent of \mathcal{M}_s^i for all $s < t$.*

Proof. Fix a player i . We start by constructing a process $(U_t^i)_{t \geq 0}$ such that each U_t^i is standard uniformly distributed and independent of U_s^i for $s < t$. Indeed, its finite-dimensional distributions satisfy

$$P(U_{t_1}^i \leq c_1, \dots, U_{t_n}^i \leq c_n) = \prod_{j=1}^n P(U_{t_j}^i \leq c_j) = \prod_{j=1}^n c_j$$

for all $t_j \in [0, \infty)$ such that $t_j \neq t_\ell$ for $j \neq \ell$, all $c_j \in [0, 1]$ and all $n \in \mathbb{N}$. Since this family of finite-dimensional distributions is consistent, Kolmogorov's existence theorem (see, for example, Theorem 36.2 of Billingsley [7]) tells us that such a process indeed exists. Independent partitions of the appropriate size can now be found as the preimage of a partition of $[0, 1]$ under U^i . Therefore, the filtration generated by U^i will serve as \mathbb{M}^i . \square

\mathbb{M}^i is the personal source of randomness that player i has available for mixing. Because these filtrations are independent, neither do players learn anything about the signal from their personal source of randomness, nor can they predict the outcome of their opponents' mixing. For every $i = 1, \dots, I$, let \mathbb{F}^i denote the augmented filtration generated by \mathbb{F} and \mathbb{M}^i .

Lemma 2.3.2. *A stochastic process γ is the instantiation of a behaviour strategy profile if and only if for every player i , the process $\gamma^i : \Omega \times [0, \infty) \rightarrow \mathcal{A}^i$ is adapted to \mathbb{F}^i and the optional projection $\mathcal{O}(\gamma^i)$ of γ^i onto the public filtration \mathbb{F} exists and is \mathbb{F} -predictable.⁸*

Remark 2.3.1. The realizations of player i 's mixed actions may, but do not have to be predictable with respect to the filtration \mathbb{M}^i . However, player i 's decision on how to mix amongst his/her pure actions has to be based on the public filtration \mathbb{F} in a predictable way. This is captured by the requirement that the optional projection $\mathcal{O}(\gamma^i)$ of γ^i onto the public filtration \mathbb{F} is predictable for

⁸The \mathbb{F} -optional projection of γ^i is defined as the unique \mathbb{F} -optional process X such that

$$\mathbb{E}[\gamma_\tau^i 1_{\{\tau < \infty\}} | \mathcal{F}_\tau] = X_\tau 1_{\{\tau < \infty\}} \text{ a.s.}$$

for every \mathbb{F} -stopping time τ ; see Section VI.2 of Dellacherie and Meyer [11] for further details on the optional projection and the related predictable projection.

every player i . Note that this condition is satisfied, for example, by a process γ such that γ^i is \mathbb{F}^i -predictable for every player i .⁹

Proof of Lemma 2.3.2. Fix a behaviour strategy profile A and define an instantiation γ by (2.7), where the partitions Ξ_t^i with $A_t^i(a^i) = P(\Xi_t^i(a^i) | \mathcal{F}_t)$ exist by Lemma 2.3.1. Since \mathbb{M}^i are defined as the filtrations generated by Ξ^i and A^i is \mathbb{F} -predictable, γ^i has the necessary measurability properties. For the converse, let γ be a stochastic process such that γ^i is \mathcal{A}^i -valued, \mathbb{F}^i -adapted and $\mathcal{O}(\gamma^i)$ is \mathbb{F} -predictable. For any $a \in \mathcal{A}$, define $A(a) := \mathcal{O}(1_{\{\gamma=a\}})$. Observe that $A(a)$ is \mathbb{F} -predictable by assumption and that $\sum_{a \in \mathcal{A}} A(a) = 1$ a.e. Because $\mathbb{M}^1, \dots, \mathbb{M}^I$ are independent of each other, it follows that

$$A_t(a) = \mathbb{E}[1_{\{\gamma_t^1=a^1\}} \cdots 1_{\{\gamma_t^I=a^I\}} | \mathcal{F}_t] = P(\gamma_t^1 = a^1 | \mathcal{F}_t) \cdots P(\gamma_t^I = a^I | \mathcal{F}_t) \text{ a.s.}$$

This means that the players' distributions are conditionally independent, given the public information. Therefore, A is indeed a behaviour strategy profile. \square

Observe that $\mu(A) = \mathcal{O}(\mu(\gamma))$ for any instantiation γ of A . Because $\mu(A)$ is \mathbb{F} -predictable, so is the density process dQ^A/dP defined in (2.1). Therefore, we immediately obtain the following consistency result.

Lemma 2.3.3. *For any behaviour strategy profile A ,*

1. Q^A agrees with P on $\mathbb{M}^1, \dots, \mathbb{M}^I$.
2. $\mathbb{F}, \mathbb{M}^1, \dots, \mathbb{M}^I$ are independent under Q^A .
3. The \mathbb{F} -optional projections under Q^A and P coincide.

⁹Indeed, the \mathbb{F} -optional and \mathbb{F} -predictable projections are constructed using càdlàg versions of the processes $\mathbb{E}[\gamma_T^i | \mathcal{F}_t]$ and $\mathbb{E}[\gamma_T^i | \mathcal{F}_{t-}]$, respectively, on $[0, T]$ for $T > 0$ (see the proof of Theorem 43 in Dellacherie and Meyer [11]). Independence of \mathcal{M}_t^i and \mathcal{F}_t implies that $\mathbb{E}[\mathbb{E}[\gamma_T^i | \mathcal{F}_{t-}] | \mathcal{F}_t] = \mathbb{E}[\gamma_T^i | \mathcal{F}_{t-}]$ and hence $\mathcal{O}(\gamma^i) = \mathcal{P}(\gamma^i)$ for \mathbb{F}^i -predictable γ^i .

The first property says that a change of measure does not affect the weight a player assigns to any pure action. By the second property, a mixed action profile remains a mixed action profile under Q^A . Finally, the last statement implies that the infinitesimal average is not affected by a change of measure.

We know from discrete-time repeated games that private strategies may lead to an increase of efficiency when players are allowed to mix their actions; see Kandori and Obara [26]. In a deviation from a PPE, however, players cannot gain anything by using private information.

Lemma 2.3.4. *Suppose that player i 's opponents play a public strategy profile A^{-i} , then player i has a best response in public strategies.*

Proof. Suppose that A^i is a best response to A^{-i} using private information in a filtration $\mathbb{G}^i \supseteq \mathbb{F}$. Similarly to Lemma 2.3.2, in an instantiation γ of A , γ^i is adapted with respect to the augmented filtration generated by \mathbb{G}^i and \mathbb{M}^i such that the \mathbb{G}^i -optional projection of γ^i is \mathbb{G}^i -predictable. Then,

$$\tilde{A}^i(a^i) := a^i \circ (1_{\{\gamma^i = a^i\}}), \quad a^i \in \mathcal{A}^i$$

is \mathbb{F} -predictable by projectivity of the optional projection onto $\mathbb{F} \subseteq \mathbb{G}^i$. Moreover, it defines a public best response to A^{-i} because \mathcal{F}_τ -conditional independence of \mathcal{M}_τ^j and \mathcal{M}_τ^i for any \mathbb{F} -stopping time τ and any $j \neq i$ implies

$$\begin{aligned} \mu(\tilde{A}^i, A^{-i})_\tau 1_{\{\tau < \infty\}} &= \mathbb{E}[\mu(\tilde{\gamma}_\tau^i, \gamma_\tau^{-i}) 1_{\{\tau < \infty\}} \mid \mathcal{F}_\tau] \\ &= \sum_{a \in \mathcal{A}^i} \mu(a) \mathbb{E} \left[\mathbb{E}^{\mathcal{O}, \mathbb{F}^i} (1_{\{\gamma^i = a^i\}})_\tau 1_{\{\tau < \infty\}} \mid \mathcal{F}_\tau \right] \prod_{j \neq i} \mathbb{E}^{\mathcal{O}} (1_{\{A^j = a^j\}})_\tau \\ &= \mu(A^i, A^{-i})_\tau 1_{\{\tau < \infty\}}, \end{aligned}$$

where we use that the \mathbb{F}^i -optional projection $\tilde{\gamma}^i = \mathcal{O}, \mathbb{F}^i(\gamma^i)$ of γ^i is an instantiation of \tilde{A}^i . This implies that $Q^{(\tilde{A}^i, A^{-i})} = Q^A$, hence it follows from (2.3) that $W(\tilde{A}^i, A^{-i}) = W(A)$ a.e. by projectivity of the conditional expectation. \square

Definition 2.3.3. A *mixed strategy* of player i is a mixture over his/her available pure strategies, that is, a probability measure κ^i on the space

$$\mathcal{P}^i = \{\gamma^i : \Omega \times [0, \infty) \rightarrow \mathcal{A}^i \mid \gamma^i \text{ is } \mathbb{F}\text{-predictable}\}^{10}$$

A *mixed strategy profile* κ is the product measure $\kappa = \kappa^1 \otimes \cdots \otimes \kappa^I$ on the product σ -algebra on $\mathcal{P} = \mathcal{P}^1 \times \cdots \times \mathcal{P}^I$ (see Footnote 10), where κ^i is a mixed strategy for player $i = 1, \dots, I$. Player i 's discounted expected future payoff of a mixed strategy profile equals

$$W_t^i(\kappa) = \int_{\mathcal{P}} W_t^i(\gamma) d\kappa(\gamma).$$

In many situations, it is easier to work with behaviour strategies rather than with mixed strategies. The advantage of mixed strategies, however, is the fact that some notions generalize straight from pure strategies to mixed strategies because they take values in pure strategies with probability one. In the remainder of this section we show in an analogue to Kuhn's theorem (see Kuhn [30]) that the two notions are essentially equivalent.

Definition 2.3.4. A mixed strategy profile κ and a behaviour strategy profile A are *realization equivalent* if they lead to the same distribution over outcomes, that is, $A(a) = \kappa(\{\gamma \in \mathcal{P} \mid \gamma = a\})$ a.e. for any $a \in \mathcal{A}$.

¹⁰Formally, κ^i is defined not on \mathcal{P}^i itself, but on the σ -algebra $\sigma\mathcal{P}^i$ on \mathcal{P}^i generated by the coordinate maps $\pi_t : \mathcal{P}^i \rightarrow (\Omega \rightarrow \mathcal{A}^i)$ given by $\pi_t(\gamma^i) := \gamma_t^i$; see also Billingsley [7, pg. 509].

Theorem 2.3.5 (Analogue of Kuhn's theorem). *Every mixed strategy profile is realization equivalent to some behaviour strategy profile. Conversely, every behaviour strategy profile has a realization equivalent mixed strategy profile.*

Proof. Let κ be a mixed strategy profile. Fix a player i and define for any $a^i \in \mathcal{A}^i$ and any $(\omega, t) \in \Omega \times [0, \infty)$,

$$A_t^i(a^i; \omega) := \kappa^i(\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\}).$$

It can be deduced that $A^i(a^i)$ is almost everywhere well defined and a predictable process for all $a^i \in \mathcal{A}^i$. Indeed, the sets

$$S(a^i) = \{(\gamma^i, \omega, t) \in \mathcal{P}^i \times \Omega \times [0, \infty) \mid \gamma_t^i(\omega) = a^i\}$$

are elements of the product σ -algebra of $\sigma\mathcal{P}^i$ (see Footnote 10) and the \mathbb{F} -predictable σ -algebra on $\Omega \times [0, \infty)$. Since the sets $\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\}$ are the (ω, t) -sections of $S(a^i)$, it follows from measurable induction that the mapping

$$(\omega, t) \rightarrow \kappa^i(\{\gamma^i \in \mathcal{P}^i \mid \gamma_t^i(\omega) = a^i\})$$

is measurable with respect to the \mathbb{F} -predictable σ -algebra, which means that $A^i(a^i)$ is predictable. Moreover, the processes $A^i(a^i)$ are nonnegative and their sum over $a^i \in \mathcal{A}^i$ is one. Since $\kappa = \kappa^1 \otimes \cdots \otimes \kappa^I$, it follows that $A(a) := A^1(a^1) \cdots A^I(a^I)$ is a realization equivalent behaviour strategy profile.

Let now A be a behaviour strategy profile and let U^1, \dots, U^I be independent processes with standard uniformly distributed marginals as in the proof of Lemma 2.3.1 on some probability space $(\tilde{\Omega}, \tilde{\mathcal{F}}, \tilde{P})$. For any player i , enumerate

$\mathcal{A}^i = \{a_1^i, \dots, a_{m_i}^i\}$ and define

$$\gamma_t^i(\tilde{\omega}, \omega) = \sum_{n=1}^{m_i} a^n 1_{\{\sum_{k=1}^{n-1} A_t^i(a_k^i; \omega) \leq U_t^i(\tilde{\omega}) < \sum_{k=1}^n A_t^i(a_k^i; \omega)\}},$$

which we consider as a mapping in $\tilde{\omega}$ from $\tilde{\Omega}$ to the set \mathcal{P}^i of predictable processes $\Omega \times [0, \infty) \rightarrow \mathcal{A}^i$. Using the σ -algebra from Footnote 10, this mapping becomes measurable, hence we can define a probability measure κ^i on \mathcal{P}^i as the preimage of γ^i under \tilde{P} , that is, $\kappa^i = \tilde{P} \circ (\gamma^i)^{-1}$. Therefore, $\kappa = \kappa^1 \otimes \dots \otimes \kappa^I$ is indeed a realization equivalent mixed strategy profile. \square

2.4 Incentive compatibility and self-generation

If a stage game is played only once, the only rational outcome is that of a Nash equilibrium, where no player can make any profits by deviating from it. When interactions are repeated indefinitely, it may be possible to support outcomes that dominate any static Nash payoff by using future payoffs to coordinate play. Indeed, players may be motivated to play actions that are not myopic best responses if they are compensated with a sufficiently high continuation payoff. Such a continuation payoff is usually associated to some form of cooperative behaviour leading to a mutually beneficial outcome, where every player is willing to forego some instantaneous profits, understanding that their opponents do the same. If players suspect that one player has deviated from such a cooperative state, they will punish the deviator by playing a strategy profile with a detrimental outcome to the deviator. The threat of such future punishments is what provides intertemporal incentives to players

and deters unilateral deviations. In games with imperfect public monitoring, punishments are necessarily attached to the public signal. We thus start by describing the dependence of players' continuation value on the public signal. The following lemma provides a stochastic differential representation of the continuation value, which is a generalization of Proposition 1 in Sannikov [37] to the multiplayer setting, when players have access to behaviour strategies and information arrives not only continuously, but also abruptly.

Lemma 2.4.1. *For an I -dimensional process W and a behaviour strategy profile A , the following are equivalent:*

- (a) W is the discounted expected payoff under A .
- (b) W is a bounded semimartingale such that for every $i = 1, \dots, I$,

$$\begin{aligned} dW_t^i &= r(W_t^i - g^i(A_t)) dt + r\beta_t^i (\sigma dZ_t - \mu(A_t) dt) \\ &\quad + r \sum_{y \in Y} \delta_t^i(y) (dJ_t^y - \lambda(y|A_t) dt) + dM_t^i \end{aligned} \quad (2.8)$$

for a martingale M^i (strongly) orthogonal to σZ and $(J^y)_{y \in Y}$ with $M^i = 0$, predictable and square-integrable processes β^i and $\delta^i(y)$ for $y \in Y$, satisfying $\mathbb{E}_{Q_t^A} \left[\int_0^T |\beta_t^i|^2 dt \right] < \infty$ and $\mathbb{E}_{Q_t^A} \left[\int_0^T |\delta_t^i(y)|^2 \lambda(y|A_t) dt \right] < \infty$ for every $T \geq 0$ and every $y \in Y$. Moreover, $M \neq 0$ if and only if players use public randomization.

The process $r\beta^i$ is the sensitivity of player i 's continuation value to the continuous component of the public signal, and the processes $r\delta^i(y)$ are the impacts on player i 's continuation value when an event of type $y \in Y$ occurs. The

intuition behind Lemma 2.4.1 is that (2.3) can be rewritten as

$$W_t^i(A) = r e^{rt} \left(\lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^A} \left[\int_0^u e^{-rs} g^i(A_s) ds \mid \mathcal{F}_t \right] - \int_0^t e^{-rs} g^i(A_s) ds \right)$$

and hence $dW_t^i = rW_t^i dt - r g^i(A_t) dt + d\text{“martingale”}$ by the product rule. However, the limiting probability measure Q_∞^A that coincides with Q_t^A on \mathcal{F}_t for every $t \geq 0$ is not equivalent to P on \mathcal{F}_∞ , hence we cannot immediately apply a martingale representation result.¹¹

Proof of Lemma 2.4.1. To show that (a) implies (b), observe first that $W(A)$ is bounded as it remains in \mathcal{V} at all times. Fix a player i and $T > 0$, abbreviate $W^i := W^i(A)$ and derive from (2.3) that

$$\begin{aligned} w_T^i &:= W_T^i - r \int_0^T (W_t^i - g^i(A_t)) dt \\ &= W_T^i + r \int_0^T g^i(A_t) dt - r \int_0^\infty \int_0^{s \wedge T} r e^{-r(s-t)} \mathbb{E}_{Q_s^A} [g^i(A_s) \mid \mathcal{F}_t] dt ds \end{aligned} \quad (2.9)$$

is a bounded \mathcal{F}_T -measurable random variable. Because $(J^y)_{y \in Y}$ are pairwise orthogonal and orthogonal to σZ , the stable subspace generated by σZ and $(J^y)_{y \in Y}$ is the space of all stochastic integrals with respect to these processes (Theorem IV.36 in Protter [36]). Therefore, we obtain the unique martingale representation property for a square-integrable martingale by Corollary 1 to Theorem IV.37 in [36]. That is, for a bounded \mathcal{F}_T -measurable random variable w_T^i , there exist an \mathcal{F}_0 -measurable c_T^i , predictable and square-integrable processes $(\beta_{t,T}^i)_{0 \leq t \leq T}$ and $(\delta_{t,T}^i(y))_{0 \leq t \leq T}$ for $y \in Y$ with $\mathbb{E}_{Q_T^A} \left[\int_0^T |\beta_{t,T}^i|^2 dt \right] < \infty$

¹¹Note that Q_∞^A is not attained by taking the limit in (2.1), but its existence is asserted by Proposition I.7.4 of Karatzas and Shreve [27].

and $\mathbb{E}_{Q_T^A} \left[\int_0^T |\delta_{t,T}^i(y)|^2 \lambda(y|A_t) dt \right] < \infty$ for all $y \in Y$ and a Q_T^A -martingale M^i orthogonal to σZ and $(J^y)_{y \in Y}$ with $M_0^i = 0$ such that

$$w_T^i = c_T^i + \int_0^T r \beta_{t,T}^i (\sigma dZ_t - \mu(A_t) dt) + \sum_{y \in Y} \int_0^T r \delta_{t,T}^i(y) (dJ_t^y - \lambda(y|A_t) dt) + M_{T,T}^i.$$

To prove that (b) holds, we need to show that c_T^i , $\beta_{t,T}^i$, $\delta_{t,T}^i(y)$ and $M_{t,T}^i$ do not depend on T . Let $\tilde{T} \leq T$ and take in (2.9) conditional expectations on $\mathcal{F}_{\tilde{T}}$ under Q_T^A to deduce that

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

Taking $\tilde{T} = 0$, this shows that $c_T^i = W_0^i$ does not depend on T . It also implies

$$w_{\tilde{T}}^i = W_0^i + \int_0^{\tilde{T}} r \beta_{t,T}^i (\sigma dZ_t - \mu(A_t) dt) + \sum_{y \in Y} \int_0^{\tilde{T}} r \delta_{t,T}^i(y) (dJ_t^y - \lambda(y|A_t) dt) + M_{\tilde{T},T}^i,$$

which yields $\beta_{\cdot,T}^i = \beta_{\cdot,\tilde{T}}^i$ and $\delta_{\cdot,T}^i(y) = \delta_{\cdot,\tilde{T}}^i(y)$ for every $y \in Y$ a.e. on $[0, \tilde{T}]$ and $M_{\tilde{T},T}^i = M_{\tilde{T},\tilde{T}}^i$ a.s. by the uniqueness of the orthogonal decomposition. Taking \mathcal{F}_t -conditional expectations, we deduce $M_{t,\tilde{T}}^i = M_{t,T}^i$ a.s. for $t \in [0, \tilde{T}]$, proving that the integral representation is independent of T , and hence $W(A)$

satisfies (b). To show the converse, derive from Itô's formula that

$$\begin{aligned} d(e^{-rt}W_t^i) &= -re^{-rt}g^i(A_t) dt + re^{-rt}\beta_t^i (\sigma dZ_t - \mu(A_t) dt) \\ &\quad + re^{-rt} \sum_{y \in Y} \delta_t^i(y) (dJ_t^y - \lambda(y|A_t) dt) + e^{-rt} dM_t^i. \end{aligned} \quad (2.10)$$

Since M^i is strongly orthogonal to σZ and $(J^y)_{y \in Y}$, it is also strongly orthogonal to the density process in (2.1), and hence it remains a martingale under the change of measure in (2.1). Integrating (2.10) from t to T and taking Q_T^A -conditional expectations on \mathcal{F}_t thus yields

$$W_t^i = \int_t^T re^{-r(s-t)} \mathbb{E}_{Q_s^A} [g^i(A_s) \mid \mathcal{F}_t] ds + e^{-r(T-t)} \mathbb{E}_{Q_T^A} [W_T^i \mid \mathcal{F}_t].$$

Since W is bounded, the second summand converges to zero a.s. as T tends to ∞ , hence W_t^i is indeed the discounted expected future value of A . \square

In discrete-time games, incentives are provided by a *continuation promise* that maps the public signal to a promised continuation payoff for every player; see, for example, Abreu, Pearce and Stacchetti [2]. The representation in (2.8) shows that in continuous-time games, the continuation value is linear in the public signal and hence, so is the continuation promise. The following incentive compatibility condition is the generalization of the respective conditions in Sannikov [37] and Sannikov and Skrzypacz [39].

Definition 2.4.1. An action profile α is *enforceable* if there exists a *continuation promise* (β, δ) with $\beta = (\beta^1, \dots, \beta^I)^\top \in \mathbb{R}^{I \times d}$ and $\delta = (\delta^1, \dots, \delta^I)^\top \in \mathbb{R}^{I \times m}$ such that for every player i , the sum of expected instantaneous payoff rate $g^i(\alpha)$

and promised continuation rate $\beta^i \mu(\alpha) + \delta^i \lambda(\alpha)$ is maximized in α^i . That is, for $i = 1, \dots, I$ and every $a^i \in \mathcal{A}^i$,

$$g^i(\alpha) + \beta^i \mu(\alpha) + \delta^i \lambda(\alpha) \geq g^i(a^i, \alpha^{-i}) + \beta^i \mu(a^i, \alpha^{-i}) + \delta^i \lambda(a^i, \alpha^{-i}) \quad \text{a.s.} \quad (2.11)$$

A behaviour strategy profile is *enforceable* if there exist processes $(\beta_t)_{t \geq 0}$, $(\delta_t)_{t \geq 0}$ such that (2.11) is satisfied a.e.¹²

Suppose that players keep their promises and the continuation promise used to enforce A are, in fact, the sensitivities of its continuation value to the public signal. Then no player has an incentive to deviate at any point in time and the strategy profile is an equilibrium. This is formalized in the following lemma, which is the continuous-time analogue to the one-shot deviation principle. It is a generalization of Proposition 2 in Sannikov [37] to our setting.

Lemma 2.4.2. *A strategy profile A is a PPE if and only if (β, δ) related to A by (2.8) enforces A .*

Proof. Fix a behaviour 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 (2.8), integrating (2.10) from t to u yields

$$\begin{aligned} W_t^i(A) &= - \int_t^u r e^{-r(s-t)} \left(\beta_s^i (\sigma dZ_s - \mu(A_s) ds) - g^i(A_s) ds - dM_s^i \right) \\ &\quad - \sum_{y \in Y} \int_t^u r e^{-r(s-t)} \delta_t^i(y) (dJ_s^y - \lambda(y|A_s) ds) + e^{-r(u-t)} W_u^i(A). \end{aligned}$$

¹²It is enough to consider deviations to pure strategies since any behaviour strategy has a realization equivalent mixed strategy by Theorem 2.3.5, and a deviation to a mixed strategy can only be profitable if it has at least one profitable pure strategy in its support.

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

$$\begin{aligned} W_t^i(\tilde{A}) &= \lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^{\tilde{A}}} \left[\int_t^u re^{-r(s-t)} g^i(\tilde{A}_s) ds \middle| \mathcal{F}_t \right] \\ &= W_t^i(A) + \lim_{u \rightarrow \infty} \mathbb{E}_{Q_u^{\tilde{A}}} \left[\int_t^u re^{-r(s-t)} \left((g^i(\tilde{A}_s) - g^i(A_s)) ds \right. \right. \\ &\quad \left. \left. + \beta_s^i (\sigma dZ_s - \mu(A_s) ds) + \sum_{y \in Y} \delta_s^i(y) (dJ_s^y - \lambda(y|A_s) ds) \right) \middle| \mathcal{F}_t \right] \text{ a.s.} \end{aligned}$$

Because the processes β and $\delta(y)$, $y \in Y$ are constructed using a martingale representation result for the bounded random variable w_T^i in (2.9), the processes $\int_t^\cdot re^{-r(s-t)} \beta_s^i (\sigma dZ_s - \mu(A_s) ds)$ and $\int_t^\cdot re^{-r(s-t)} \delta_s^i(y) (dJ_s^y - \lambda(y|A_s) ds)$ are bounded mean oscillation (BMO) martingales under the probability measure Q_u^A up to any time $u > t$. Since Assumption 2.1.1 implies that the jumps of $(\lambda(y|A_s) - 1) \Delta J_s^y$ in (2.1) are bounded from below by $-1 + \varepsilon$ for any $y \in Y$, it follows from Remark 3.3 and Theorem 3.6 in Kazamaki [28] that $\int_t^\cdot re^{-r(s-t)} \beta_s^i (\sigma dZ_s - \mu(\tilde{A}_s) ds)$ and $\int_t^\cdot re^{-r(s-t)} \delta_s^i(y) (dJ_s^y - \lambda(y|\tilde{A}_s) ds)$ are BMO-martingales under $Q_u^{\tilde{A}}$. Together with Fubini's theorem, this implies

$$\begin{aligned} W_t^i(\tilde{A}) - W_t^i(A) &= \int_t^\infty re^{-r(s-t)} \mathbb{E}_{Q_s^{\tilde{A}}} \left[g^i(\tilde{A}_s) - g^i(A_s) \right. \\ &\quad \left. + \beta_s^i (\mu(\tilde{A}_s) - \mu(A_s)) + \delta_s^i (\lambda(\tilde{A}_s) - \lambda(A_s)) \right] ds \text{ a.s.} \end{aligned} \tag{2.12}$$

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

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

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

on Ξ . Set $\tilde{A}^i := \hat{A}^i 1_{\Xi} + A^i 1_{\Xi^c}$. Because β and δ are predictable, we can and do choose Ξ and \hat{A} to be predictable as well. In particular, \tilde{A}^i is a behaviour strategy for player i . For such an \tilde{A} , the expectation in (2.12) is strictly positive for $t = 0$, which is a contradiction. \square

Lemmas 2.4.1 and 2.4.2 motivate how we construct equilibrium profiles in continuous time — as solution to (2.8) subject to the enforceability constraint in (2.11). However, because we consider repeated games with an infinite time horizon and a terminal payoff does not exist, we cannot apply results from the theory of backward stochastic differential equations to find a solution. Instead, we use time-homogeneity of repeated games to construct forward solutions similarly the techniques in discrete-time repeated games. Abreu, Pearce and Stacchetti [2] introduced the notion of self-generating payoff sets, these are, sets \mathcal{W} of payoffs that can be attained by an incentive compatible continuation promise that remains within the set. If promises are kept, the payoff in the next period is in \mathcal{W} , implying that it is attainable again by an incentive compatible continuation promise that remains in \mathcal{W} by self-generation. This gives rise to an iterative procedure of constructing continuation values and associated strategy profiles. Because actions are enforceable in every period, the resulting strategy profile is a PPE. The following is the definition of a self-generating payoff set in a continuous-time setting.




Figure 2.4: Since \mathcal{W} is self-generating, the drift rate $w - g(a) - \delta\lambda(a)$ has to point towards the interior of the set, that is, $N_w^\top (g(a) + \delta\lambda(a) - w) \geq 0$. Moreover, the diffusion $r\beta_t (dZ_t - \mu(\alpha) dt)$ has to be tangential to $\partial\mathcal{W}$ as the continuation value would escape \mathcal{W} immediately otherwise. Finally, an event of type $y \in Y$ incurs a jump in the continuation value of size $r\delta(y)$. Since W cannot jump outside of \mathcal{W} , it is necessary that $w + r\delta(y) \in \mathcal{W}$.

Definition 2.4.2. A set $\mathcal{W} \subseteq \mathbb{R}^I$ is *self-generating* for discount rate $r > 0$ if for every $w \in \mathcal{W}$ there exists a solution $(W, A, \beta, \delta, Z, (J^y)_{y \in Y}, M)$ to (2.8) such that (β, δ) enforces A , $W_0 = w$ a.s. and $W_\tau \in \mathcal{W}$ a.s. for every stopping time τ .

Because there is no minimal time step, this definition does not immediately give rise to an iterative procedure as in discrete time. In addition, we need to construct a suitable sequence of stopping times $(\tau_n)_{n \geq 0}$ with $\tau_n \rightarrow \infty$ as n approaches ∞ , at which the solutions are concatenated. On each of the sets $([\tau_n, \tau_{n+1}])$, self-generation imposes certain restrictions on possible incentives at the boundary \mathcal{W} .¹³ Indeed, motivated by (2.8) and illustrated in Figure 2.4, at any $w \in \partial\mathcal{W}$ these restrictions are:

1. Inward-pointing drift: $N_w^\top (g(a) + \delta\lambda(a) - w) \geq 0$,
2. Tangential volatility: $N_w^\top \beta = 0$,
3. Jumps within the set: $w + r\delta(y) \in \mathcal{W}$ for every $y \in Y$.

¹³ $([\tau_n, \tau_{n+1}]) := \{(w, t) \in \Omega \times [0, \infty) \mid \tau_n(w) < t \leq \tau_{n+1}(w)\}$ is called a *stochastic interval*.

Lemma (2.4.2) implies that any self-generating payoff set is contained in $\mathcal{E}(r)$. This fact is used in the proofs of the folk theorems in Chapter 3 because it guarantees that a construction satisfying conditions 1–3 yields equilibrium profiles. For the explicit characterization of $\mathcal{E}(r)$ for two-player games in Chapter 4, we additionally need to show that $\mathcal{E}(r)$ is self-generating itself, making it the largest bounded self-generating set.

Proposition 2.4.3. *The set $\mathcal{E}(r)$ is the largest bounded self-generating set.*

This result is the equivalent of Theorem 1 in Abreu, Pearce and Stacchetti [2]. The intuition behind the result is that continuation strategies of PPE are incentive compatible after any history. Because the continuation game after any time τ is equivalent to the game starting at time 0 as seen in (2.3), it follows that $W_\tau \in \mathcal{E}(r)$ a.s. This last conclusion, however, is subject to some subtle measurability issues that we need to address before a formal proof of Proposition 2.4.3. Without restrictions on β and δ , solutions to (2.8) are weak solutions. That is, the components of the public signal Z , $(J^y)_{y \in Y}$ and the probability space are part of the solution. The probability space thus depends on the payoff that is being attained, i.e.,

$$\mathcal{E}(r) := \left\{ w \in \mathcal{V} \left| \begin{array}{l} \text{There exists } (\Omega, \mathcal{F}, \mathbb{F}, P) \text{ containing a Brownian} \\ \text{motion } Z, \text{ independent Poisson processes } (J^y)_{y \in Y} \\ \text{and a PPE } A \text{ with } W_0(A) = w \text{ } P\text{-a.s.} \end{array} \right. \right\}.$$

We call $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$ a *stochastic framework* for A .

Remark 2.4.1. Technically, this weak definition is necessary to ensure existence of the solutions. But even aside from the technical advantages, the weak

solution concept is appropriate here. From an interpretation standpoint, the difference between a strong solution and a weak solution to a stochastic differential equation lies in the causality of the noise. If the noise is defined exogenously and not affected by players' actions, this corresponds to a strong solution. In games of imperfect information, however, the noise is induced by players' strategies and thus cannot be fixed at the beginning. This is in line with discrete-time games, where we only care about the distribution of the public signal and not on what probability space the distribution is realized.

At any stopping time τ , the value $W_\tau(A)$ is a random variable. Because of the weak formulation, the probability space depends on the payoff $w \in \mathcal{E}(r)$, and hence it is not clear what measurability conditions a random variable in $\mathcal{E}(r)$ should satisfy. This is clarified by the following lemma, whose proof relies on the construction of regular conditional probabilities.

Lemma 2.4.4. *For an \mathcal{F}_0 -measurable random variable W^* in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$ the following are equivalent:*

- (a) $W^* \in \mathcal{E}(r)$ a.s.,
- (b) *There exist a strategy profile A , square-integrable and predictable processes β and $\delta(y)$ for $y \in Y$, a martingale M (strongly) orthogonal to σZ and $(J^y)_{y \in Y}$, and a bounded semimartingale W such that β enforces A , $W_0 = W^*$ a.s. and the processes $W, A, \beta, \delta, Z, (J^y)_{y \in Y}$ and M are related by the stochastic differential equation (2.8).*

Proof. Let first $W^* \in \mathcal{E}(r)$ a.s. Although a PPE may exist on different probability spaces in (a) for each realization $W^* = w$, we can use the fact that

the models all share the same path space to construct a regular conditional probability on that space. The path space of a behaviour strategy profile A and its stochastic framework is given by $\Delta(\mathcal{A})^{[0,\infty)} \times \mathcal{D}^d$, where \mathcal{D}^d is the space of càdlàg functions from $[0, \infty)$ into \mathbb{R}^d . By Theorem A.2.2 in Kallenberg [24], there exists a metric on \mathcal{D}^d that induces the Skorohod topology, under which \mathcal{D}^d is complete and separable. Since \mathcal{V} is a closed subset of \mathbb{R}^I and $\Delta(\mathcal{A})^{[0,\infty)}$ is compact by Tychonoff's theorem (see Theorem 37.3 in Munkres [32]), it follows that $\Omega = \mathcal{V} \times \Delta(\mathcal{A})^{[0,\infty)} \times \mathcal{D}^d$ is complete and separable as well. Therefore, by Theorem V.3.19 in Karatzas and Shreve [27], there exists a regular conditional probability $P_x(F) : \mathcal{V} \times \mathcal{F} \rightarrow [0, 1]$, which means

1. for each $w \in \mathcal{V}$, P_w is a probability measure on (Ω, \mathcal{F}) ,
2. for each $F \in \mathcal{F}$, the mapping $x \mapsto P_w(F)$ is $Borel(\mathcal{V})$ -measurable,
3. for each $F \in \mathcal{F}$, $P_w(F) = P(F | W^* = w)$ for ν -a.e. $w \in \mathcal{V}$, where ν is the distribution of W^* .

We know that for each $w \in \mathcal{E}(r)$, there exists a PPE A^w achieving w . Let A be the process defined pointwise by A^w on $\{W^* = w\}$. It follows from the properties of a regular conditional probability that A is a PPE achieving W^* . Indeed, for any player i and any behaviour strategy profile \tilde{A} with $\tilde{A}^{-i} = A^{-i}$,

$$\begin{aligned} P(W_0^i(A) \geq W_0^i(\tilde{A})) &= \int_{\mathcal{E}(r)} P(W_0^i(A) \geq W_0^i(\tilde{A}) | W^* = w) d\nu(w) \\ &= \int_{\mathcal{E}(r)} P_w(W_0^i(A^w) \geq W_0^i(\tilde{A})) d\nu(w) = 1 \end{aligned}$$

and in the same way $P(W_0(A) = W^*) = \int_{\mathcal{E}(r)} P_w(W_0(A^w) = w) d\nu(w) = 1$.

To show the converse, suppose that $W^* \notin \mathcal{E}(r)$ on an \mathcal{F}_0 -measurable set Ξ with $\nu(\Xi) > 0$. Then there exists an \mathcal{F}_0 -measurable set $\tilde{\Xi}$ with $\nu(\tilde{\Xi}) > 0$ such that some player i can improve his/her strategy to $\tilde{A}^{w,i}$ for $w \in \tilde{\Xi}$. Letting $\hat{A}^i := A^i 1_{\Xi^c}(w) + \tilde{A}^{w,i} 1_{\Xi}(w)$, it follows that

$$P(W_0^i(A) \geq W_0^i(\hat{A}^i, A^{-i})) = \int_{\mathcal{V}} P_w(W_0^i(A) \geq W_0^i(\tilde{A}^{w,i}, A^{-i})) d\nu(w) < 1,$$

contradicting the assumption that A is a PPE. \square

Proof of Proposition 2.4.3. By Lemma 2.4.2, any bounded self-generating set \mathcal{W} is contained in $\mathcal{E}(r)$. Since $\mathcal{E}(r)$ is bounded, it remains to show that $\mathcal{E}(r)$ is self-generating. Take $w \in \mathcal{E}(r)$ so that Lemma 2.4.1 yields the existence of a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$, a behaviour strategy profile A enforced by (β, δ) , a martingale M orthogonal to σZ and a bounded semimartingale W satisfying (2.8) with $W_0 = w$ a.s. We now fix a stopping time τ and show that $W_\tau \in \mathcal{E}(r)$ a.s. To do so, we set $W^* = W_\tau$, $\tilde{\mathcal{F}}_t = \mathcal{F}_{\tau+t}$, $\tilde{Z}_t = Z_{\tau+t} - Z_\tau$, $\tilde{J}_t^y := J_{\tau+t}^y - J_\tau^y$ for every $y \in Y$, $\tilde{M}_t = M_{\tau+t} - M_\tau$, $\tilde{W}_t = W_{\tau+t}$, $\tilde{\beta}_t = \beta_{\tau+t}$, $\tilde{\delta}_t = \delta_{\tau+t}$ and $\tilde{A}_t = A_{\tau+t}$. Because the tilde-processes and -filtrations satisfy condition (b) in Lemma 2.4.4, we obtain that $W_\tau = W^* \in \mathcal{E}(r)$ a.s. \square

2.5 Execution speed and the use of public randomization

In this section we derive quantitative relations between the informativeness of the public signal, players' patience, the speed at which players execute their strategies and public randomization. While these relations are expected

from discrete time, most of them do not have a closed-form representation in a discrete-time setting. Previously, the drift function μ and the volatility σ of the continuous component of the public signal and the intensities $\lambda(y|\cdot)$ of events of type $y \in Y$ were fixed. In this section we view the PPE payoff set $\mathcal{E}(r, \mu, \sigma, \lambda)$ as an object depending on all four game primitives and show how the game primitives can be exchanged for each other. We show that $\mathcal{E}(r, \mu, \sigma, \lambda)$ depends on these game primitives only through the informativeness of the public signal relative to players' discounting $\gamma = \sigma^\top(\sigma\sigma^\top)^{-1}\mu/\sqrt{r}$ and λ/r . Moreover, the equilibrium payoff set $\mathcal{E}(r, \mu, \sigma, \lambda)$ increases as this relative informativeness increases componentwise. A corollary to this result is the fact that $\mathcal{E}(r)$ is monotonically decreasing in the discount rate r .

As a first result, we prove the following law for the exchangeability of game primitives, showing that scaling the discount rate by κ has the same effect on equilibrium payoffs as dividing the signal-to-noise ratio $\sigma^\top(\sigma\sigma^\top)^{-1}\mu$ by $\sqrt{\kappa}$ and the intensities λ of infrequent events by κ . For two-player games, this feature can also be observed from the differential equation characterizations of $\mathcal{E}(r)$ in Theorem 2 in Sannikov [37] and our Theorem 4.2.1.¹⁴ In games with an entirely continuous public signal, this implies a square-root law that is similar in spirit to Corollary 1 in Faingold and Sannikov [12], where a square-root law is obtained for continuous-time games between one long-lived player and a continuum of short-lived players. Contrary to their result, however, we also show that strategies can be transformed by a time-change to attain the exact same paths of the continuation value at a different speed.

¹⁴Indeed, the denominator in (4.3) depends on (r, μ, σ) only through $\sigma^\top(\sigma\sigma^\top)^{-1}\mu/\sqrt{r}$ and the numerator of (4.3) is unaffected: the change in λ is compensated by the restriction that $w + r\delta(y) \in \mathcal{E}(r)$, that is, δ is scaled by $1/\kappa$ when r and λ are scaled by κ .

Lemma 2.5.1. *Let $\tilde{\mu} : \mathcal{A} \rightarrow \mathbb{R}^{d_c}$ and $\tilde{\sigma} \in \mathbb{R}^{d_c \times d_z}$ be such that $\tilde{\sigma}\tilde{\sigma}^\top$ is invertible and $\tilde{\sigma}^\top(\tilde{\sigma}\tilde{\sigma}^\top)^{-1}\tilde{\mu} = \sqrt{\kappa}\sigma^\top(\sigma\sigma^\top)^{-1}\mu$ for some $\kappa > 0$. Then a strategy profile A is a PPE for the game primitives $(r, \mu, \sigma, \lambda)$ if and only if $(A_{\kappa t})_{t \geq 0}$ is a PPE with respect to the game primitives $(\kappa r, \tilde{\mu}, \tilde{\sigma}, \kappa\lambda)$. Moreover, for every $t \geq 0$,*

$$\tilde{W}_t((A_{\kappa s})_{s \geq 0}; \kappa r, \tilde{\mu}, \tilde{\sigma}, \kappa\lambda) = W_{\kappa t}(A; r, \mu, \sigma, \lambda) \quad a.s., \quad (2.13)$$

where \tilde{W} is the discounted expected future payoff with respect to the time-changed filtration $(\mathcal{F}_{\kappa t})_{t \geq 0}$ and adjusted reference probability measure.¹⁵ In particular, the equilibrium payoff set depends on $(r, \mu, \sigma, \lambda)$ only through the ratios $\sigma^\top(\sigma\sigma^\top)^{-1}\mu/\sqrt{r}$ and λ/r .

For $\kappa = 1$, the result says that the continuation value depends on μ and σ only through the informativeness of the continuous component of the public signal, the signal-to-noise-ratio $\sigma^\top(\sigma\sigma^\top)^{-1}\mu$. Because the induced probability measure in (2.1) depends on μ and σ only through that quantity, such a transformation leads to the same distribution over possible signals. For $\kappa < 1$, Lemma 2.5.1 says that players becoming more patient has the same effect as increasing the informativeness of the signal. As time becomes less valuable to the players, a longer interval of observations of the public signal becomes available at the same cost, hence players can better estimate the drift rate μ of its continuous component and the intensities $\lambda(y|\cdot)$ of the rare events. Players are considered to be more patient when their discount rate is lower because they value future payoffs more. In this class of games, being more patient can be taken very literally, by executing the same strategy profile at a slower

¹⁵The reference probability measure is adjusted such that $(J_{\kappa t}^y)_{t \geq 0}$ all have intensity 1.

speed. Similarly, players being less patient has the same effect as a decrease in the informativeness of the signal. Note that $\tilde{\mu}$ and $\tilde{\lambda}$ are invertible linear transformations of μ and λ , hence (2.11) is solvable for μ and λ if and only if it is solvable for $\tilde{\mu}$ and $\tilde{\lambda}$. Thus, the same action profiles are enforceable.

Proof of Lemma 2.5.1. For a strategy profile A in a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$, define the time-changed processes

$$\tilde{A}_t := A_{\kappa t}, \quad \tilde{Z}_t := \frac{1}{\sqrt{\kappa}} Z_{\kappa t}, \quad \tilde{J}_t^y := J_{\kappa t}^y \quad \text{for } y \in Y, \quad \frac{d\tilde{P}_t}{dP} := \prod_{y \in Y} \kappa^{-J_{\kappa t}^y} e^{(\kappa-1)t}.$$

Observe that \tilde{A} is predictable with respect to the time-changed filtration $\tilde{\mathbb{F}} = (\tilde{\mathcal{F}}_t)_{t \geq 0}$. For any $T > 0$, \tilde{P}_T defines a probability measure equivalent to P such that on $[0, T]$, $(\tilde{J}^y)_{y \in Y}$ are Poisson processes with intensity 1 under \tilde{P}_T . Arguing as in Proposition I.7.4 of Karatzas and Shreve [27], there exists a unique probability measure \tilde{P} that agrees with \tilde{P}_t on $\tilde{\mathcal{F}}_t$ for every $t > 0$. Thus, \tilde{P} will serve as new reference probability measure, under which $(\tilde{J}^y)_{y \in Y}$ have intensity 1 on $[0, \infty)$. By the scaling property of Brownian motion, \tilde{Z} is an $\tilde{\mathbb{F}}$ -Brownian motion under both P and \tilde{P} . Define a family $(\tilde{Q}_t^{\tilde{A}})_{t \geq 0}$ of probability measures induced by \tilde{A} with respect to $\tilde{\mu}$, $\tilde{\sigma}$, $\tilde{\lambda} = \kappa\lambda$, \tilde{Z} , $(\tilde{J}^y)_{y \in Y}$ and \tilde{P} analogously as in (2.1). Because \tilde{P} is equivalent to P on \mathcal{F}_t for any $t > 0$, $(\tilde{Q}_t^{\tilde{A}})_{t \geq 0}$ can also be represented via a density process with respect to P . Since $(\tilde{J}_t^y)_{t \geq 0}$ are Poisson processes with intensity κ under P , it follows that

$$\begin{aligned} \frac{d\tilde{Q}_t^{\tilde{A}}}{dP} &= \mathcal{E}_t \left(\int_0^\cdot \tilde{\mu}(\tilde{A}_s)^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\sigma} d\tilde{Z}_s + \sum_{y \in Y} \int_0^\cdot \left(\frac{\tilde{\lambda}(y|\tilde{A}_s)}{\kappa} - 1 \right) (d\tilde{J}_s^y - \kappa ds) \right) \\ &= \mathcal{E}_t \left(\int_0^\cdot \tilde{\mu}(A_{\kappa s})^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\sigma} \frac{dZ_{\kappa s}}{\sqrt{\kappa}} + \sum_{y \in Y} \int_0^\cdot (\lambda(y|A_{\kappa s}) - 1) (dJ_{\kappa s}^y - \kappa ds) \right). \end{aligned}$$

Since $\tilde{\sigma}^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\mu} / \sqrt{\kappa} = \sigma^\top (\sigma \sigma^\top)^{-1} \mu$ by assumption, substituting $d\tilde{s} = \kappa ds$ leads to

$$\frac{d\tilde{Q}_t^{\tilde{A}}}{dP} = \mathcal{E}_{\kappa t} \left(\int_0^\cdot \mu(A_{\tilde{s}})^\top (\sigma \sigma^\top)^{-1} \sigma dZ_{\tilde{s}} + \sum_{y \in Y} \int_0^\cdot (\lambda(y|A_{\tilde{s}}) - 1) (dJ_{\tilde{s}}^y - d\tilde{s}) \right) = \frac{dQ_{\kappa t}^A}{dP}$$

and hence $\tilde{Q}_t^{\tilde{A}}$ coincides with $Q_{\kappa t}^A$ on $\tilde{\mathcal{F}}_t = \mathcal{F}_{\kappa t}$. Observe that the expected flow payoff $g^i(a) = f^i(a^i, \mu(a), \lambda(a)) = f^i(a^i, \sigma \tilde{\sigma}^\top (\tilde{\sigma} \tilde{\sigma}^\top)^{-1} \tilde{\mu}(a) / \sqrt{\kappa}, \tilde{\lambda}(a) / \kappa)$ still depends on a^{-i} only through $\tilde{\mu}(a)$ and $\tilde{\lambda}(a)$. Substituting $d\tilde{s} = \kappa ds$ again, we obtain for every $t \geq 0$,

$$\begin{aligned} \tilde{W}_t^i(\tilde{A}; \kappa r, \tilde{\mu}, \tilde{\sigma}, \kappa \lambda) &:= \int_t^\infty \kappa r e^{-\kappa r(s-t)} \mathbb{E}_{\tilde{Q}_s^{\tilde{A}}} [g^i(\tilde{A}_s) | \tilde{\mathcal{F}}_t] ds \\ &= \int_{\kappa t}^\infty r e^{-r(\tilde{s}-\kappa t)} \mathbb{E}_{Q_s^A} [g^i(A_{\tilde{s}}) | \mathcal{F}_{\kappa t}] d\tilde{s} \\ &= W_{\kappa t}^i(A; r, \mu, \sigma, \lambda) \text{ a.s.} \end{aligned} \tag{2.14}$$

where we used that $\tilde{Q}_s^{\tilde{A}}$ and Q_s^A coincide on $\tilde{\mathcal{F}}_s$.¹⁶ Because all unilateral deviations of $(A_{\kappa t})_{t \geq 0}$ correspond to unilateral deviations of A , it follows from (2.14) that $(A_{\kappa t})_{t \geq 0}$ is a PPE with respect to $(\kappa r, \tilde{\mu}, \tilde{\sigma}, \kappa \lambda)$ if and only if A is a PPE with respect to $(r, \mu, \sigma, \lambda)$. \square

The quantities $\gamma = \sigma^\top (\sigma \sigma^\top)^{-1} \mu / \sqrt{r}$ and λ/r are measures of the informativeness of the public signal adjusted for the patience of players. It follows

¹⁶To ensure that the sets of measure 0 integrate to a set of measure 0, consider

$$\int_t^\infty \kappa r e^{-\kappa r(s-t)} \mathbb{E}_{\tilde{Q}_s^{\tilde{A}}} [g^i(\tilde{A}_s) | \tilde{\mathcal{F}}_t] ds = \sum_{k=1}^\infty \kappa r \mathbb{E}_{\tilde{Q}_{t+k}^{\tilde{A}}} \left[\int_{t+k-1}^{t+k} e^{-\kappa r(s-t)} g^i(\tilde{A}_s) ds \mid \tilde{\mathcal{F}}_t \right],$$

hence we are summing over only countably many sets of measure 0.

from (2.13) at time 0 that the equilibrium payoff set depends on the game primitives $(r, \mu, \sigma, \lambda)$ only through γ and λ/r . For games with a Brownian signal, this is similar to the equilibrium analysis in Daley and Green [10], where an exogenous news process is observed that is driven by Brownian motion. The quality of the news process, a quantity corresponding to $\gamma^\top \gamma$, plays a central role in their equilibrium analysis.

The following result shows how to transform a PPE A when the continuous component of the public signal becomes more informative relative to players' discounting, that is, $\hat{\gamma} \geq \gamma$ in every component. Essentially, players generate artificial noise through public randomization to reduce the relative informativeness of the signal back to γ .

Theorem 2.5.2. *Let A be a PPE with respect to $(r, \mu, \sigma, \lambda)$. Let $\hat{\mathbb{F}} = (\hat{\mathcal{F}}_t)_{t \geq 0}$ denote the filtration generated by the public filtration \mathbb{F} and the addition of a public randomization device. Denote by $\hat{W}_t(\cdot)$ the discounted expected future payoff conditional on $\hat{\mathcal{F}}_t$. Then there exists a PPE \hat{A} with respect to the transformed game parameters below by adding public randomization to¹⁷*

(a) A with $\hat{W}_t(\hat{A}; r, \mu, \sigma \Lambda, \lambda) = W_t(A; r, \mu, \sigma, \lambda)$ a.s. for any symmetric Λ in $\mathbb{R}^{d_z \times d_z}$ with eigenvalues in $[-1, 1]$ such that $\ker(\sigma \Lambda) = \ker(\sigma)$.

(b) A with $\hat{W}_t(\hat{A}; r, \Lambda^{-1} \mu, \sigma, \lambda) = \hat{W}_t(\hat{A}; r, \mu, \Lambda \sigma, \lambda) = W_t(A; r, \mu, \sigma, \lambda)$ a.s. for any symmetric $\Lambda \in \mathbb{R}^{d_c \times d_c}$ with eigenvalues in $[-1, 1] \setminus \{0\}$ and $\Lambda \sigma \sigma^\top = \sigma \sigma^\top \Lambda$.

(c) $(A_{\kappa t})_{t \geq 0}$ with $\hat{W}_t(\hat{A}; \kappa r, \mu, \sigma, \kappa \lambda) = W_{\kappa t}(A; r, \mu, \sigma, \lambda)$ a.s. for $\kappa \in (0, 1)$.

¹⁷By adding public randomization to a strategy profile A , we mean that disregarding the information of the public randomization device in \hat{A} is identical to playing A . Formally, $\mathcal{O}\hat{A} = A$, where $\mathcal{O}(\cdot)$ is the optional projection onto the public filtration \mathbb{F} .

This is a remarkable result because public randomization allows one to transform PPE continuously over games with a higher quality of the signal, achieving exactly the same path of continuation values. For discrete-time games, Proposition 1 in Kandori [25] shows that the one-period decomposition of payoffs can be attained with the same expected continuation payoff if the Blackwell-informativeness of the signal is increased. However, in discrete time there is no way to link this result to the strategy profile or the evolution of the continuation value over time. Returning to the question of how to achieve equilibrium payoffs also under a smaller discount rate, we see in (c) that this is achieved by performing a time change and then adding public randomization to the time-changed strategy profile. We obtain the following corollary.

Corollary 2.5.3. *If players have access to a public randomization device, then it follows that $\mathcal{E}(r) \subseteq \mathcal{E}(r')$ for any $0 < r' < r$.*

Proof of Theorem 2.5.2. To show (a), we need to show that a PPE A with respect to $(r, \mu, \sigma, \lambda)$ can be transformed to a PPE with respect to $(r, \mu, \sigma\Lambda, \lambda)$. Let $(\Omega, \mathcal{F}, \mathbb{F}, P, Z, (J^y)_{y \in Y})$ be the stochastic framework of A and let $\beta, \delta, Z, (J^y)_{y \in Y}$ and M be the processes from Lemma 2.4.1 that satisfy (2.8) for A and $W = W(A)$ with respect to μ and σ . Let Z^\perp be an d_z -dimensional Brownian motion orthogonal to both Z and M and denote by $\hat{\mathbb{F}}$ the augmented filtration generated by \mathbb{F} and Z^\perp . Write $\Lambda = Q^\top DQ$ for orthogonal Q and diagonal

$$D = \begin{pmatrix} \mathbb{I}_k & 0 & 0 \\ 0 & -\mathbb{I}_n & 0 \\ 0 & 0 & \tilde{D} \end{pmatrix}$$

such that \tilde{D} has diagonal entries in $(-1, 1)$. Define $\tilde{\Lambda} := Q^\top \sqrt{\mathbb{I}_{d_z} - \tilde{D}^2} Q$ so that $\Lambda^2 + \tilde{\Lambda}^2 = \mathbb{I}_{d_z}$. Then $\hat{Z} := \Lambda Z + \tilde{\Lambda} Z^\perp$ is a Brownian motion with respect to $\hat{\mathbb{F}}$. Set

$$\hat{\Lambda} := Q^\top \begin{pmatrix} 0 & 0 \\ 0 & \left(\sqrt{\mathbb{I}_{d_z - k - n} - \tilde{D}^2}\right)^{-1} \end{pmatrix} Q$$

and define $\hat{Z}^\perp := \hat{\Lambda}(Z - \Lambda Z)$. It follows from

$$d\langle \hat{Z}^\perp, \hat{Z}^\perp \rangle_t = \begin{pmatrix} 0 & 0 \\ 0 & \mathbb{I}_{d_z - k - n} \end{pmatrix} dt$$

and $d\langle \hat{Z}^\perp, \hat{Z} \rangle_t = 0$ that \hat{Z}^\perp is a martingale orthogonal to \hat{Z} . This gives us the decomposition $Z = \Lambda \hat{Z} + \tilde{\Lambda} \hat{Z}^\perp$ and hence

$$\begin{aligned} dW_t &= r(W_t - g(A_t) - \delta_t \lambda(A_t)) dt + r \sum_{y \in Y} \lambda(y|A_t) dJ_t^y \\ &\quad + r\beta_t(\sigma dZ_t - \mu(A_t) dt) + dM_t \\ &= r(W_t - g(A_t) - \delta_t \lambda(A_t)) dt + r \sum_{y \in Y} \lambda(y|A_t) dJ_t^y \\ &\quad + r\beta_t(\sigma \Lambda d\hat{Z}_t - \mu(A_t) dt) + r\beta_t \sigma \tilde{\Lambda} d\hat{Z}_t^\perp + dM_t. \end{aligned}$$

Since $\ker(\sigma \Lambda) = \ker(\sigma)$, it follows that there exists a matrix $\Delta \in \mathbb{R}^{d_c \times d_c}$ such that $\Delta \sigma = \sigma \Lambda$.¹⁸ Therefore, $\sigma \Lambda \hat{Z} = \Delta^2 \sigma Z + \sigma \Lambda \tilde{\Lambda} Z^\perp$ is orthogonal to M and

¹⁸Indeed, σ is an isomorphism from $\ker(\sigma)^\perp$ to \mathbb{R}^{d_c} with inverse $\sigma^\top(\sigma\sigma^\top)^{-1}$ and hence $b_i := \sigma^\top(\sigma\sigma^\top)^{-1}e_i$ for $i = 1, \dots, d_c$ form a basis of $\ker(\sigma)^\perp = \ker(\sigma\Lambda)^\perp$, where e_i denotes the i^{th} standard basis vector in \mathbb{R}^{d_c} . Define the vectors $\delta_i := \sigma \Lambda b_i$ for $i = 1, \dots, d_c$ and let

hence also to $\hat{M} := M + \int r\beta_s\sigma\tilde{\Lambda} d\hat{Z}_s^\perp$. It follows that W also fulfills (2.8) for $\sigma\Lambda$ with processes β , \hat{Z} , $(J^y)_{y \in Y}$ and \hat{M} . Since β enforces A , the strategy profile A is also a PPE in the continuous-time game with volatility $\sigma\Lambda$ by Lemma 2.4.1. Note, however, that A is a PPE as an $\hat{\mathbb{F}}$ -predictable process, that is, when players use the new orthogonal information suitably. We derive from Itô's formula that

$$\begin{aligned} d(e^{-rt}W_t) &= -re^{-rt}(g(A_t) + \delta_t\lambda(A_t)) dt + \sum_{y \in Y} re^{-rt}\lambda(y|A_t) dJ_t^y \\ &\quad + re^{-rt}\hat{\beta}_t(\sigma\Lambda d\hat{Z}_t - \mu(A_t) dt) + e^{-rt} d\hat{M}_t. \end{aligned} \quad (2.15)$$

$\ker(\sigma\Lambda) = \ker(\sigma)$ implies $\text{rank}(\sigma\Lambda) = d_c$, hence we can define a family $(\tilde{Q}_t^A)_{t \geq 0}$ of probability measures as in (2.1) with respect to μ and $\sigma\Lambda$. Integrating (2.15) from t to T and taking \tilde{Q}_T^A -conditional expectations on $\hat{\mathcal{F}}_t$ yields

$$W_t = \int_t^T e^{-r(s-t)} \mathbb{E}_{\tilde{Q}_s^A} [g(A_s) | \hat{\mathcal{F}}_t] ds + e^{-r(T-t)} \mathbb{E}_{\tilde{Q}_T^A} [W_T | \hat{\mathcal{F}}_t].$$

Taking the limit as $T \rightarrow \infty$ yields $W_t(A; r, \mu, \sigma, \lambda) = \hat{W}_t(A; r, \mu, \sigma\Lambda, \lambda)$ a.s. since W is bounded. This concludes the proof of (a).

For statement (b), consider first $\tilde{\sigma} = \Lambda\sigma$ and define the $d_z \times d_z$ matrix $\Lambda' := \sigma^\top(\sigma\sigma^\top)^{-1}\Lambda\sigma$ so that $\sigma\Lambda' = \Lambda\sigma$. Note that $\Lambda'^\top = \sigma^\top\Lambda(\sigma\sigma^\top)^{-1}\sigma = \Lambda'$, i.e., Λ' is symmetric. Every vector in the kernel of σ is an eigenvector of Λ' with eigenvalue 0. Let κ be an eigenvalue of Λ' to an eigenvector v which is not in the kernel of σ . Then $\Lambda\sigma v = \sigma\Lambda'v = \kappa\sigma v$, that is, κ is an eigenvalue

$\Delta = (\delta_1, \dots, \delta_{d_c})$ be the matrix with column vectors δ_i . By construction, $\Delta\sigma b_i = \delta_i = \sigma\Lambda b_i$. Since (b_1, \dots, b_{d_c}) can be completed to a basis of \mathbb{R}^{d_z} with any basis of $\ker(\sigma\Lambda)^\perp$ and $\ker(\sigma)^\perp = \ker(\sigma\Lambda)^\perp$, it follows that $\Delta\sigma = \sigma\Lambda$.

of Λ for eigenvector σv . Since this applies to all eigenvalues of Λ' outside the kernel of σ , the eigenvalues of Λ' lie in $[-1, 1]$ and $\ker(\sigma\Lambda') = \ker(\sigma)$. Moreover, $\sigma\Lambda' = \Lambda\sigma$ has rank d_c because Λ is invertible and σ has rank d_c . The statement now follows by applying (a) to Λ' . Observe that the change $\tilde{\mu} = \Lambda^{-1}\mu$ is completely equivalent since

$$\sigma^\top \Lambda^\top (\Lambda \sigma \sigma^\top \Lambda^\top)^{-1} \mu = \sigma^\top \Lambda \Lambda^{-1} (\sigma \sigma^\top)^{-1} \Lambda^{-1} \mu = \sigma^\top (\sigma \sigma^\top)^{-1} \Lambda^{-1} \mu$$

and hence the induced probability measures coincide at all times. For (c), Lemma 2.5.1 implies $\tilde{W}_t((A_{\kappa s})_{s \geq 0}, \kappa r, \mu, \sigma/\sqrt{\kappa}, \lambda/\kappa) = W_{\kappa t}(A, r, \mu, \sigma, \lambda)$ a.s. for $\kappa \in (0, 1)$. The statement follows from (a) for the matrix $\Lambda = \text{diag}_{d_z}(\sqrt{\kappa})$ applied to $\sigma/\sqrt{\kappa}$. \square

Chapter 3

Folk Theorems in Games with Continuous Information

Folk theorems constitute key results in the theory of repeated games, providing sufficient conditions under which players can achieve socially efficient outcomes as they get increasingly patient. Results of this kind are called folk theorems because their statement is folklore in games with perfect observation. Indeed, if players perfectly observe each others' actions, deviations of a strategy profile are immediately detected and the deviator can be punished in future periods. If these punishments are sufficiently severe and players are patient enough, such a deviation is unprofitable and becomes impossible in equilibrium.

In games with imperfect information, deviations are not unambiguously detectable and coordination becomes more difficult. Punishments are necessarily attached to outcomes of the public signal, and hence punishments will occur from time to time even if no player has deviated. As a result, players will never achieve perfect efficiency in equilibrium, but one may wonder whether

they achieve asymptotic efficiency as players become arbitrarily patient, that is, as the discount rate r goes to 0. In discrete time, this question has been answered affirmatively by Fudenberg, Levine and Maskin [16], who show that a folk result holds if deviations of any two players can be statistically distinguished by observing the public signal. In this chapter, we show how the techniques in [16] can be extended to a continuous-time setting with Brownian information, that is, when $Y = \emptyset$. The restriction to continuous monitoring is natural for a folk theorem as we briefly discuss in Section 3.5.

The standard version of the folk theorem is the minmax folk theorem, stating that any feasible and individually rational payoff is attainable in equilibrium if players are sufficiently patient, that is, $\mathcal{E}(r) \rightarrow \mathcal{V}^*$ as $r \rightarrow 0$. This strong version of the folk theorem may not be necessary to ensure asymptotic efficiency, however, as \mathcal{V}^* includes many Pareto-inefficient payoffs. In many applications, a Nash-threat folk theorem is sufficient, stating that $\mathcal{E}(r)$ extends to the set \mathcal{V}^0 of all Pareto-efficient payoffs dominating a static Nash payoff.

The importance of folk theorems lies in the ability to verify whether a system is well designed. If efficiency is impossible even when players are arbitrarily patient, one may have to consider redesigning the game. In the climate agreement example of Section 2.2, payoffs on the Pareto-efficient frontier correspond to payoffs where the terms of the agreement are upheld by its signatories. Because this is ultimately the goal of signing such an agreement, a Nash-threat folk theorem gives the possibility to check whether the observed information is sufficient to enforce the agreement in the long run.

Corollary 2.5.3 shows that $\mathcal{E}(r)$ is increasing in players' patience if players have access to a public randomization device. Despite the fact that Corol-

lary 2.5.3 requires public randomization, the folk theorem does not rely on public randomization. This is because we make a statement about $\mathcal{E}(r)$ only indirectly, by showing that any smooth payoff set in the interior of \mathcal{V}^* is self-generating for sufficiently small discount rates. By taking smooth inner approximations $\mathcal{W}_n \rightarrow \mathcal{V}^*$, the folk theorem follows from Proposition 2.4.3. To show that a compact payoff set \mathcal{W} is self-generating, we construct equilibrium strategies with continuation values in \mathcal{W} with the following steps:

1. For any payoff $w \in \mathcal{W}$, there exist W^w, A^w, β^w related by (2.8) such that $W_0^w = w$, β^w enforces A^w and W^w remains in \mathcal{W} for a short but positive amount of time τ_w . The discount rate r_w may depend on the payoff w .
2. Solutions to (2.8) are uniform in a neighbourhood U_w of w , that is, there exists $\tilde{r} > 0$ such that for any $r \in (0, \tilde{r})$, the time τ and the strategy profile A can be chosen uniformly across U_w .
3. By compactness, a concatenation of these local solutions is a global solution to (2.8), which is a PPE by Lemma 2.4.2.

As we have explained in Section 2.4, for payoffs w on the boundary $\partial\mathcal{W}$, a solution W^w to (2.8) remains in \mathcal{W} only if the drift rate points towards the interior of \mathcal{W} and the volatility is tangential to the set \mathcal{W} . This means we have to find sufficient conditions for incentives β to be constructed on tangent hyperplanes. The uniformity condition in Step 2 means that these solutions exist on a fixed probability space with a given Brownian motion for the entire neighbourhood U_w , that is, locally, these are strong solutions to (2.8). This is important when we concatenate these local solutions to a

global solution in Step 3: by compactness of \mathcal{W} , we have to deal with only finitely many probability spaces. Indeed, for any finite subcover U_1, \dots, U_n , we enlarge the associated probability spaces and we concatenate the local solutions at a suitable sequence of stopping times $(\tau_n)_{n \geq 0}$. By independence and stationarity of Brownian increments, we can choose $(\tau_{n+1} - \tau_n)_{n \geq 0}$ to be independent and identically distributed as $\tau = \min_{k=1, \dots, n} \tau_{U_k}$. Because $\tau > 0$ a.s., this countable concatenation yields a global solution.

Requiring strong solutions to (2.8) entails that the constructed equilibrium profiles are constant on each of the intervals $(\tau_n, \tau_{n+1}]$. Because the concatenation is countable, the resulting strategies have bounded oscillation. This is a very desirable feature from both an implementation and an interpretation standpoint as agents can switch actions only finitely many times on finite time intervals. It also shows that, despite the continuous-time framework, it is not necessary to act infinitesimally fast to attain efficiency in the limit. This provides a partial rebuttal to the concern that continuous-time models lead to strategies that cannot be implemented because they possibly exhibit unbounded oscillation.

The remainder of this chapter is structured as follows. In Section 3.1, we introduce the notion of enforceability on tangent hyperplanes, which is necessary for Steps 1 and 2 above. In Section 3.6, we carry out the above construction of equilibrium profiles in detail. A set \mathcal{W} for which this is possible is called *uniformly decomposable on tangent hyperplanes*. Folk theorems are thus reduced to finding sufficient conditions on the game primitives such that any smooth $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes for sufficiently small discount rates. We do this in Section 3.2 when players are restricted to

pure strategies. If players are allowed to mix their actions, we show in Section 3.3 how approximation results can significantly lower the conditions on game primitives if the public signal is sufficiently high dimensional. We elaborate on the finite-variation property of our constructed equilibrium strategies in Section 3.4 and discuss extensions to abrupt information in Section 3.5.

3.1 Pairwise identifiability and enforceability on hyperplanes

Definition 3.1.1.

1. Let $T \in \mathbb{R}^{I \times (I-1)}$ be a matrix whose column vectors T_1, \dots, T_{I-1} span a hyperplane $H \subseteq \mathbb{R}^I$. An action profile α is *enforceable on hyperplane H* if there exists a matrix $B \in \mathbb{R}^{(I-1) \times d_c}$ such that α is enforced by $\beta = TB$.
2. A matrix $\beta \in \mathbb{R}^{I \times d_c}$ *enforces α orthogonal* to vector $N \in \mathbb{R}^I$ if it enforces α and satisfies $N^\top \beta = 0$.

Observe that the two notions of enforceability are equivalent, i.e., α is enforceable on a hyperplane H if and only if it is enforceable orthogonal to the normal vector N of H . Indeed, if $\beta = TB$, then $N^\top \beta = 0$. Conversely, if $N^\top \beta = 0$, then all column vectors β_j lie in H , which means they can be written as linear combinations of the T_j . This is equivalent to $\beta = TB$.

Enforceability on a hyperplane means that players transfer continuation value amongst each other to compensate players for which α is not a best response. The matrix β determines the rate at which these values are transferred. We distinguish two types of hyperplanes.

Definition 3.1.2. A hyperplane H is said to be *coordinate* if it is orthogonal to a coordinate axis. H is *regular* if it is not coordinate.

For an enforceable action profile α , the additional requirement to be enforceable on a coordinate hyperplane means that the corresponding player does not make any transfers. Such an action profile α thus necessarily involves a best response of player i . Indeed, the system (2.11) has a solution with $\beta^i = 0$ if and only if α^i is a best response to α^{-i} . We state this as a lemma.

Lemma 3.1.1. *An enforceable action profile α is enforceable on a hyperplane orthogonal to the i^{th} coordinate axis if and only if α satisfies the best response property for player i , that is, $g^i(\alpha) \geq g^i(a^i, \alpha^{-i})$ for all $a^i \in \mathcal{A}^i$.*

For an action profile α to be enforceable on regular hyperplanes, players' impacts on the distribution of the public signal need to be sufficiently identifiable. Let $M^i(\alpha)$ denote the $(d_c \times |\mathcal{A}^i|)$ -dimensional matrix, whose column vectors $\mu(a^i, \alpha^{-i}) - \mu(\alpha)$, $a^i \in \mathcal{A}^i$ are given by the impact on the drift rate of the public signal that player i 's deviation from α^i to a^i has. Observe that $\text{rank } M^i(\alpha) \leq |\mathcal{A}^i| - 1$ since multilinearity implies

$$\sum_{a^i \in \mathcal{A}^i} \alpha^i(a^i) (\mu(a^i, \alpha^{-i}) - \mu(\alpha)) = 0.$$

Definition 3.1.3. A mixed action profile α is *pairwise identifiable* if for any two players i and $j \neq i$, it holds that $\text{span } M^i(\alpha) \cap \text{span } M^j(\alpha) = \{0\}$.

Pairwise identifiability means that deviations of any two players lead to linearly independent impacts on the drift rate of the public signal. Therefore, deviations of any two players can be statistically distinguished. The next

result is the analogue of Lemma 5.5 in Fudenberg, Levine and Maskin [16]. The proof shows that under the assumption of pairwise identifiability, any two players' incentives are isolated by an orthogonal decomposition.

Lemma 3.1.2. *Suppose that an enforceable action profile α is pairwise identifiable. Then it is enforceable on all regular hyperplanes.*

Proof. We show that α is enforceable orthogonal to the normal vector N of the hyperplane. Because the hyperplane is regular, N has at least 2 non-zero entries and we will assume that these are the first two. Let $\beta \in \mathbb{R}^{I \times d_c}$ enforce α . Pairwise identifiability implies that $\mathbb{R}^{d_c} = (\Lambda^i(\alpha) \cap \Lambda^j(\alpha))^\perp = \Lambda^i(\alpha)^\perp + \Lambda^j(\alpha)^\perp$ for all $i \neq j$ and hence $\beta^i = \beta^{\perp,i} + \tilde{\beta}^i$, where $\beta^{\perp,i} \perp \Lambda^i(\alpha)$ for all i and

$$\tilde{\beta}^1 \perp \Lambda^2(\alpha) \quad \text{and} \quad \tilde{\beta}^i \perp \Lambda^1(\alpha), \quad i = 2, \dots, I.$$

Let $G^i(\alpha)$ denote the row vector of losses $g^i(\alpha) - g^i(a^i, \alpha^{-i})$ in player i 's expected flow payoff when she switches from α^i to a^i , so that α is enforceable if and only if $G^i(\alpha) \geq \beta^i M^i(\alpha)$ holds componentwise for every player $i = 1, \dots, I$. We construct $B = (B^1, \dots, B^I)^\top$ enforcing α on H by setting

$$B^1 = \tilde{\beta}^1 - \sum_{i=2}^I \frac{N^i}{N^1} \tilde{\beta}^i, \quad B^2 = \tilde{\beta}^2 - \frac{N^1}{N^2} \tilde{\beta}^1 \quad \text{and} \quad B^i = \tilde{\beta}^i, \quad i = 3, \dots, I. \quad (3.1)$$

Indeed, since $\beta^{\perp,1}$ and $\tilde{\beta}^2, \dots, \tilde{\beta}^I$ are orthogonal to $\Lambda^1(\alpha)$, it follows that

$$B^1 M^1(\alpha) = \tilde{\beta}^1 M^1(\alpha) - \sum_{i=2}^I \frac{N^i}{N^1} \tilde{\beta}^i M^1(\alpha) = (\beta^1 - \beta^{\perp,1}) M^1(\alpha) \leq G^1(\alpha),$$

The inequalities for players $i = 2, \dots, I$ are verified in the same manner. Finally, note that $N^\top B = 0$ by construction. \square

A strong solution to (2.8) entails that it is possible to find (W, A, β, M) solving (2.8) for a fixed Brownian motion Z . The conditions for existence of strong solutions are quite stringent, requiring that $M \equiv 0$, that is, players do not use public randomization, $A \equiv \alpha$, i.e., players do not change their action profile and β is a Lipschitz continuous functional of W . Because we choose \mathcal{W} to be smooth (a non-empty convex set with C^2 boundary), $w \mapsto N_w$ is Lipschitz on the boundary $\partial\mathcal{W}$, where N_w is the unique outward-pointing normal vector at w . Since the concatenation of (locally) Lipschitz continuous maps is again (locally) Lipschitz continuous, it remains to find sufficient conditions such that the construction $N \mapsto \beta$ in Lemma 3.1.2 is locally Lipschitz continuous. As we can see from (3.1), this may be tricky where the tangent hyperplane changes from being regular to being coordinate. The following lemma states various conditions such that β is locally bounded and Lipschitz continuous.

Lemma 3.1.3. *Let $N \in \mathbb{R}^I \setminus \{0\}$ and let α be an enforceable action profile. Suppose that one of the following conditions holds true:*

1. α is pairwise identifiable and N is not parallel to any coordinate axis,
2. α is pairwise identifiable and enforceable orthogonal to N ,
3. α is enforceable orthogonal to e_i and α^i is a unique best response to α^{-i} , that is, $\alpha^i \in \mathcal{A}^i$ and $g^i(\alpha) > g^i(a^i, \alpha^{-i})$ for every $a^i \in \mathcal{A}^i \setminus \{\alpha^i\}$,
4. α is a static Nash equilibrium.

Then there exist a neighbourhood U_N of N and a bounded, Lipschitz continuous map $\beta_\alpha : U_N \rightarrow \mathbb{R}^{I \times d_c}$ such that $\beta_\alpha(x)$ enforces α orthogonal to x .

Proof. Statement 1: Let $\beta_\alpha^i(N) = B^i$ as in (3.1), which is locally Lipschitz continuous in N . The statement holds by choosing U_N such that the first two coordinates are bounded away from zero.

Statement 2: By statement 1, it is enough to consider N coordinate. Suppose $N = e_1$, let $\tilde{\beta}^1, \dots, \tilde{\beta}^I$ be defined as in the proof of Lemma 3.1.2 and set

$$\beta_\alpha^1(x) := - \sum_{i=2}^I \frac{x^i}{x^1} \tilde{\beta}^i, \quad \beta_\alpha^i(x) = \tilde{\beta}^i, \quad i = 2, \dots, I. \quad (3.2)$$

Along the lines of the proof of Lemma 3.1.2, it follows that $\beta_\alpha(x)$ enforces α orthogonal to x if $x_1 \neq 0$. The statement follows by choosing U_N bounded away from $\{x^1 = 0\}$.

Statement 3: Suppose $N = e_1$ and that $\alpha^1 \in \mathcal{A}^1$ is a unique best response to α^{-1} . Let $\beta_\alpha^i(x)$ as in (3.2), except that $\tilde{\beta}^i$ are replaced by β^i . Then, clearly, (2.11) is fulfilled for players $i = 2, \dots, I$. Because of the unique best response property, there exists an $\varepsilon > 0$ such that $g^1(\alpha^1) \geq g^1(a^1, \alpha^{-1}) + \varepsilon$ for every $a^1 \in \mathcal{A}^1 \setminus \{\alpha^1\}$. Let $B = \max_{i=2, \dots, I} \max_{a^1 \in \mathcal{A}^1} |\beta^i(\mu(a^1, \alpha^{-1}) - \mu(\alpha))|$, which is finite because β is fixed. If $B = 0$, then α is a Nash equilibrium and the result holds by statement 4. Suppose therefore that $B > 0$. Then for all x in

$$U_{e_1} := \left\{ x \in \mathbb{R}^I \mid \|x - e_1\| \leq \frac{\varepsilon}{B(I-1) + \varepsilon} \right\},$$

x^1 is bounded away from 0 and hence $|x^i|/x^1 \leq \varepsilon/(B(I-1))$. It follows that

$$|\beta_\alpha^1(x)(\mu(a^1, \alpha^{-1}) - \mu(\alpha))| = \sum_{i=2}^I \frac{|x^i|}{x^1} |\beta^i(\mu(\alpha) - \mu(\alpha))| \leq \varepsilon$$

for every $a^1 \in \mathcal{A}^1$. Together with the unique best response property for player 1,




Figure 3.1: Fudenberg, Levine and Maskin [16] show how to decompose payoffs $w \in \partial\mathcal{W}$ in discrete time into a current-period payoff $g(\alpha)$ and a continuation payoff $u(X)$ parallel to the tangent hyperplane T_w . In contrast to continuous time, payoffs v in a neighbourhood of w are decomposable with respect to the same hyperplane.

this shows that $\beta_\alpha(x)$ enforces α for $x \in U_{e_1}$. By construction, $\beta_\alpha(x)^\top x = 0$ and β_α is Lipschitz continuous and bounded since x_1 is bounded away from 0.

Statement 4: This is clear because $\beta_\alpha(x) = 0$ for all $x \in \mathbb{R}^I$. \square

Because continuation payoffs are bounded away from the separating hyperplane T_w in discrete time, an action profile can be enforced on a nearby hyperplane by moving the continuation payoff by a small (and constant) amount; see also Figure 3.1. In continuous time, Lemma 3.1.3 is necessary to ensure that an action profile can be enforced on nearby hyperplanes without changing value transfers between players significantly. This additional requirement is the reason why the condition of decomposability on tangent hyperplanes (cf. Fudenberg, Levine and Maskin [16]) needs to be strengthened to the following notion of uniform decomposability.

Definition 3.1.4. A smooth payoff set \mathcal{W} is *uniformly decomposable on tangent hyperplanes* if for any $w \in \partial\mathcal{W}$ with outward normal N_w , there exists an enforceable action profile α with $g(\alpha)$ strictly separated from \mathcal{W} by the tangent hyperplane T_w so that (α, N_w) satisfies one of the conditions of Lemma 3.1.3.




Figure 3.2: At every point $w \in \partial\mathcal{W}$ the tangential diffusion $r\beta_t(\sigma dZ_t - \mu(\alpha) dt)$ leads to an outward-pointing drift of order $r^2 \text{tr}(\beta\sigma\sigma^\top\beta^\top) dt$. For sufficiently small r , this term is dominated by the inward-pointing drift $r(W_t - g(A_t)) dt$.

The following proposition shows that uniform decomposability on tangent hyperplanes is indeed a sufficient condition for payoff set to be self-generating. It is the continuous-time analogue to Theorem 4.1 of Fudenberg, Levine and Maskin [16]. Its proof is deferred to Section 3.6.

Proposition 3.1.4. *Suppose that a smooth, compact set $\mathcal{W} \subseteq \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes. Then there exists a discount rate \tilde{r} such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for any $r \in (0, \tilde{r})$.*

Folk theorems are thus obtained by finding sufficient conditions on the game primitives such that any smooth and compact payoff set $\mathcal{W} \subseteq \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes. We do this in the next two sections for pure and behaviour strategies, respectively.

Remark 3.1.1. One may wonder whether it is possible to show self-generation of a smooth payoff set \mathcal{W} without locally constructing strong solutions to (2.8). At first sight, such a proof could potentially work with the weaker notion of decomposability on tangent hyperplanes as in discrete time. Note, however, that a closed set \mathcal{W} is self-generating only if β is tangential at the bound-

ary $\partial\mathcal{W}$. It follows from Itô's formula that the tangential diffusion leads to an outward-pointing drift of order $r^2 \text{tr}(\beta\sigma\sigma^\top\beta^\top) dt$; see the proof of Proposition 3.1.4 in Section 3.6 for details. To ensure that \mathcal{W} is self-generating, this outward-pointing drift has to be compensated by the inward pointing drift of order $r(W_t - g(A_t)) dt$ as indicated in Figure 3.2. This is possible for r sufficiently small if β is bounded. Because the construction in (3.1) is locally Lipschitz continuous where it is locally bounded, locally constructing strong solutions comes at no cost.

3.2 Folk theorems in pure strategies

For a payoff set \mathcal{W} to be uniformly decomposable on tangent hyperplanes, there has to exist an enforceable action profile a such that $g(a)$ is separated from \mathcal{W} by T_w and (a, N_w) satisfies one of the conditions in Lemma 3.1.3. For regular payoffs, this means that enforceable action profiles have to be pairwise identifiable. If suitable conditions are met for coordinate payoffs, we can thus decompose smooth payoff sets \mathcal{W} in the interior of $\mathcal{V}^* \cap \mathcal{V}^\dagger$, where

$$\mathcal{V}^\dagger := \text{conv } g(\{a \in \mathcal{A} \mid a \text{ is enforceable and pairwise identifiable}\} \cup \mathcal{A}^N).$$

Let $\mathcal{A}^{(i)} \subseteq \mathcal{A}$ denote the pure action profiles that maximize player i 's expected flow payoff over \mathcal{A} . The continuous-time folk theorems differ from their discrete-time counterparts in the additional assumption that for every player i , one element of $\mathcal{A}^{(i)}$ and the minmax profile $\underline{a}_i = (\underline{a}_i^i, \underline{a}_i^{-i})$ against player i are either pairwise identifiable or satisfy the unique best response

property for player i . These conditions ensure that Conditions 2 and 3 of Lemma 3.1.3 are satisfied at coordinate payoffs.

Theorem 3.2.1. *Suppose that for every player i ,*

1. *the minmax profile \underline{a}_i is enforceable and it is either pairwise identifiable or satisfies the unique best response property for player i ,*
2. *there exists an enforceable action profile $a_i^* \in \mathcal{A}^{(i)}$ that is either pairwise identifiable or satisfies the unique best response property for player i .*

Then for any compact, smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^ \cap \mathcal{V}^\dagger$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Corollary 3.2.2 (Minmax folk theorem in pure strategies). *Suppose that Conditions 1 and 2 of Theorem 3.2.1 are met and that all pure action profiles achieving extremal payoffs are enforceable and pairwise identifiable. Then for any compact, smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Proof of Theorem 3.2.1. Since \mathcal{W} is contained in the interior of \mathcal{V}^\dagger , at any payoff $w \in \partial\mathcal{W}$ there exists an enforceable and pairwise identifiable action profile a such that $g(a)$ is separated from \mathcal{W} by T_w . If T_w is regular, then (a, N_w) satisfies Condition 1 of Lemma 3.1.3. If T_w is coordinate to the i^{th} axis, then w either maximizes or minimizes player i 's payoff on \mathcal{W} by convexity. If w maximizes player i 's payoff, then $g(a_i^*)$ is separated from \mathcal{W} by T_w because a_i^* maximizes g^i over \mathcal{A} . Moreover, (a_i^*, e_i) satisfies either Condition 2 or 3 of Lemma 3.1.3 by assumption. If w minimizes player i 's payoff, then $g(\underline{a}_i)$ is

separated from \mathcal{W} by T_w because $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$. Again by assumption, (\underline{a}_i, e_i) satisfies either Condition 2 or 3 of Lemma 3.1.3. \square

In the remainder of this section we establish the weaker Nash-threat folk theorem, stating sufficient conditions for players to attain asymptotic efficiency dominating static Nash behaviour.

Definition 3.2.1.

1. An action profile α *Pareto-dominates* a profile $\tilde{\alpha}$ if $g^i(\alpha) \geq g^i(\tilde{\alpha})$ for every player i and $g^j(\alpha) > g^j(\tilde{\alpha})$ for at least one player j .
2. An action profile is *Pareto-efficient* if it is not Pareto-dominated by any other action profile.

Lemma 3.2.3. *Suppose that $g^i(a) = b^i(a^i)\mu(a) - c^i(a^i)$, i.e., g is affine in μ . Then any Pareto-efficient pure action profile is enforceable.*

Proof. Fix a Pareto-efficient pure action profile $a \in \mathcal{A}$. Because its payoff is on the “upper right” boundary of \mathcal{V} , there exists a direction $N \in \mathbb{R}^I$ with $N^i > 0$ for $i = 1, \dots, I$ such that $g(a) = \arg \max_{v \in \mathcal{V}} N^\top v$. Then β with row vectors $\beta^i := \sum_{j \neq i} b^j(a^j)N^j/N^i$ enforces a . Indeed, for every $\tilde{a}^i \in \mathcal{A}^i$, we have

$$\begin{aligned}
g^i(\tilde{a}^i, a^{-i}) + \beta^i \mu(\tilde{a}^i, a^{-i}) &= g^i(\tilde{a}^i, a^{-i}) + \frac{1}{N^i} \sum_{j \neq i} (N^j g^j(\tilde{a}^i, a^{-i}) + N^j c^j(a^j)) \\
&= \frac{1}{N^i} \sum_{j=1}^I N^j g^j(\tilde{a}^i, a^{-i}) + \frac{1}{N^i} \sum_{j \neq i} N^j c^j(a^j) \\
&\leq \frac{1}{N^i} \sum_{j=1}^I N^j g^j(a) + \frac{1}{N^i} \sum_{j \neq i} N^j c^j(a^j) \\
&= g^i(a) + \beta^i \mu(a). \quad \square
\end{aligned}$$

Because in continuous-time games incentives can be provided only linearly to the public signal, it is necessary to have an affine payoff structure for Pareto-efficient payoffs to be enforceable. Note that Pareto-dominance is transitive and irreflexive, hence for every player i there exists at least one Pareto-efficient action profile globally maximizing player i 's Payoff. Lemmas 3.1.1 and 3.2.3 thus imply that it is enforceable on the corresponding coordinate hyperplane.

Corollary 3.2.4. *Suppose that g is affine in μ . Then, for every player i , there exists an enforceable Pareto-efficient pure action profile $a_i^p \in \mathcal{A}^{(i)}$. In particular, it is enforceable on the hyperplane coordinate to the i^{th} axis.*

Proof. Observe that action profiles in $\mathcal{A}^{(i)}$ are Pareto-dominated only by other action profiles in $\mathcal{A}^{(i)}$ because Pareto-dominance of some $\tilde{a} \in \mathcal{A}$ over $a \in \mathcal{A}^{(i)}$ entails $g^i(\tilde{a}) \geq g^i(a)$. Since the relation of Pareto-dominance is transitive and irreflexive, there cannot be any circular relations on $\mathcal{A}^{(i)}$. Because there are only finitely many elements in $\mathcal{A}^{(i)}$, at least one element is not dominated, hence Pareto-efficient. Such an action profile a_i^p is enforceable by Lemma 3.2.3 and because it is a static best response for player i , it is enforceable on the hyperplane orthogonal to the i^{th} axis due to Lemma 3.1.1. \square

Theorem 3.2.5 (Nash-threat folk theorem in pure strategies). *Suppose that g is affine in μ , that there exists a Nash equilibrium a_e in pure actions and that Pareto-efficient action profiles are pairwise identifiable. Let \mathcal{V}^0 denote the convex hull of $g(a_e)$ and the Pareto-efficient payoffs Pareto-dominating $g(a_e)$. Then for any compact, smooth set \mathcal{W} in the interior of \mathcal{V}^0 , there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Proof. It follows from Lemma 3.2.3 that $\mathcal{V}^0 \subseteq \mathcal{V}^\dagger$. Because \mathcal{V}^0 Pareto-dominates $g(a_e) \in \mathcal{V}^*$, it follows that also $\mathcal{V}^0 \subseteq \mathcal{V}^*$. The proof thus works in the same

way as for Theorem 3.2.1, with the following changes for coordinate payoffs. If w maximizes player i 's payoff on \mathcal{W} , then a_i^p from Corollary 3.2.4 decomposes w with (a_i^p, e_i) satisfying Condition 3 of Lemma 3.1.3. For w minimizing player i 's payoff, $(a_e, -e_i)$ satisfies Condition 4 of Lemma 3.1.3 instead. \square

3.3 Folk theorems in behaviour strategies

For the standard folk theorem to hold in pure strategies, we essentially need that all action profiles attaining extremal payoffs in \mathcal{V}^* are enforceable and pairwise identifiable. By considering strategies in mixed actions, it is possible to approximate these pure action payoffs $g(a)$ by a sequence of mixed action payoffs $(g(\alpha_n))_{n \geq 0}$, where each α_n has the desired properties while its limit a does not. More specifically, for uniform decomposability we need that an approximation of the minmax profile is either pairwise identifiable or has the unique best response property. While the unique best response property of the minmax profile carries over to approximations by linearity of the expectation, pairwise identifiability of the approximation requires the stronger notion of pairwise full rank; see Lemma 3.3.4 below for details.

Definition 3.3.1.

1. An action profile α has *individual full rank for player i* if $M^i(\alpha)$ has rank $|\mathcal{A}^i| - 1$. It has *individual full rank* if this is true for every player.
2. An action profile α is said to have *pairwise full rank for players i and j* if $M^{ij}(\alpha) = [M^i(\alpha), M^j(\alpha)]$ has rank $|\mathcal{A}^i| + |\mathcal{A}^j| - 2$. An action profile has *pairwise full rank* if this is true for all pairs of players $j \neq i$.

As we show in the proof of the next lemma, individual full rank implies that the system of inequalities (2.11) can be solved with equality, thus any action profile with individual full rank is enforceable. Pairwise full rank is equivalent to individual full rank and pairwise identifiability.

Lemma 3.3.1. *An action profile α is enforceable if for every player i one of the following conditions holds:*

1. α has individual full rank for player i ,
2. α^i is a best response to α^{-i} .

The enforceability condition (2.11) imposes I systems of linear inequalities, one for each player $i = 1, \dots, I$. Because $\text{rank } M^i(\alpha) \leq |\mathcal{A}^i| - 1$, we cannot simply solve the system i by applying the left-inverse of $M^i(\alpha)$, but we need to additionally exploit that the linear dependence amongst the columns of $G^i(\alpha)$ and $M^i(\alpha)$ is the same, where $G^i(\alpha)$ denotes the row vector of expected losses in player i 's instantaneous payoff rate by switching from α^i to $a^i \in \mathcal{A}^i$.

Proof. Fix a player i and suppose first that Condition 1 is satisfied for action profile α . Let $a^i \in \mathcal{A}^i$ be an action with $\alpha^i(a^i) > 0$ and enumerate $\mathcal{A}^i = \{a_1^i, \dots, a_{m_i}^i\}$ such that $a^i = a_{m_i}^i$ is the last element. For the sake of brevity, denote by $M_j^i(\alpha)$ the column of $M^i(\alpha)$ corresponding to action a_k^i . Because of the linear dependence amongst the columns of $M^i(\alpha)$, we obtain

$$M_{m_i}^i(\alpha) = - \sum_{k=1}^{m_i-1} \frac{\alpha(a_k^i)}{\alpha(a_{m_i}^i)} M_k^i(\alpha).$$

Condition 1 implies that there is no other linear dependence amongst the columns of $M^i(\alpha)$ and thus the $d_c \times (|\mathcal{A}^i| - 1)$ -dimensional submatrix $\tilde{M}^i(\alpha)$

consisting of the first $|\mathcal{A}^i| - 1$ columns has full column rank. In particular, $\tilde{M}^i(\alpha)$ has a left-inverse $\tilde{M}_L^i(\alpha)$ and

$$\beta^i = G^i(\alpha) \begin{pmatrix} \tilde{M}_L^i(\alpha) \\ 0 \end{pmatrix}$$

solves the system for player i with equality. Indeed,

$$\begin{aligned} G^i(\alpha) \begin{pmatrix} \tilde{M}_L^i(\alpha) \\ 0 \end{pmatrix} M^i(\alpha) &= G^i(\alpha) \begin{pmatrix} \mathbb{I}_{m_i-1} & \sum_{k=1}^{m_i-1} \frac{\alpha(a_k^i)}{\alpha(a_{m_i}^i)} \tilde{M}_L^i(\alpha) \tilde{M}_k^i(\alpha) \\ 0 & 0 \end{pmatrix} \\ &= \left(G_1^i(\alpha), \dots, G_{m_i-1}^i(\alpha), - \sum_{k=1}^{m_i-1} \frac{\alpha(a_k^i)}{\alpha(a_{m_i}^i)} G_k^i(\alpha) \right), \end{aligned}$$

where we used that $\tilde{M}_L^i(\alpha) \tilde{M}_k^i(\alpha) = e_k$ for every $k = 1, \dots, m_i$. The claim under Condition 1 follows since

$$G_{m_i}^i(\alpha) = - \sum_{k=1}^{m_i-1} \frac{\alpha(a_k^i)}{\alpha(a_{m_i}^i)} G_k^i(\alpha).$$

Under Condition 2, $\beta^i = 0$ solves the inequalities for player i . □

Let \mathcal{V}^Δ denote the set of payoffs achievable in mixed actions. While it may be strictly smaller than \mathcal{V} for some games (see Figure 3.3), the extremal payoffs always correspond to pure action profiles and hence are contained in \mathcal{V}^Δ .

Lemma 3.3.2. *Suppose that for every pair of players i, j , there exists a mixed action profile α^{ij} having ij -pairwise full rank. Then the set of payoffs C of action profiles with pairwise full rank for all pairs of players is dense in \mathcal{V}^Δ .*

	L	R
T	(1, 5)	(2, 2)
B	(0, 0)	(5, 1)




Figure 3.3: The right panel shows \mathcal{V}^Δ and \mathcal{V}^* for the stage game with payoffs given in the table to the left. In the proof of the folk theorems, we need the approximation of extremal payoffs in \mathcal{V}^* and the minmax payoffs only, which all lie in \mathcal{V}^Δ .

Proof. Let $E \subseteq \Delta(\mathcal{A})$ denote the set of mixed action profiles with pairwise full rank. By Lemma 6.2 of Fudenberg, Levine and Maskin [16], E is dense in $\Delta(\mathcal{A})$. Because $g : \Delta(\mathcal{A}) \rightarrow \mathcal{V}^\Delta$ is continuous and surjective, $C \supseteq g(E)$ is dense in \mathcal{V}^Δ . \square

We are now ready to formulate a Nash-threat folk theorem.

Theorem 3.3.3 (Nash-threat folk theorem). *Let \mathcal{V}^0 be the convex hull of $g(\alpha_e)$ and the Pareto-efficient payoff vectors Pareto-dominating $g(\alpha_e)$. Suppose that either g is affine in μ and every Pareto-efficient pure action profile is pairwise identifiable, or*

1. *for every pair of players i and j , there exists at least one profile α_{ij} with pairwise full rank for that pair of players, and*
2. *for every player i , there exists an enforceable action profile $a_i^* \in \mathcal{A}^{(i)}$ that is either pairwise identifiable or satisfies the best response property for player i .*

Then for any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^0$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.

Proof. The statement follows from Proposition 3.1.4 and Lemma 3.6.2 once we show that \mathcal{W} is uniformly decomposable on tangent hyperplanes. Suppose first that all Pareto-efficient pure action profiles a_1, \dots, a_N are pairwise identifiable and that the expected flow payoff is affine in μ . Then the profiles a_1, \dots, a_N are enforceable by Lemma 3.2.3. Since \mathcal{W} is contained in the interior of $\text{conv}(g(\alpha_e), g(a_1), \dots, g(a_N))$, at any point $w \in \partial\mathcal{W}$ there exists an enforceable action profile $\alpha \in \{\alpha_e, a_1, \dots, a_N\}$ such that $g(\alpha)$ is separated from \mathcal{W} by T_w . Suppose first that T_w is regular. Then α and N_w satisfy Conditions 1 or 4 of Lemma 3.1.3.

If T_w is coordinate to the i^{th} axis, then w either maximizes or minimizes player i 's payoff on \mathcal{W} by convexity. If w maximizes player i 's payoff, then a_i^p from Corollary 3.2.4 maximizes g^i over \mathcal{A} and it is pairwise identifiable by assumption, hence (a_i^p, e_i) satisfies Condition 2 of Lemma 3.1.3. If w minimizes player i 's payoff, then $g(\alpha_e)$ is separated from \mathcal{W} by T_w by assumption and $(\alpha_e, -e_i)$ satisfy condition 4 of Lemma 3.1.3.

Suppose now that Conditions 1 and 2 hold instead. In this case, we can actually decompose smooth payoff sets \mathcal{W} in the interior of the slightly larger set $\mathcal{V}^\circ := \{v \in \mathcal{V}^* \mid v^i \geq g^i(\alpha_e)\}$. Denote by $\tilde{a}_1, \dots, \tilde{a}_K$ pure action profiles with extremal payoffs such that \mathcal{V}° is contained in $\text{conv}(g(\alpha_e), g(\tilde{a}_1), \dots, g(\tilde{a}_K))$. For each \tilde{a}_k there exists a mixed action profile α_k with pairwise full rank by Lemma 3.3.2, such that $g(\alpha_k)$ is arbitrarily close to $g(\tilde{a}_k)$. By Lemma 3.3.1, $\alpha_1, \dots, \alpha_K$ are all enforceable and pairwise identifiable. Choose $\alpha_1, \dots, \alpha_K$ such that \mathcal{W} is contained in the interior of $\text{conv}(g(\alpha_e), g(\alpha_1), \dots, g(\alpha_K))$. Therefore, for every $w \in \mathcal{W}$, there exists an enforceable action profile such that its expected flow payoff is strictly separated from \mathcal{W} by T_w .

If T_w is regular or if w minimizes the payoff of a player i on \mathcal{W} , the statement works in the same way as before. If w maximizes the payoff of player i over \mathcal{W} , we use a_i^* instead of a_i^p . It follows from Lemma 3.1.1 that a_i^* is enforceable orthogonal to e_i and (a_i^*, e_i) satisfies Condition 2 or 3 of Lemma 3.1.3 by assumption. \square

To be able to show a minmax version of the folk theorem, we need one more approximation result. Its first part is identical to Lemma 6.3 of Fudenberg, Levine and Maskin [16]. However, we also need that the resulting action profile leads to a locally uniform decomposition as in Lemma 3.1.3.

Lemma 3.3.4. *Suppose that every pure action profile has individual full rank. Then for any $\varepsilon > 0$ and any player $i = 1, \dots, I$, there exists an enforceable action profile α with best response property for player i and $|g^i(\alpha) - \underline{v}^i| < \varepsilon$. Moreover, if either*

1. *every pure action profile is pairwise identifiable, or*
2. *the best response to the minmax profile $\underline{\alpha}_{-i}^i$ is unique,*

then the pair $(\alpha, -e_i)$ satisfies condition 2 or 3 of Lemma 3.1.3 respectively.

Proof. Fix a player i and let $\underline{\alpha}_i^{-i}$ denote a minmax profile against player i . By assumption, (a^i, a^{-i}) has individual full rank for every $a^i \in \mathcal{A}^i$ and every $a^{-i} \in \mathcal{A}^{-i}$. Therefore, similarly as in the proof of Lemma 6.2 in Fudenberg, Levine and Maskin [16], one can find a sequence of profiles $(\underline{\alpha}_{(n)}^{-i})_{n \geq 0}$ converging to $\underline{\alpha}_i^{-i}$ such that $(a^i, \underline{\alpha}_{(n)}^{-i})$ has individual full rank for every $a^i \in \mathcal{A}^i$ and all n .

Let a_n^i be a best response for player i to $\underline{\alpha}_{(n)}^{-i}$. The profiles $(a_n^i, \underline{\alpha}_{(n)}^{-i})$ are enforceable orthogonal to $-e_i$ by Lemmas 3.3.1 and 3.1.1. Let $\underline{a}^i \in \mathcal{A}^i$ be an

accumulation point of $(a_n^i)_{n \geq 0}$ and choose a subsequence $(n_k)_{k \geq 0}$ such that $a_{n_k}^i = \underline{a}^i$ for all $k \in \mathcal{N}$. Observe that \underline{a}^i is also a best response to $\underline{\alpha}_i^{-i}$ because

$$g^i(\underline{a}^i, \underline{\alpha}_i^{-i}) = \lim_{k \rightarrow \infty} g^i(a_{n_k}^i, \underline{\alpha}_{(n_k)}^{-i}) \geq \lim_{k \rightarrow \infty} g^i(\tilde{a}^i, \underline{\alpha}_{(n_k)}^{-i}) = g^i(\tilde{a}^i, \underline{\alpha}_i^{-i}),$$

hence $\underline{a}^i \in \arg \max g^i(\cdot, \underline{\alpha}_i^{-i})$ and $g^i(\underline{a}^i, \underline{\alpha}_i^{-i}) = \underline{v}^i$. Therefore, for any $\varepsilon > 0$ we can find k large enough such that $|g^i(\underline{a}^i, \underline{\alpha}_{(n_k)}^{-i}) - \underline{v}^i| < \varepsilon$.

Under Condition 1, $(\underline{\alpha}_{(n)}^{-i})_{n \geq 0}$ can be chosen in a way that $(a^i, \underline{\alpha}_{(n)}^{-i})$ has pairwise full rank for every $a^i \in \mathcal{A}^i$ and all n . Therefore, $((a^i, \underline{\alpha}_{(n)}^{-i}), -e_i)$ satisfies the second condition of Lemma 3.1.3. Under Condition 2, it follows from multilinearity that there exists a ν large enough such that \underline{a}^i is also a unique best response to $\underline{\alpha}_{(n_k)}^{-i}$. Therefore, Condition 3 of Lemma 3.1.3 is fulfilled for the pair $((\underline{a}^i, \underline{\alpha}_{(n_k)}^{-i}), -e_i)$. \square

Theorem 3.3.5 (Minmax folk theorem). *Suppose that*

1. *every pure action profile has individual full rank,*
2. *for every pair of players i and j , there exists an action profile α_{ij} with pairwise full rank for these players,*
3. *for every player i , there exists an action profile $a_i^* \in \mathcal{A}^{(i)}$ that is either pairwise identifiable or has the unique best response property for player i ,*
4. *for every player i , best responses to the minmax profile $\underline{\alpha}_i^{-i}$ are unique.*

Then for any compact, smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Observe that the second condition is satisfied if there exists at least one pairwise identifiable pure action profile because then it has pairwise full rank by Condition 1. A sufficient condition for the folk theorem to hold is that all pure action profiles have pairwise full rank. Then Conditions 1, 2 and 3 are clearly fulfilled and the fourth condition can be circumvented by Lemma 3.3.4.

Proof of Theorem 3.3.5. Condition 2 and Lemma 3.3.2 imply that any extremal payoff can be approximated by the payoff of a mixed action profile with pairwise full rank. All points w where T_w is regular can thus be dealt with as in the proof of Theorem 3.3.3 under Conditions 1 and 2. In the case where w maximizes player i 's payoff on \mathcal{W} , Condition 1 and Lemmas 3.1.1 and 3.3.1 show that a_i^* is enforceable orthogonal to e_i . Because of Condition 3, (a_i^*, e_i) satisfies either Condition 2 or 3 of Lemma 3.1.3. If w minimizes player i 's payoff on \mathcal{W} , Condition 4 and Lemma 3.3.4 ensure that there exists an enforceable action profile α_i with best response property for player i such that $g(\alpha_i)$ is strictly separated from \mathcal{W} by T_w and $(\alpha_i, -e_i)$ satisfies Condition 3 of Lemma 3.1.3. □

Because pairwise identifiability of action profiles is essential for the folk result to hold, it is worth mentioning a special class of games where this assumption is always satisfied. A game is said to be of a *product structure* if the impacts of players' deviations on the drift are orthogonal, that is, $\text{span } M^i(a) \perp \text{span } M^j(a)$ for all $i \neq j$ and all pure action profiles $a \in \mathcal{A}$. Clearly, this implies pairwise identifiability of all pure action profiles, hence a Nash-threat folk theorem holds for any game, in which Pareto-efficient action profiles are enforceable.

Corollary 3.3.6. *Consider a game with a product structure such that g is an affine function of μ . For any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}_0$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

Since pairwise identifiability and individual full rank are equivalent to having pairwise full rank, we obtain the minmax folk theorem for games with a product structure in the following form.

Corollary 3.3.7. *Suppose that in a game with a product structure every pure action profile has individual full rank. Then for any smooth set $\mathcal{W} \subseteq \text{int } \mathcal{V}^*$, there exists a discount rate $\tilde{r} > 0$ such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for all $r \in (0, \tilde{r})$.*

3.4 Finite variation of equilibrium profiles

Because the constructed equilibrium profiles are concatenations of locally constant strategy profiles at i.i.d. copies of a positive stopping time τ , the resulting equilibrium profiles exhibit finitely many changes on every finite time interval. This is a very desirable feature for implementation because it seems unrealistic that agents can adapt their strategy profiles arbitrarily often. In this section we present an example of such a strategy profile and compare it to the techniques used in Sannikov [37]. Consider a Partnership between two players as in Section 2 of [37], where each player $i = 1, 2$ continuously chooses an effort level in $\mathcal{A}^i = \{0, 1\}$. Players cannot observe whether their partner is working and only see a stochastic output $dX_t^i = A_t^i dt + dZ_t^{A,i}$, where $Z^{A,1}, Z^{A,2}$ are independent Brownian motions under the probability measure Q^A induced by players' strategy profile A . Players split the total output $4(X^1 + X^2)$ and each




Figure 3.4: Shown is a cover of $\partial\mathcal{W}$ (bold black line) into four overlapping sets (solid coloured lines), such that payoffs in a band of width ε around the sets (dashed lines in colour) can be decomposed with respect to the same pure action profile for discount rate $r = 0.1$. The cover of \mathcal{W} is completed by playing the static Nash equilibrium in the interior of \mathcal{W} . Also depicted is $\partial\mathcal{E}(0.1)$ (thin black line) constructed with the techniques in Sannikov [37].

player i pays a cost of effort $3A_t^i$ so that player i receives

$$\int_t^\infty e^{-r(s-t)} (2 dX_t^1 + 2 dX_t^2 - 3A_t^i dt).$$

This leads to an expected flow payoff of $g^i(a) = 2a^{-i} - a^i$. To illustrate that the finite-variation property does not depend on players' ability of mixing, we restrict the example to pure strategies.

Figure 3.4 shows a possible cover for a smooth payoff set \mathcal{W} in the interior of \mathcal{V}^* , such that on each element of the cover, the SDE (2.8) admits a strong solution for discount rate 0.1. To ensure that the stopping times can be chosen strictly positive uniformly on each element of the cover, payoffs in a band of

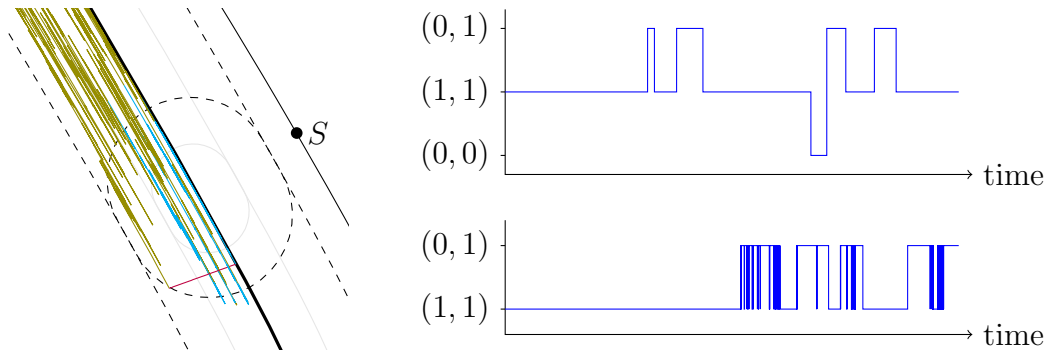


Figure 3.5: The left panel shows the simulation of the continuation value of a PPE in a zoom-in of Figure 3.4. Lines in olive, cyan and red mean that action profiles $(1, 1)$, $(0, 1)$ and $(0, 0)$, respectively, are played. When the continuation value leaves the band around the cover of $\partial\mathcal{W}$, the static Nash equilibrium is played until the boundary of \mathcal{W} is reached. The upper right panel shows the corresponding strategy profile. The lower right panel shows a strategy profile constructed with the techniques in Sannikov [37] with unbounded oscillation when the continuation value crosses the switching point S (left panel).

width ε around the element of the cover need to be decomposable with respect to the same pure action profile. The strategy profile is changed only when the continuation value leaves this band, ensuring that the strategy profile remains constant for a small but positive amount of time.¹⁹

In comparison, the construction of equilibrium profiles in Sannikov [37] works even on the boundary $\partial\mathcal{E}(r)$, where the constructed strategy profiles are constant up to a finite number of “switching points”. However, due to the unbounded variation of Brownian motion, the players switch between action profiles an infinite number of times during a finite time interval when the continuation value crosses a switching point; see also Figure 3.5. As shown

¹⁹The stopping times $(\tau_n)_{n \geq 0}$ constructed in the proof are functionals of the signal such that the continuation value would not escape the band of width ε if it were to start at any point of *any* neighbourhood. For all practical purposes (including this example), one may think of these times as the times $\tilde{\tau}_n$ at which the continuation value leaves the band of width ε around $U_{W_{\tilde{\tau}_{n-1}}}$. Because $\tilde{\tau}_n \geq \tau_n$ a.s. for any n , the concatenation still extends to ∞ .

in Sannikov [37], the chosen action profiles are unique outside these switching points. The equilibrium payoff set $\mathcal{E}(0.1)$ is thus not uniformly decomposable on tangent hyperplanes for $r = 0.1$ as no action profile can be tangentially enforced in a neighbourhood of these switching points.²⁰ This means that the construction of such a cover is possible only in the interior of $\mathcal{E}(0.1)$ and the amount of inefficiency is related to the reaction times of players: the closer $\partial\mathcal{W}$ is to $\partial\mathcal{E}(r)$, the smaller are the overlaps of the elements of the cover and hence the shorter are the times $\tau_{n+1} - \tau_n$ in distribution. While our approach of constructing equilibrium profiles is more general in the sense that it is applicable to any finite number of players, this example shows that it can have advantages even in two-player games: at the cost of only ε of efficiency, it is possible to use equilibrium profiles with locally bounded oscillation.

If players are not restricted to pure strategies, the realizations of their strategies are drawn continuously. Therefore, players switch actions infinitely often on finite time intervals even for constant (but mixed) strategy profiles. However, mixing is done individually for each player, and because of the multilinearity in (2.5), the public signal is not affected by the different realizations of a player's mixed action as long as his/her strategy profile remains constant. The strategies of a player's opponents are therefore not affected by the realizations of his/her mixed strategy, hence a change of actions within a constant strategy profile is a less complicated operation than a change of strategy profile. Moreover, because μ and g are extended to mixed action profiles by multilinearity, continuous-time mixing may be interpreted as a division of effort amongst the pure actions in its support. This is a common formulation

²⁰If $\mathcal{E}(r)$ is smooth, it is uniformly decomposable for any discount rate $\tilde{r} < r$.

in continuous-time games of strategic experimentation; see, for example, the papers by Bolton and Harris [8] as well as Keller and Rady [29].

On such a uniformly decomposable payoff set \mathcal{W} , players adapt their strategies only at stopping times $(\tau_n)_{n \geq 0}$. This leads to the interpretation of a continuously repeated game as a discretely repeated game where the length of the periods is not fixed but random. Indeed, the fact that there exists a strong solution to (2.8) on these intervals of random lengths is consistent with a discrete-time interpretation, where the public signal is a random variable sampled at time τ_n , hence necessarily fixed over the entire period $((\tau_n, \tau_{n+1}])$. In such a discrete-time game with random time intervals any payoff in \mathcal{W} can be attained in equilibrium and its canonical embedding into continuous time is a continuous-time equilibrium profile.

This leads to interesting questions for future research on the connection between equilibria in discrete- and continuous-time models. Indeed, by considering a monotonically increasing sequence of uniformly decomposable payoff sets $(\mathcal{W}_n)_{n \geq 0}$ with $\mathcal{W}_n \rightarrow \mathcal{E}(r)$ and their induced discrete-time games, it may be possible to obtain a sequence of discrete-time games such that equilibrium profiles of the discrete-time games converge to continuous-time equilibria. The convergence of equilibrium strategies remains an open question, and attempts using discrete-time games with fixed time intervals have not been entirely satisfactory. While Staudigl and Steg [42] show the convergence of a suitable sequence of discrete-time games to a continuous-time game with Brownian information, they need to consider weaker notions of convergence for equilibrium profiles. Considering games with random time intervals instead of fixed ones may thus be the key to solving this interesting question.

3.5 Comments on abrupt information

The continuous-monitoring setting is a natural setting to study folk theorems in continuous time: To attain efficiency in the limit, it is necessary that Pareto-efficient action profiles be enforceable without value burning, which restricts incentives to be tangential to the self-generating payoff sets \mathcal{W} . This is the only way in which the continuous information can be used on the boundary of \mathcal{W} anyway. The use of the discontinuous component of the public signal, however, is severely restricted. Because incentives provided through the observation of rare events result in a jump of the continuation value, tangential incentives on the boundary of a self-generating set can be provided only on segments where $\partial\mathcal{W}$ is flat. Indeed, otherwise a tangential jump would necessarily escape \mathcal{W} in contradiction to the fact that $\mathcal{E}(r)$ is self-generating; see also Figure 3.6.

In many ways, the discontinuous component of the public signal has similar features as the public signal in a discrete-time game. There are countably many observations on $[0, \infty)$ and upon the arrival of such an observation, the continuation payoff jumps to its new expected value. One may thus wonder whether the techniques of Fudenberg Levine and Maskin [16] are applicable, that is, whether continuation payoffs can be constructed parallel to the tangent hyperplane, rather than on the tangent hyperplane itself. However, such a construction requires that the arrival time of an event is bounded away from zero so that the expected inward movement due to the drift rate is strictly positive. Indeed, the constructed incentives lie below the tangent hyperplane by an amount that corresponds to the inward drift as illustrated in Figure 3.6.




Figure 3.6: At payoffs where $\partial\mathcal{W}$ is curved, abrupt information cannot be used to provide incentives without value burning. Because the arrival time of an event could be arbitrarily soon, tangential incentives have to be arbitrarily small for the jumps to land in \mathcal{W} .

Finally, we mention the obvious conclusion that the conditions in our folk theorems of Sections 3.2 and 3.3 are sufficient also when, in addition, abrupt information is observed. Indeed, if a strategy profile A is enforced by a process β in the continuous-monitoring game, then $(\beta, 0)$ enforces A in the game including abrupt information. As a result, the set $\mathcal{E}_c(r)$ of equilibrium payoffs with continuous monitoring is contained in $\mathcal{E}(r)$, the equilibrium payoff set when the full information is observed. Therefore, $\mathcal{E}_c(r) \rightarrow \mathcal{V}^*$ implies that also $\mathcal{E}(r) \rightarrow \mathcal{V}^*$. Note that for a fixed discount rate $r > 0$, the addition of abrupt information enlarges $\mathcal{E}(r)$, and hence increases the efficiency of equilibrium payoffs. Sometimes this increase can be quite large precisely because of the additional possibility to burn value as we show in the next chapter. For approximation of the efficient frontier on $\partial\mathcal{V}^*$, however, this impact becomes insignificant because of the restrictions on incentives described above.

3.6 Proof of Proposition 3.1.4

The proof of Proposition 3.1.4 is done in two stages. First, we show in Lemma 3.6.1 that a smooth, uniformly decomposable payoff set is locally self-

generating. In Lemma 3.6.2, we show that a compact, locally self-generating set is self-generating for a sufficiently small discount rate.

Definition 3.6.1. A set $\mathcal{W} \subseteq \mathbb{R}^n$ is called *locally self-generating* if for every point $w \in \mathcal{W}$, there exist an open neighbourhood U_w of w , a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$, an enforceable strategy profile A , a martingale M orthogonal to σZ and $\tilde{r} > 0$ such that for every discount rate $r \in (0, \tilde{r})$ there exists a stopping time $\tau > 0$ such that for all $v \in U_w \cap \mathcal{W}$, there exist W^v with $W_0^v = v$ a.s. and β^v such that $v \mapsto W^v$ and $v \mapsto \beta^v$ are Borel measurable and on $0 \leq t \leq \tau$, the processes W^v, β^v, A, M, Z are related by (2.8), β_t^v enforces A_t a.s. and $W_t^v \in \mathcal{W}$ a.s.

Lemma 3.6.1. *Suppose that a smooth set $\mathcal{W} \subseteq \mathcal{V}^*$ is uniformly decomposable on tangent hyperplanes. Then \mathcal{W} is locally self-generating.*

Proof. Suppose first that w is in the interior of \mathcal{W} . Let $U_w = B_\varepsilon(w)$, where $\varepsilon > 0$ is chosen such that the open ball $B_{2\varepsilon}(w)$ is contained in \mathcal{W} . For a static Nash equilibrium α_e , the constant strategy profile $A \equiv \alpha_e$ is enforced by $\beta \equiv 0$. For any $r > 0$ and any $v \in U_w$, let W^v be a strong solution to

$$dW_t^v := r(W_t^v - g(\alpha_e)) dt$$

with initial condition $W_0^v = v$ a.s. Its solution $W_t^v = v + (e^{rt} - 1)(v - g(\alpha_e))$ is clearly measurable in v . Let $t_0(r) := \log(1 + \varepsilon/(\|w - g(\alpha_e)\| + \varepsilon))/r$, then for any $v \in B_\varepsilon(w)$ it follows that

$$\|W_t^v - v\| \leq (e^{rt} - 1)(\varepsilon + \|w - g(\alpha_e)\|) \leq \varepsilon$$

on $[0, t_0]$. Therefore, in the interior of \mathcal{W} , we can support any discount rate $r > 0$ by choosing $U_w = B_\varepsilon(w)$ and the deterministic time $\tau \equiv t_0(r) > 0$.

For any $w \in \partial\mathcal{W}$, denote by N_w the outward unit normal to $\partial\mathcal{W}$ in w and by T_w the tangent hyperplane to $\partial\mathcal{W}$ in w . By smoothness of \mathcal{W} , both of these are unique and continuous in $w \in \partial\mathcal{W}$. Fix now a payoff $w \in \partial\mathcal{W}$. It will be convenient to work in a coordinate system with origin in w and a basis consisting of an orthonormal basis of T_w and N_w , where we choose the I^{th} coordinate in the direction of N_w . Since $\partial\mathcal{W}$ is a C^2 submanifold, we can locally parametrize it by a twice differentiable function φ . Let $\hat{v} = (v_1, \dots, v_{I-1})$ denote the projection onto the first $I-1$ components so that $\partial\mathcal{W}$ is parametrized by $(\hat{v}, \varphi(\hat{v}))$. By assumption, there exists an enforceable action profile α such that $g(\alpha)$ is strictly separated from \mathcal{W} by T_w . Let β_α be the locally Lipschitz continuous function from Lemma 3.1.3, which assigns to any vector $x \in \mathbb{R}^I$ a matrix β enforcing α orthogonal to x . Choose $\varepsilon > 0$ such that

1. $N_w^\top N_w > 0$ for all $v \in B_{2\varepsilon}(w) \cap \partial\mathcal{W}$.
2. For all $v \in B_{2\varepsilon}(w)$, $\|\nabla\varphi(\hat{v})\| \leq p_1$ and $|\Delta_{ij}\varphi(\hat{v})| \leq p_2$ for $i, j = 1, \dots, I$ and constants $p_1, p_2 > 0$, where $\Delta_{ij}\varphi$ denotes the second partial derivative of φ with respect to \hat{v}_i and \hat{v}_j .
3. There exists a constant B such that on $B_{2\varepsilon}(w)$, $\beta_\alpha(-\nabla\varphi(\hat{\cdot}), 1)$ is Lipschitz continuous and $|(\beta_\alpha \sigma \sigma^\top \beta_\alpha^\top)_{ij}| \leq B$ for $i, j = 1, \dots, I$.
4. $c(\varepsilon) := N_w^\top(g(\alpha) - w) - 2\varepsilon(1 + p_1) - p_1\|g(\alpha) - w\| > 0$.

The first condition makes sure that a local parametrization φ exists with bounded gradient as in Condition 2. Since $\partial\mathcal{W}$ is assumed to be C^2 , the

first two derivatives of φ are continuous, hence locally bounded. In particular, by letting ε small enough we get the first two conditions to hold. For the third condition, observe that φ is continuous with bounded derivative by Condition 2, hence Lipschitz continuous. Since the projection $\hat{\cdot}$ is Lipschitz continuous with Lipschitz constant 1 and the composition of Lipschitz continuous functions is Lipschitz again, the third condition holds in a small neighbourhood of $N_w = (-\nabla\varphi(\hat{w}), 1)$ by Lemma 3.1.3. Finally, $N_w^\top(g(\alpha) - w)$ in Condition 4 is positive by strict separation of $g(\alpha)$ from \mathcal{W} . Because $\nabla\varphi$ is continuous and $\nabla\varphi(\hat{w}) = 0$, p_1 can be made arbitrarily small by choosing a small ε . This implies that $c(\varepsilon) > 0$ for sufficiently small ε .

Fix a stochastic framework $(\Omega, \mathcal{F}, \mathbb{F}, P, Z)$ and an ε satisfying all of the above conditions. Let $U_w := B_\varepsilon(w)$ and $\tilde{r} := 2c(\varepsilon)/((I - 1)^2 p_2 B)$, fix a discount rate $r \leq \tilde{r}$ and set $A \equiv \alpha$. For all $v \in B_\varepsilon(w)$, let W^v denote the strong solution to

$$dW_t^v = r(W_t^v - g(\alpha))dt + r\beta_\alpha(-\nabla\varphi(\hat{W}_t^v), 1)(\sigma dZ_t - \mu(\alpha) dt)$$

on $\llbracket 0, \tau_v \rrbracket$ with $W_0 = v$ a.s., where $\tau_v := \inf\{t > 0 \mid W_t^v \notin B_{2\varepsilon}(w)\}$. Using that $\beta_\alpha \circ (-\nabla\varphi(\hat{\cdot}), 1)$ is uniformly bounded and Lipschitz continuous on $B_{2\varepsilon}(w)$ by Condition 3, a strong solution to this stochastic differential equation exists by Theorem 5.2.1 of Øksendal [34].²¹ Note that $\beta_t := \beta_\alpha(-\nabla\varphi(\hat{W}_t^v), 1)$ is predictable on $\llbracket 0, \tau_v \rrbracket$ as concatenation of a predictable process with a Borel measurable function. Moreover, it enforces A and it is bounded on $\llbracket 0, \tau_v \rrbracket$ by Lemma 3.1.3, hence locally square-integrable.

²¹Uniform boundedness implies the linear growth condition needed for the existence result. The function $f(x) = r(x - g(\alpha))$ is linear, hence Lipschitz continuous and of linear growth.




Figure 3.7: The distance of W_t^v to $\partial\mathcal{W}$ is measured by D_t^v in the direction of N_w for all v in a neighbourhood of w . The continuation value W_t^v remains in \mathcal{W} if and only if $D_t^v \leq 0$. The sensitivity β_α of W^v to Z is chosen orthogonal to $(-\nabla\varphi(\hat{W}_t^v), 1)$, the normal vector to $\partial\mathcal{W}$ at the projection $(\hat{W}_t^v, \varphi(\hat{W}_t^v))$ of W_t^v onto $\partial\mathcal{W}$.

Let $D_t^v := W_t^{v,I} - \varphi(\hat{W}_t^v)$ measure the distance from W_t to $\partial\mathcal{W}$ in the direction of N_w as shown in Figure 3.7. By Itô's formula,

$$\begin{aligned}
dD_t^v &= (-\nabla\varphi(\hat{W}_t^v), 1)^\top dW_t^v - \frac{1}{2} \sum_{i,j=1}^{I-1} \frac{\partial^2\varphi(\hat{W}_t^v)}{\partial x_i \partial x_j} d\langle W^{v,i}, W^{v,j} \rangle_t \\
&= r \left((-\nabla\varphi(\hat{W}_t^v), 1)^\top (W_t^v - g(\alpha)) - \frac{r}{2} \sum_{i,j=1}^{I-1} \frac{\partial^2\varphi(\hat{W}_t^v)}{\partial x_i \partial x_j} (\beta_t \sigma \sigma^\top \beta_t^\top)_{ij} \right) dt \\
&\quad + r (-\nabla\varphi(\hat{W}_t^v), 1)^\top \beta_\alpha (-\nabla\varphi(\hat{W}_t^v), 1) (\sigma dZ_t - \mu(\alpha) dt) \\
&\leq r \left(\frac{r}{2} (n-1)^2 p_2 B - c(\varepsilon) \right) dt,
\end{aligned}$$

where we used that $x^\top \beta_\alpha(x) = 0$ for all $x \in B_{2\varepsilon}(w)$ and that Conditions 2 and 4 imply

$$\begin{aligned}
&((- \nabla\varphi(\hat{W}_t^v), 0) + N_w)^\top (W_t^v - w + w - g(\alpha)) \\
&\leq \| \nabla\varphi(\hat{W}_t^v) \| (2\varepsilon + \|g(\alpha) - w\|) + 2\varepsilon - N_w^\top (g(\alpha) - w) \leq -c(\varepsilon).
\end{aligned}$$

This implies that for any discount rate $r \in (0, \tilde{r})$, D^v is absolutely continuous with $dD_t^v/dt \leq 0$ on $\llbracket 0, \tau_v \rrbracket$, where τ_v depends on r . Since $D_0^v \leq 0$ for all $v \in U_w \cap \mathcal{W}$, it follows that $D_t^v \leq 0$ on $\llbracket 0, \tau_v \rrbracket$. Next, we show that the τ_v are uniformly bounded from below by a stopping time $\tau > 0$. The idea is that this stochastic differential equation is sufficiently nice such that the flow $v \mapsto W^v$ is continuous and thus W^v can be approximated by $W^{\bar{v}}$ for \bar{v} close to v . This leads to a cover of $B_\varepsilon(w)$ with a finite subcover, over which the minimum of stopping times is still positive. Denote $V_t^v := e^{-rt}(W_t^v - v)$ and derive from the product rule that it satisfies

$$dV_t^v = re^{-rt}(v - g(\alpha)) dt + re^{-rt}\beta_\alpha(-\nabla\varphi(v + \widehat{e^{rt}V_t^v}), 1)(\sigma dZ_t - \mu(\alpha) dt).$$

Fix a time horizon $T > 0$ to make e^{rt} bounded and Lipschitz continuous. To apply Theorem V.37 of Protter [36] we write V^v in its integrated form

$$V_t^v = r(1 - e^{-rt})(v - g(\alpha)) + \int_0^t F(V^v)_s(\sigma dZ_s - \mu(\alpha) ds), \quad t \leq T,$$

where $F(V^v)_s = re^{-rs}\beta_\alpha(-\nabla\varphi(v + \widehat{e^{rs}V_s^v}), 1)$. Both the finite variation part and F are Lipschitz,²² hence Theorem V.37 of Protter [36] applies and we deduce that the flow $v \mapsto V^v(\omega)$ is continuous for almost all ω .²³ Define the stopping time $\sigma_v := \inf\{t > 0 \mid e^{rt}\|V_t^v\| > \varepsilon/2\}$, which is strictly positive

²²The result requires that F is functional Lipschitz, which is satisfied for any operator induced by a Lipschitz function; see Protter [36, pg. 251].

²³Here, $V^v(\omega)$ is to be understood as an element of the space \mathcal{D}^I of càdlàg functions from $[0, \infty)$ to \mathbb{R}^I with the topology of uniform convergence on compacts. A compatible metric is given by

$$d(f, g) := \sum_{k=1}^{\infty} \frac{1}{2^k} \left(1 \wedge \sup_{0 \leq s \leq k} \|f(s) - g(s)\| \right);$$

see Protter [36, pg. 220].

by continuity of V^v . For any $v, \bar{v} \in \overline{B_\varepsilon(w)}$, define the auxiliary process $\Lambda_t^{v, \bar{v}} := e^{rt} V_t^v 1_{\llbracket 0, \sigma_{\bar{v}} \rrbracket} + e^{r\sigma_{\bar{v}}} V_{\sigma_{\bar{v}}}^v 1_{(\sigma_{\bar{v}}, \infty)}$. For fixed $\bar{v} \in \overline{B_\varepsilon(w)}$, the map $v \mapsto \Lambda^{v, \bar{v}}$ is continuous for almost all ω . Because $\|\Lambda^{\bar{v}, \bar{v}}\| \leq \varepsilon/2$, there exists an $\varepsilon_{\bar{v}} > 0$ such that $\|\Lambda^{v, \bar{v}}\| < \varepsilon$ for all $v \in B_{\varepsilon_{\bar{v}}}(\bar{v})$. This implies $\|W^v - v\| < \varepsilon$ on $\llbracket 0, \sigma_{\bar{v}} \rrbracket$ for all $v \in B_{\varepsilon_{\bar{v}}}(\bar{v}) \cap B_\varepsilon(w)$ and thus $\tau_v > \sigma_{\bar{v}}$ a.s. Since $\overline{B_\varepsilon(w)}$ is compact there exists a finite subcover $B_{\varepsilon_{\bar{v}_1}}(\bar{v}_1), \dots, B_{\varepsilon_{\bar{v}_n}}(\bar{v}_n)$ of $\overline{B_\varepsilon(w)}$ and thus $\tau := \inf_{k=1, \dots, n} \sigma_{\bar{v}_k}$ is strictly positive. Since $v \mapsto W^v$ is continuous, it is Borel measurable and because β is a continuous functional of W , so is $v \mapsto \beta^v$. \square

Remark 3.6.1. Observe that the existence of a static Nash equilibrium is not necessary to attain payoffs in the interior of \mathcal{W} , and hence Proposition 3.1.4 remains valid if players are restricted to pure strategies. Indeed, the argument is similar to Proposition 9.2.2 of Mailath and Samuelson [31], where payoffs in the interior are attained with the same techniques as on the boundary. For $w \in \text{int } \mathcal{W}$, choose an $\varepsilon > 0$ such that $B_{2\varepsilon}(w) \subseteq \text{int } \mathcal{W}$ and let the role of T_w be taken by any hyperplane H through w . Let N be orthogonal to H and let a be any enforceable action profile such that $g(a) \notin H$ and (a, N) satisfy a condition of Lemma 3.1.3. Proceeding in the same way as on the boundary, we obtain a continuation value W that remains in $B_{2\varepsilon}(w) \subseteq \mathcal{W}$ a.s.

If a uniformly decomposable payoff set is compact, these local solutions can be concatenated to a global solution according to the following lemma.

Lemma 3.6.2. *Let $\mathcal{W} \subseteq \mathbb{R}^I$ be a compact locally self-generating set. Then there exists a discount rate \tilde{r} such that $\mathcal{W} \subseteq \mathcal{E}(r)$ for any $r \in (0, \tilde{r})$.*

Proof. The family of open neighbourhoods $(U_w)_{w \in \mathcal{W}}$ forms an open cover of \mathcal{W} , hence by compactness there exists a finite subcover $(U^k)_{k=1, \dots, N}$. By

making these sets disjoint, we obtain a finite, Borel measurable cover of \mathcal{W} . On each of these (now disjoint) sets U^k , there exists a stochastic framework $(\Omega^k, \mathcal{F}^k, \mathbb{F}^k, P^k, Z^k)$. Define $(\Omega, \mathcal{F}, \mathbb{F}, P)$ as the product space, that is, $\Omega := \Omega^1 \times \dots \times \Omega^N$, $\mathcal{F} := \mathcal{F}^1 \otimes \dots \otimes \mathcal{F}^N$ and similarly for $\mathcal{F}_t, t \geq 0$ and $P(B) := P(B^1) \dots P(B^N)$ for $B = B^1 \times \dots \times B^N \in \mathcal{F}$. Choose any discount rate r smaller than $\tilde{r} = \min(\tilde{r}^1, \dots, \tilde{r}^N) > 0$. Then, dependent on r , for every U_k there exist a strategy profile A^k , a martingale M^k orthogonal to σZ^k and a stopping time $\tau^k > 0$ such that for every $v \in U^k \cap \mathcal{W}$, there exist β^v, W^v measurable in v satisfying the appropriate conditions. Define $\tau(\omega^1, \dots, \omega^N) := \min(\tau^1(\omega^1), \dots, \tau^N(\omega^N), 1)$, which is positive P -a.s.

Fix $v \in \mathcal{W}$ and let κ be the index such that $v \in U^\kappa$. On $\llbracket 0, \tau \rrbracket$ we set $Z = Z^\kappa$, $A = A^\kappa$, $M = M^\kappa$, $W = W^v$ and $\beta = \beta^v$ and we know that they have the desired properties. In particular, $W_\tau \in \mathcal{W}$, hence we can concatenate the solution with the processes related to the neighbourhood U^k in which W_τ falls into. More precisely, let $\tilde{\tau}^1, \dots, \tilde{\tau}^N$ be stopping times such that for every $v \in U^k \cap \mathcal{W}$, there exist solutions to (2.8) starting at time τ , meeting all the necessary properties of local self-generation. Because Brownian motions have independent and stationary increments, $\tilde{\tau}^k - \tau$ can be chosen to have the same distribution as τ^k for every $k = 1, \dots, N$. This implies that $\tilde{\tau} - \tau$ is identically distributed as τ , where $\tilde{\tau} = \min(\tilde{\tau}^1, \dots, \tilde{\tau}^N, 1)$. Moreover, Theorem 32 in Protter [36] implies that W^1, \dots, W^N are strong Markov processes and hence $\tilde{\tau} - \tau$ is independent of \mathcal{F}_τ . On $\llbracket \tau, \tilde{\tau} \rrbracket$, set

$$Z_t(\omega) := Z_{\tau(\omega)}^\kappa(\omega^\kappa) + \sum_{k=1}^N Z_{t-\tau(\omega)}^k(\omega^k) 1_{\{W_{\tau(\omega)} \in U^k\}}$$

and similarly for M . It follows from the strong Markov property of each Z^k that Z is a (P, \mathbb{F}) -Brownian motion. Since τ is bounded, the optional stopping theorem implies that M is a (P, \mathbb{F}) -martingale. Because P is defined as the product measure and each M^k is orthogonal to σZ^k , M is orthogonal to σZ . Further, define the process

$$A_t(\omega) := \sum_{k=1}^N A_{t-\tau(\omega)}^k(\omega^k) 1_{\{W_{\tau(\omega)} \in U^k\}},$$

which is $\Delta(\mathcal{A})$ -valued and \mathbb{F} -predictable and hence it is a valid behaviour strategy profile. Similarly, $\beta_t(\omega) := \beta_{t-\tau(\omega)}^{W_{\tau(\omega)}}(\omega)$ and $W_t(\omega) := W_{t-\tau(\omega)}^{W_{\tau(\omega)}}(\omega)$ are \mathbb{F} -predictable since $v \mapsto \beta^v$ and $v \mapsto W^v$ are measurable by assumption. Finally, it follows by construction that β enforces A on $\llbracket 0, \tilde{\tau} \rrbracket$ and all the processes are related by (2.8).

An iteration of this procedure leads to a sequence of stopping times $(\tau_n)_{n \geq 0}$ and solutions $(W^n, A^n, \beta^n, Z^n, M^n)_{n \geq 0}$ to (2.8) on $(\tau_n, \tau_{n+1}]$ such that $\tau_{n+1} - \tau_n$ are independent and identically distributed as τ . By the following Lemma 3.6.3, τ_n diverges to ∞ a.s., hence a countable concatenation of $(W^n, A^n, \beta^n, Z^n, M^n)$ yields a solution to (2.8) on $[0, \infty)$ attaining v . Since v was arbitrary, \mathcal{W} is self-generating for discount rate r and hence $\mathcal{W} \subseteq \mathcal{E}(r)$ by Proposition 2.4.3. The statement follows since $r \in (0, \tilde{r})$ was arbitrary. \square

For a given subcover of \mathcal{W} , we can fix the probability space at the beginning of the concatenation. Note, however, that this does not imply that the resulting global solution is a strong solution to (2.8) because we cannot find a fixed Brownian motion at the beginning. Indeed, we need to concatenate it at $(\tau_n)_{n \geq 0}$, based on which neighbourhood W_{τ_n} fell into.

Lemma 3.6.3. *Let $(\tau_n)_{n \geq 0}$ be a sequence of random variables with $\tau_0 = 0$ such that $\tau_{n+1} - \tau_n$ are strictly positive and i.i.d. Then $\tau_n \rightarrow \infty$ a.s.*

Proof. Let $\tilde{\tau}_n := \tau_n - \tau_{n-1}$ for $n \geq 1$. Then $\tilde{\tau}_n$ are i.i.d. and $\tau_n = \sum_{k=1}^n \tilde{\tau}_k$. Therefore, the strong law of large numbers implies that

$$\frac{1}{n} \tau_n = \frac{1}{n} \sum_{k=1}^n \tilde{\tau}_k \rightarrow \mathbb{E}[\tau_1] \text{ a.s.}$$

That is, for all $\varepsilon > 0$ there exists n_0 such that for any $n \geq n_0$ it holds that $|\tau_n/n - \mathbb{E}[\tau_1]| < \varepsilon$ a.s. Letting $\varepsilon = \mathbb{E}[\tau_1]/2$, we obtain $\tau_n > n \mathbb{E}[\tau_1]/2$ a.s. for all $n \geq n_0$. This lower bound, and hence also τ_n , diverges to ∞ a.s. since $\mathbb{E}[\tau_1] > 0$ because $\tau_1 = \tau_1 - \tau_0 > 0$ a.s. □

Chapter 4

Explicit Characterization of $\mathcal{E}(r)$ in Two-Player Games

One of the main advantages of continuous-time models is that they often give rise to explicit results that do not exist in discrete time. We have seen a flavour of these results in Section 2.5 and we will see more of them in this chapter. In his introduction of continuous-time repeated games with imperfect information, Sannikov [37] shows that a continuous-time formulation of repeated games makes it possible to relate the continuation value of equilibrium strategy profiles to the boundary of the equilibrium payoff set $\mathcal{E}(r)$. For a class of two-player games with Brownian information, he obtains a differential equation characterizing the boundary of the equilibrium payoff set $\mathcal{E}(r)$ for any discount rate $r > 0$. Such an explicit description of the equilibrium payoff set also for impatient players is a result without analogue in discrete time. In many situations, however, the assumption that the observed information is entirely Brownian is not suitable.

Consider, for example, a joint venture between two parties, where each party secretly chooses their effort levels. Each party observes only the total revenue of the partnership, which depends on the chosen effort levels and on the demand for the produced good. Demand is stochastic so that the total revenue conveys only imperfect information about players' efforts. The demand fluctuates with normal market behaviour and from time to time, a demand shock occurs due to a scandal involving one of the parties. The total revenue consists of the instantaneous revenue due to production and sale, and it contains infrequent but large expenses due to settlements of lost lawsuits or payments for equipment upkeep. In this example, the assumption of Brownian information is suitable only for the instantaneous revenue due to normal market fluctuations. The arrival of infrequent events — such as demand shocks, lawsuits and equipment failure — are better modelled by the jump times of Poisson processes.

In this chapter, we extend the class of games for which an explicit characterization of $\mathcal{E}(r)$ is known to two-player games involving both continuous and abrupt information as in the general model of Chapter 2. We show that the presence of abrupt information may have a drastic impact on the equilibrium payoff set, even if the infrequent events are of purely informational nature and do not themselves affect players' payoffs. This is due to the fundamentally different ways in which continuous and abrupt information are used to provide incentives at the boundary of $\mathcal{E}(r)$: while the continuous information can be used only to transfer value tangentially amongst players, the arrival of rare events can be used to both transfer and destroy value; see also the left panel of Figure 4.1. This fundamental difference has first been observed in Sannikov and




Figure 4.1: While unbounded amounts of value can be transferred or destroyed upon the arrival of an infrequent event in Sannikov and Skrzypacz [39] (right panel), this is not the case in our setting (left panel) because the jumps have to land in the bounded set $\mathcal{E}(r)$. In contrast, the continuous information allows unbounded transfers of value, but the transfers have to be tangential (left panel).

Skrzypacz [39] in their treatment of discrete-time games with frequent actions. In [39], these informational restrictions are used to establish a payoff bound \mathcal{M} for discrete-time games when the length of the time period approaches zero. By working directly in a continuous-time setting, we are able to relate these informational restrictions to the boundary $\partial\mathcal{E}(r)$ and find a precise description of $\mathcal{E}(r)$ for any discount rate $r > 0$. Our main result is thus a generalization of Theorem 2 in Sannikov [37] to a more general information structure that includes abrupt information, as well as an improvement from an asymptotic payoff bound in Sannikov and Skrzypacz [39] to a characterization of $\mathcal{E}(r)$.

In terms of provided incentives, this chapter differs from [39] by the fact that only bounded amounts of value can be transferred or destroyed upon the arrival of a rare event. This is not a difference in the underlying model, but rather in the result that we prove: Instead of characterizing $\mathcal{E}(r)$ explicitly, Sannikov and Skrzypacz [39] construct a payoff bound \mathcal{M} as intersection of dominating half-spaces $H(N)$, for which value is destroyed relative to the

direction N . For our characterization, it is necessary that the jumps remain in the bounded set $\mathcal{E}(r)$; see also Figure 4.1. Because the jumps are of size $r\delta(y)$ for $y \in Y$ as seen in (2.8), this means that only bounded amounts of value can be transferred or destroyed. As the discount rate r increases and players get more impatient, fewer incentives can be provided through the observation of rare events — a feature that does not arise in [39].

The main idea behind the proof is similar to Sannikov [37]: Because $\mathcal{E}(r)$ is self-generating, the continuation value of equilibrium strategy profiles can never leave $\mathcal{E}(r)$. This restricts the possible incentives that can be provided at the boundary of $\mathcal{E}(r)$, which, in turn, is related to the curvature of the set. In our model, these restrictions at $w \in \partial\mathcal{E}(r)$ with normal vector N_w are:

1. Inward-pointing drift: $N_w^\top (g(a) + \delta\lambda(a) - w) \geq 0$,
2. Tangential volatility: $N_w^\top \beta = 0$,
3. Jumps within the set: $w + r\delta(y) \in \mathcal{E}(r)$ for every $y \in Y$.

Sannikov [37] shows that when information arrives continuously only, the boundary of the equilibrium payoff set is explicitly characterized by an ordinary differential equation using Restrictions 1 and 2. Such an explicit form of the differential equation arises because these restrictions are local restrictions on the use of information and depend on the geometry of $\mathcal{E}(r)$ only through the normal vector N_w at w . In the general setting involving abrupt information, restriction 3 creates a non-trivial fixed point problem: in order to solve a differential equation involving Restriction 3 at some point $w \in \partial\mathcal{E}(r)$, one would have to know its entire solution $\partial\mathcal{E}(r)$ already. We solve this fixed

point problem with an iterative procedure over the arrival of rare events, where we relax the condition on the jumps to land in some fixed payoff set \mathcal{W} . The resulting set $\mathcal{B}(r, \mathcal{W})$ is self-generating up to the arrival of the first event, at which point the continuation value jumps to \mathcal{W} . This is a continuous-time analogue of the standard set operator \mathcal{B} in Abreu, Pearce and Stacchetti [2]. Similarly to [2], a successive application of \mathcal{B} to \mathcal{V}^* converges to $\mathcal{E}(r)$. Contrary to its discrete counterpart, however, $\mathcal{B}(r, \mathcal{W})$ can be explicitly characterized through a differential equation describing its boundary.

Our main result generalizes Theorem 2 in Sannikov [37] along two directions. Not only do we allow for a more general information structure, but we also require less stringent assumptions on the game primitives. In particular, our characterization of $\mathcal{E}(r)$ is new also for a continuous public signal if the signal is one-dimensional. This is a non-trivial extension that contains important games such as Cournot duopolies in a homogeneous good, where only the market price is observed, or partnership games, where only the joint output/total revenue is observed. The weakening of the assumptions and the introduction of jumps lead to a loss of regularity of the equilibrium payoff set. Nevertheless, we are able to show that the boundary is locally twice differentiable except at finitely many points, which is important for the computation of $\mathcal{E}(r)$ in specific games with numerical procedures.

The remainder of the chapter is organized as follows. We motivate and state our main result in Sections 4.1 and 4.2. In Section 4.3, we discuss and interpret our result by computing $\mathcal{E}(r)$ for two specific examples for which the characterization of the equilibrium payoff set is new: the climate agreement example of Section 2.2 and a Cournot duopoly, where the publicly observable

market price is one-dimensional. In Section 4.4, we formally introduce the set $\mathcal{B}(r, \mathcal{W})$ and establish the continuous-time analogue of the algorithm in Abreu, Pearce and Stacchetti [2] for the approximation of $\mathcal{E}(r)$. We discuss our main assumption in Section 4.5 and elaborate on alternative conditions that may be met for our main result to apply. We present a natural candidate for $\mathcal{B}(r, \mathcal{W})$ and show how to verify this candidate. A description of how to implement the numerical solution of our main result is presented in Section 4.6, together with several examples. Section 4.7 proves our main result and Section 4.8 contains the verification argument for our result in Section 4.5.2. Finally, Appendix 4.A contains minor results of mainly technical nature that are isolated to improve the flow of the main argument in Section 4.7.

4.1 Motivation

The arrival of rare events carries many similarities to the public signal in discrete-time games. Indeed, there exists a canonical embedding of the abrupt information into a discrete setting with periods given by $((\sigma_{n-1}, \sigma_n])$ for σ_n indicating the n^{th} jump amongst $(J^y)_{y \in Y}$. At time σ_n , the signal y is observed if and only if an event y happens at that time. Because Y is finite in our setting, the discrete information does not satisfy a bang-bang result and the jump may go into the interior of $\mathcal{E}(r)$. However, since two or more jumps of independent Poisson processes happen at the same time with probability 0, it is always possible to use public randomization between events to jump back to the boundary. It is therefore sufficient to check Restriction 3 above only at the boundary of the set, which makes a characterization via $\partial\mathcal{E}(r)$ possible.

On each of the sets $([\sigma_n, \sigma_{n+1}])$, we construct equilibrium profiles on $\partial\mathcal{E}(r)$ that remain in $\mathcal{E}(r)$ until the arrival of the next event. Restrictions 1–3 above thus motivate the following definition of restricted-enforceable incentives that have to be satisfied on the boundaries of $\mathcal{E}(r)$ and $\mathcal{B}(r, \mathcal{W})$, respectively.

Definition 4.1.1. An action profile $a \in \mathcal{A}$ is *restricted-enforceable* for any w, N, r and $\mathcal{W} \subseteq \mathbb{R}^2$ if it is enforced by $(T\phi, \delta)$ with $N^\top (g(a) + \delta\lambda(a) - w) \geq 0$ as well as $w + r\delta(y) \in \mathcal{W}$ and $N^\top \delta(y) \leq 0$ for every $y \in Y$, where ϕ is a row vector in \mathbb{R}^{d_c} and $T \in S^1$ is orthogonal to N . Let $\Xi_a(r, N, w, \mathcal{W})$ denote the set of all pairs (ϕ, δ) that restricted-enforce a .

Remark 4.1.1. The condition $N^\top \delta(y) \leq 0$ means that upon the arrival of event $y \in Y$, value can only be transferred tangentially or burned relative to the direction N . For the equilibrium payoff set $\mathcal{E}(r)$, these are the only ways how the information of events can be used because $\mathcal{E}(r)$ is convex. Imposing this condition also for the use of incentives on $\partial\mathcal{B}(r, \mathcal{W})$ both simplifies the characterization of the boundary of $\mathcal{B}(r, \mathcal{W})$ and improves the speed of convergence of the algorithm to $\mathcal{E}(r)$.

In Chapter 3, we have seen that any payoff on the boundary of a compact, smooth and self-generating payoff set \mathcal{W} can be attained by an enforceable strategy profile whose continuation value does not escape \mathcal{W} . In this section we will motivate that maximality of $\mathcal{E}(r)$ implies that a solution to (2.8) attaining $w \in \partial\mathcal{E}(r) \cap \mathcal{E}(r)$ cannot fall into the interior of $\mathcal{E}(r)$ at payoffs w where $\partial\mathcal{E}(r)$ is continuously differentiable. Let $\mathcal{C} \subseteq \partial\mathcal{E}(r) \cap \mathcal{E}(r)$ be continuously differentiable, fix w_0 in the relative interior of \mathcal{C} and let $\hat{\cdot}$ denote the projection onto \mathcal{C} in the direction of $N = N_{w_0}$. Because payoffs in \mathcal{C} are attained by a PPE,




Figure 4.2: The distance of W to \mathcal{C} in the direction N is measured by D . The continuation value remains on \mathcal{C} if and only if \mathcal{C} is a solution to (4.2).

there has to exist a restricted-enforceable action profile at any $w \in \mathcal{C}$, that is, $\Xi_a(w, N_w, r, \mathcal{E}(r)) \neq \emptyset$ for at least one $a \in \mathcal{A}$. Let $a^* : \mathcal{C} \mapsto \mathcal{A}$, $\phi^* : \mathcal{C} \mapsto \mathbb{R}^{I \times d_c}$ and $\delta^* : \mathcal{C} \mapsto \mathbb{R}^{I \times m}$ be a selection with $(\phi^*(w), \delta^*(w)) \in \Xi_{a^*(w)}(w, N_w, r, \mathcal{E}(r))$ for every $w \in \mathcal{C}$. Consider a solution $(W, A, \beta, \delta, M, Z, (J^y)_{y \in Y})$ to (2.8) starting at $w_0 \in \mathcal{C}$ with $A = a^*(\hat{W})$, $\beta = T_{\hat{W}} \phi^*(\hat{W})$, $\delta = \delta^*(\hat{W})$ and $M \equiv 0$. Let $D = N^\top W - \hat{W}$ be a measure of distance of W to \mathcal{C} as illustrated in Figure 4.2. Itô's formula implies that

$$\begin{aligned} dD_t &= \left(r N_{\hat{W}_t}^\top (W_t - g(A_t) - \delta_t \lambda(A_t)) + \frac{r^2}{2} \kappa(\hat{W}_t) |\phi_t|^2 \right) dt \\ &\quad + r N_{\hat{W}_t}^\top \beta_t (dZ_t - \mu(A_t) dt) + r \sum_{y \in Y} N_{\hat{W}_t}^\top \delta(y) dJ_t^y. \end{aligned} \quad (4.1)$$

Because $N_{\hat{W}_t}^\top \beta_t = 0$, on each interval $[\sigma_n, \sigma_{n+1})$, equation (4.1) reduces to

$$dD_t = \left(r N_{\hat{W}_t}^\top (W_t - g(A_t) - \delta_t \lambda(A_t)) + \frac{r^2}{2} \kappa(\hat{W}_t) |\phi_t|^2 \right) dt.$$

Therefore, $D_t \equiv 0$ and the continuation value remains on \mathcal{C} if and only if at every $w \in \mathcal{C}$, the curvature of \mathcal{C} is given by

$$\kappa(w) = \frac{2N_w^\top \left(g(a^*(w)) + \delta^*(w) \lambda(a^*(w)) - w \right)}{r \|\phi^*(w)\|^2}. \quad (4.2)$$




Figure 4.3: If a solution \mathcal{C} to (4.2) falls into the interior of $\mathcal{E}(r)$ in a neighbourhood of $w_0 \in \partial\mathcal{E}(r)$, continuity in initial conditions implies that a solution \mathcal{C}' to (4.2) starting at $w_0 + \varepsilon N_{w_0}$ has end points v_L, v_R in $\mathcal{E}(r)$. Since any $w \in \mathcal{C}'$ is attainable by a solution W to (2.8) that remains on \mathcal{C}' until $\tau := \sigma_1 \wedge \inf\{t \geq 0 \mid W \in \{v_L, v_R\}\}$, a concatenation at time τ with a PPE attaining $W_\tau \in \mathcal{E}(r)$ is itself a PPE.

Note that even though this argument is motivated with $\mathcal{C} \subseteq \partial\mathcal{E}(r) \cap \mathcal{E}(r)$, it works for any curve \mathcal{C} satisfying $\bigcup_{a \in \mathcal{A}} \Xi_a(w, N_w, r, \mathcal{E}(r)) \neq \emptyset$ at every $w \in \mathcal{C}$. Thus, more generally, any payoff w on a curve \mathcal{C} with curvature given by (4.2) is attainable by an enforceable strategy profile such that its continuation value remains on \mathcal{C} until either an endpoint of \mathcal{C} is reached or a rare event occurs, at which point W jumps to $\mathcal{E}(r)$.

Now suppose that there exist a^*, ϕ^*, δ^* such that the solution \mathcal{C} to (4.2) starting at $w_0 \in \partial\mathcal{E}(r)$ falls strictly into the interior of $\mathcal{E}(r)$. Then, by moving initial conditions a little bit outside of $\mathcal{E}(r)$, one can construct a self-generating set that is larger than $\mathcal{E}(r)$ as indicated in Figure 4.3. This stands in contradiction to Proposition 2.4.3, which states that $\mathcal{E}(r)$ is the largest bounded self-generating payoff set. Observe that this enlargement argument requires continuity of (4.2) in initial conditions, which we establish in Section 4.7.2.

In summary, the curvature of $\partial\mathcal{E}(r)$ at $w \in \partial\mathcal{E}(r)$ is given by taking the maximum in (4.2) over all action profiles and all restricted-enforceable incentives available at $w \in \partial\mathcal{E}(r)$, where $\partial\mathcal{E}(r)$ is continuously differentiable. At corners of $\mathcal{E}(r)$, the continuous information cannot be used to provide incen-

tives because Brownian motion has infinite variation. In Sannikov [37], corners thus have to be contained in \mathcal{V}^N . In our setting, this does not have to be true a priori as the discontinuous component of the public signal has finite variation, and hence non-trivial incentives can be provided through the observation of rare events. The characterization of $\partial\mathcal{E}(r)$ is thus completed by describing all payoffs, where incentives can be provided through the abrupt information exclusively, that is, by characterizing the set

$$\mathcal{G}(r) := \left\{ w \in \partial\mathcal{E}(r) \left| \begin{array}{l} \exists (a, N, \delta) \text{ with } (0, \delta) \in \Xi_a(w, N, r, \mathcal{E}(r)), \\ \text{where } N^\top(w - \tilde{w}) \geq 0 \text{ for all } \tilde{w} \in \mathcal{E}(r) \end{array} \right. \right\}.$$

4.2 Main result

For any action profile $a \in \mathcal{A}$ and any player $i = 1, 2$, let $M^i(a)$ and $\Lambda^i(a)$ denote the matrices containing column vectors $\mu(\tilde{a}^i, a^{-i}) - \mu(a)$ and $\lambda(\tilde{a}^i, a^{-i}) - \lambda(a)$, respectively. $M^i(a)$ and $\Lambda^i(a)$ contain the informational changes induced by player i 's deviations from a^i to any other action. We make the following two assumptions, which we motivate in Section 4.5 after the introduction of the algorithm that approximates $\mathcal{E}(r)$.

Assumption 4.2.1. Suppose that $\text{span } \Lambda^i(a)^\top \subseteq \text{span } M^i(a)^\top Q^{-i}(a)^\top$, where $Q^i(a)$ is any matrix whose row vectors form a basis of $\ker M^i(a)$.

Assumption 4.2.1 is the conjunction of an identifiability assumption together with the assumption that the information driving the continuous component of the public signal dominates the information underlying the rare events. Indeed, if the continuous component of the public signal makes all action profiles pairwise identifiable in the sense of Chapter 3, then there exist

matrices Q^1, Q^2 isolating incentives for players 1 and 2 by Lemma 1 of Sannikov [37]. For such a game, Assumption 4.2.1 reduces to $\text{span } \Lambda^i(a) \subseteq \text{span } M^i(a)$ for $i = 1, 2$, that is, any change of incentives through the rare events can be compensated by incentives arising from the continuous information.

Assumption 4.2.2. Suppose that

- (i) best responses are unique, and
- (ii) if $a \in \mathcal{A}$ is enforced by (β, δ) with $N^\top \beta = 0$ and $N^\top \delta(y) \leq 0$ for a coordinate direction $N = \pm e_i$, then (\hat{a}^i, a^{-i}) is enforceable, where \hat{a}^i is player i 's best response to a^{-i} .

Remark 4.2.1. Assumptions 4.2.1 and 4.2.2 are generalizations of the assumptions in Sannikov [37] to our framework. Note, however, that our assumptions are less stringent when $Y = \emptyset$, that is, no rare events are observed. Indeed, in that case Assumption 4.2.1 is always fulfilled and Assumption 4.2.2 reduces to Assumption 4.2.2.(i), corresponding to Assumption 2.(i) in Sannikov [37].

Theorem 4.2.1. *Under Assumptions 4.2.1 and 4.2.2, $\mathcal{E}(r)$ is the largest closed subset of \mathcal{V}^* such that $\partial \mathcal{E}(r) \setminus \mathcal{G}(r)$ is continuously differentiable with curvature at almost every payoff w given by*

$$\kappa(w) = \max_{a \in \mathcal{A}} \max_{(\phi, \delta) \in \Xi_a(w, r, N_w, \mathcal{E}(r))} \frac{2N_w^\top (g(a) + \delta\lambda(a) - w)}{r \|\phi\|^2}, \quad (4.3)$$

where we set $\kappa(w) = 0$ if the maxima are taken over empty sets. Moreover, $\mathcal{G}(r)$ consists of straight line segments and isolated points only, and all corners of $\mathcal{E}(r)$ are contained in the set of static Nash payoffs.

Compared to games where information arrives continuously only, the observation of rare events enlarges the equilibrium payoff set $\mathcal{E}(r)$. Indeed, if incentives provided through the continuous information exclusively are sufficient to enforce a strategy profile, these incentives are still sufficient when additional events are observed. At points on the boundary, where incentives are provided through both continuous and discontinuous information, the optimality equation looks similar to that in Sannikov [37]. However, the possible tradeoffs between incentives provided through the continuous and abrupt information increases efficiency and enlarges $\mathcal{E}(r)$: The numerator of (4.3) is a measure for the inefficiency, that is, the amount that $\partial\mathcal{E}(r)$ is below the stage game payoff $g(a)$. Even though value burning can only increase this amount of inefficiency, the tradeoff between incentives may significantly reduce the amount of tangential volatility required to provide sufficient incentives, leading to an increase in overall efficiency.

Theorem 4.2.1 characterizes $\mathcal{E}(r)$ as the largest fixed point of the differential equation (4.3). As we have motivated in Section 4.1, we will construct strategy profiles, whose continuation value remains on curves with curvature given by the optimality equation. To ensure that the maxima are attained and these strategy profiles indeed exhibit the desired properties, we need that $\Xi_a(w, N_w, r, \mathcal{E}(r))$ is compact, which is true if and only if $\mathcal{E}(r)$ is compact. Because this is not a priori clear, the algorithm presented in Section 4.4 is used also for the proof of Theorem 4.2.1 and not only for the computation of $\mathcal{E}(r)$ in specific games. By choosing the compact set \mathcal{V}^* as starting point of the algorithm and choosing sufficient conditions to ensure that $\mathcal{B}(r, \mathcal{W})$ preserves compactness, the above argument goes through.

Even though the curvature is characterized only at almost every payoff $\partial\mathcal{E}(r) \setminus \mathcal{G}(r)$, a solution is unique with the additional requirement that it be continuously differentiable. This implies that $\partial\mathcal{E}(r)$ is twice continuously differentiable almost everywhere, which is important for the numerical solution of (4.3) as numerical procedures rely on discretizations. We will elaborate on the numerical implementation in Section 4.6.

4.3 Interpretation of results

4.3.1 Climate agreement

Consider the climate agreement example of Section 2.2, where normalize the volatility of the public signal as explained in Remark 2.1.2. When only the continuous components are observed, i.e., the industrial production indices and the atmospheric greenhouse gas concentration, the boundary is characterized by Theorem 2 in Sannikov [37]. When, in addition, economic and political events are observed, the equilibrium payoff set is characterized by our main result, Theorem 4.2.1. Figure 4.4 shows and compares the computed payoff sets for different discount rates $r = 0.1$ and $r = 0.2$.

We see that self-regulation of the two countries is more efficient with the observation of the political and economic events. This is not surprising since the additional information makes it more difficult for countries to cheat on the agreement undetectedly. What may be surprising, however, is the amount of gained efficiency. As shown in the left panel of Figure 4.4, $\mathcal{E}(0.1)$ extends a lot closer to the efficient frontier than its continuous counterpart, and this,




Figure 4.4: Comparison of $\mathcal{E}(r)$ to the set $\mathcal{E}_c(r)$ of all equilibrium payoffs when only the continuous information is observed. The climate agreement is honoured by both countries while the continuation value remains in the region of $\partial\mathcal{E}(r)$ that is drawn with solid lines. This region becomes larger as more information is available.

even though the necessary value burnt at payoff A is minimal. This effect is even more pronounced for discount rate 0.2 in the right panel of Figure 4.4. In that case, no strategy profile is enforceable without value burning with the exception of the static Nash profile. Nevertheless, non-trivial equilibria exist and their associated payoffs can be computed with our Theorem 4.2.1.

The strength of our result lies in its quantitative nature. By comparing the gained efficiency, it is possible to precisely quantify the value of abrupt information, which may have important policy implications. A mechanism designer, for example, may compare the value of observing certain events to the cost of reporting them. If the gained efficiency outweighs the cost of reporting, requiring participants to report these events will increase overall efficiency. In this example of the climate agreement, these considerations should be taken into account when drafting the agreement as the efficient frontier corresponds to payoffs, where the climate agreement is enforced; see Figure 4.4.




Figure 4.5: The sum of payoffs in every possible equilibrium dominates the static Nash payoff, even though collusion of the firms in the monopoly supply $(4, 4)$ is impossible as shown by Sannikov and Skrzypacz [38]. Indeed, the right panel shows the optimal action profiles on the boundary $\partial\mathcal{E}(r)$: since $(4, 4)$ is not enforceable on the negative diagonal, collusion of the firms is impossible.

4.3.2 Cournot duopoly

In this section we present the example of a Cournot duopoly in a single homogeneous good, which shows that even in the absence of jumps, our equilibrium characterization is applicable to a wider class of games than previously known from Sannikov [37]. At any point in time, firm $i = 1, 2$ chooses its individual supply rate from the sets $\mathcal{A}^i = \{0, 1, 2, 3, 4, 5\}$ and the demand for the product is stochastic. Since the products from the two firms are indistinguishable, their market prices coincide and depend on the firms' actions only through the total supply. Under strategy profile (A^1, A^2) , it is given by

$$dP_t = (20 - 2(A_t^1 + A_t^2)) dt + dZ_t^A,$$

where Z^A is a one-dimensional Q^A -Brownian motion. We suppose that the production costs are linear in the supply rate and given by $c_i(A_t^i) = 3A_t^i$ so

that the discounted future payoff of firm i at time t equals

$$\int_t^\infty r e^{-r(s-t)} (A_s^i dP_s - 3A_s^i ds).$$

In this game, the public information is one-dimensional, and hence the pairwise-identifiability assumption in Sannikov [37] fails. Nevertheless, one can calculate $\partial\mathcal{E}(r)$ as shown in Figure 4.5, based on our Theorem 4.2.1.

4.4 Algorithm

In this section we present a continuous-time analogue of the algorithm in Abreu, Pearce and Stacchetti [2] and show how it can be used to solve the aforementioned fixed point problem. We begin by defining the set-valued operator $\mathcal{B}(r, \mathcal{W})$ in our setting.

Definition 4.4.1. For $\mathcal{W} \subseteq \mathbb{R}^2$, let $\mathcal{B}(r, \mathcal{W})$ denote the largest bounded payoff set that is self-generating up to the arrival of the first rare event such that the jump of the continuation value lands in \mathcal{W} and is directed inwards at $\partial\mathcal{B}(r, \mathcal{W})$.

What stands out in comparison to discrete time is the condition that the jumps are directed inwards at the boundary of $\mathcal{B}(r, \mathcal{W})$. This is related to the discussion in Remark 4.1.1. Because we do not know a priori whether $\mathcal{B}(r, \mathcal{W})$ is closed or not, a strategic definition of $\mathcal{B}(r, \mathcal{W})$ is rather technical, and hence deferred to Section 4.7.1. The following lemma formalizes the algorithm that converges to $\mathcal{E}(r)$.

Lemma 4.4.1. *Let $\mathcal{W}_0 = \mathcal{V}^*$ and $\mathcal{W}_n = \mathcal{B}(r, \mathcal{W}_{n-1})$ for $n \geq 1$. Then $(\mathcal{W}_n)_{n \geq 0}$ is decreasing in the set-inclusion sense with $\bigcap_{n \geq 0} \mathcal{W}_n = \mathcal{E}(r)$.*

Unlike its discrete-time counterpart, the boundary of the resulting set at each step of the iteration is characterized by a differential equation similarly to $\mathcal{E}(r)$. Since the condition on the incentives provided through jumps is fixed, such a characterization of $\mathcal{B}(r, \mathcal{W})$ is explicit. Let

$$\mathcal{G}(r, \mathcal{W}) := \left\{ w \in \partial\mathcal{B}(r, \mathcal{W}) \left| \begin{array}{l} \exists (a, \delta) \text{ with } (0, \delta) \in \Xi_a(w, N, r, \mathcal{W}) \\ \text{for all normals } N \text{ to } \partial\mathcal{B}(r, \mathcal{W}) \text{ at } w \end{array} \right. \right\}$$

denote the set of all payoffs $w \in \partial\mathcal{B}(r, \mathcal{W})$ where incentives can be provided through the observation of rare events only. We say that such a pair (a, δ) *decomposes* w , and, in turn, that w is *decomposable*. We obtain the following characterization of $\partial\mathcal{B}(r, \mathcal{W})$.

Proposition 4.4.2. *Suppose that Assumptions 4.2.1 and 4.2.2 hold and let $\mathcal{W} \subseteq \mathcal{V}^*$ be a compact and convex set with non-empty interior. Then $\mathcal{B}(r, \mathcal{W})$ is the largest closed subset of \mathcal{V}^* such that $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ is continuously differentiable with curvature at almost every payoff w given by*

$$\kappa(w) = \max_{a \in \mathcal{A}} \max_{(\phi, \delta) \in \Xi_a(w, r, N_w, \mathcal{W})} \frac{2N_w^\top (g(a) + \delta\lambda(a) - w)}{r\|\phi\|^2}, \quad (4.4)$$

where we set $\kappa(w) = 0$ if the maxima are taken over empty sets. Moreover, $\mathcal{G}(r, \mathcal{W})$ consists of straight line segments and isolated points only, and all corners of $\mathcal{B}(r, \mathcal{W})$ are contained in the set of static Nash payoffs. If a straight line segment is decomposed by some action profile a , its infinite continuation goes through $g(a)$.

Observe that Theorem 4.2.1 follows from Lemma 4.4.1 and Proposition 4.4.2. Indeed, because $\mathcal{B}(r, \mathcal{W})$ preserves compactness, Lemma 4.4.1 shows that $\mathcal{E}(r)$

is compact. Therefore, an application of Proposition 4.4.2 for $\mathcal{W} = \mathcal{E}(r)$ provides a description of the boundary of $\mathcal{B}(\mathcal{E}(r)) = \mathcal{E}(r)$.

Knowing that $\mathcal{B}(r, \mathcal{W})$ preserves compactness under Assumptions 4.2.1 and 4.2.2, a strategic definition of $\mathcal{B}(r, \mathcal{W})$ becomes simpler than in the general case. For any convex set \mathcal{W} and every $w \in \partial\mathcal{W}$, let $\mathcal{N}_{\mathcal{W}}(w)$ denote the set of all outward normal vectors to $\partial\mathcal{W}$ at w . Define the *normal bundle* of \mathcal{W} as

$$\mathcal{N}_{\mathcal{W}} := \{(w, N) \mid w \in \partial\mathcal{W}, N \in \mathcal{N}_{\mathcal{W}}(w)\}.$$

Under Assumptions 4.2.1 and 4.2.2, any payoff $w \in \mathcal{B}(r, \mathcal{W})$ can be attained by a solution $(W, A, \beta, \delta, M, Z, (J^y)_{y \in Y})$ such that on $\llbracket 0, \sigma_1 \rrbracket$, W remains in $\mathcal{B}(r, \mathcal{W})$, (β, δ) enforces A , $W + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and on $\{W \in \partial\mathcal{B}(r, \mathcal{W})\}$, $N^\top \delta(y) \leq 0$ for every $N \in \mathcal{N}_{\mathcal{B}(r, \mathcal{W})}(W)$ and every $y \in Y$. With this strategic definition it is now possible to prove Lemma 4.4.1. We start by showing that \mathcal{B} is monotonic in the set-inclusion sense.

Lemma 4.4.3. *Let $\mathcal{W} \subseteq \mathcal{W}'$. Then $\mathcal{B}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W}')$.*

Proof. By definition, any payoff $w \in \mathcal{B}(r, \mathcal{W})$ can be attained by an enforceable strategy profile with continuation value W such that $W \in \mathcal{B}(r, \mathcal{W})$ up to time σ_1 with $W_{\sigma_1} \in \mathcal{W}$. Moreover, if $W_{\sigma_1-} \in \partial\mathcal{B}(r, \mathcal{W})$ then the jump goes towards the interior of $\mathcal{B}(r, \mathcal{W})$, i.e., W_{σ_1} is contained in the lower half space

$$H(W_{\sigma_1-}, N_{W_{\sigma_1-}}) := \left\{ w \in \mathbb{R}^2 \mid N_{W_{\sigma_1-}}^\top (w - W_{\sigma_1-}) \leq 0 \right\}.$$

Since $\mathcal{W} \subseteq \mathcal{W}'$ implies that $W_{\sigma_1} \in \mathcal{W}'$ or $W_{\sigma_1} \in \mathcal{W}' \cap H(W_{\sigma_1-}, N_{W_{\sigma_1-}})$, respectively, it follows that $\mathcal{B}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W}')$ by maximality of $\mathcal{B}(r, \mathcal{W}')$. \square

The next lemma shows that similarly to discrete time, the operator \mathcal{B} is closely related to self-generation.

Lemma 4.4.4. *Let $\mathcal{W} \subseteq \mathcal{V}$. If $\mathcal{W} \subseteq \mathcal{B}(r, \mathcal{W})$, then $\mathcal{B}(r, \mathcal{W})$ is self-generating. Conversely, if \mathcal{W} is self-generating and convex, then $\mathcal{W} \subseteq \mathcal{B}(r, \mathcal{W})$.*

Proof. Suppose first that $\mathcal{W} \subseteq \mathcal{B}(r, \mathcal{W})$ and fix $w \in \mathcal{B}(r, \mathcal{W})$ arbitrary. By definition, there exists a solution W to (2.8) such that W remains in $\mathcal{B}(r, \mathcal{W})$ on $\llbracket 0, \sigma_1 \rrbracket$ and $W_{\sigma_1} \in \mathcal{W}$ a.s. This implies that $W_{\sigma_1} \in \mathcal{B}(r, \mathcal{W})$ and hence W_{σ_1} is attained by another solution W' to (2.8) that remains in $\mathcal{B}(r, \mathcal{W})$ up to the second jump time σ_2 with $W'_{\sigma_2} \in \mathcal{W}$ a.s. Therefore, the concatenation of W and W' at time σ_1 is a solution on $\llbracket 0, \sigma_2 \rrbracket$. Because Poisson processes have only countably many jumps, a countable iteration of this procedure leads to an enforceable strategy profile, whose continuation value remains in $\mathcal{B}(r, \mathcal{W})$ forever. This shows that $\mathcal{B}(r, \mathcal{W})$ is self-generating.

For the second statement, self-generation implies that for any $w \in \mathcal{W}$, there exists an enforceable strategy profile with continuations that remain in \mathcal{W} . In particular, $W_{\sigma_1} \in \mathcal{W}$ a.s., and by convexity, at the boundary the jump cannot be directed outwards. Therefore, $\mathcal{W} \subseteq \mathcal{B}(r, \mathcal{W})$ by maximality of $\mathcal{B}(r, \mathcal{W})$. \square

We are now ready to show the convergence of the algorithm in Lemma 4.4.1.

Proof of Lemma 4.4.1. Since $(\mathcal{W}_n)_{n \geq 0}$ is decreasing and bounded from below by the empty set, it must converge and its limit satisfies $\mathcal{W}_\infty = \mathcal{B}(r, \mathcal{W}_\infty)$. Because this implies that \mathcal{W}_∞ is self-generating by Lemma 4.4.4, it follows from Proposition 2.4.3 that $\mathcal{W}_\infty \subseteq \mathcal{E}(r)$. To show that also $\mathcal{E}(r) \subseteq \mathcal{W}_\infty$, observe that $\mathcal{E}(r) \subseteq \mathcal{W}_n$ for some $n \geq 0$ implies $\mathcal{E}(r) \subseteq \mathcal{B}(\mathcal{E}(r)) \subseteq \mathcal{B}(r, \mathcal{W}_n) = \mathcal{W}_{n+1}$

by self-generation and Lemmas 4.4.3 and 4.4.4. Since $\mathcal{E}(r) \subseteq \mathcal{W}_0$ it follows that $\mathcal{E}(r)$ is contained in \mathcal{W}_n for every $n \geq 0$ and hence $\mathcal{W}_\infty = \mathcal{E}(r)$. \square

4.5 Discussion of Assumption 4.2.1

In this section we outline the proof of Proposition 4.4.2 in broad strokes and discuss in which steps Assumption 4.2.1 is needed. We show in Section 4.5.1 that Assumption 4.2.1 implies a transferability of incentives from the discontinuous component to the continuous component of the public signal, which is a powerful tool in the proof. It also implies that the characterizations in Theorem 4.2.1 and Proposition 4.4.2 are valid if players do not have access to a public randomization device. In Section 4.5.2, we discuss conditions under which the characterization in Theorem 4.2.1 and Proposition 4.4.2 remain valid in games that violate Assumption 4.2.1. These conditions are not conditions on the game primitives, but rather a verification that a solution to (4.4) with straight line segments in $\mathcal{G}(r, \mathcal{W})$ still describes $\mathcal{B}(r, \mathcal{W})$ in specific games. The characterization of $\partial\mathcal{B}(r, \mathcal{W})$ can be divided into three main parts:

1. Characterization through (4.4) outside of $\mathcal{G}(r, \mathcal{W})$,
2. Characterization of $\mathcal{G}(r, \mathcal{W})$,
3. Closedness of $\mathcal{B}(r, \mathcal{W})$.

Analogous to the motivation of $\partial\mathcal{E}(r) \setminus \mathcal{G}(r)$ in Section 4.1, the characterization of $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ requires only that $\bigcup_{a \in \mathcal{A}} \Xi_a(w, N_w, r, \mathcal{W})$ is non-empty for $w \in \partial\mathcal{B}(r, \mathcal{W})$ and as such does not depend on Assumption 4.2.1. For ease of reference, we state it as a lemma.

Lemma 4.5.1. *Suppose that Assumption 4.2.2 is satisfied and that $\mathcal{W} \subseteq \mathcal{V}^*$ is compact and convex. Let \mathcal{C} oriented by $w \mapsto N_w$ be a solution to (4.4) with positive curvature throughout. Then any payoff $w \in \mathcal{C}$ is attainable by a solution to (2.8) with $(\beta, \delta) \in \Xi_A(W, N_w, r, \mathcal{W})$ that remains on \mathcal{C} until the arrival of the first event or an end point of \mathcal{C} is reached.*

Because the characterization of $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ does not depend on Assumption 4.2.1, the set characterized in Proposition 4.4.2 is a natural candidate for $\mathcal{B}(r, \mathcal{W})$. Let $\tilde{\mathcal{B}}(r, \mathcal{W})$ denote this candidate, that is, it is the largest closed payoff set in \mathcal{V}^* , for which the boundary is given by (4.4) outside of

$$\tilde{\mathcal{G}}(r, \mathcal{W}) := \left\{ w \in \partial\tilde{\mathcal{B}}(r, \mathcal{W}) \left| \begin{array}{l} \exists (a, \delta) \text{ with } (0, \delta) \in \Xi_a(w, N, r, \mathcal{W}) \\ \text{for all normals } N \text{ to } \partial\tilde{\mathcal{B}}(r, \mathcal{W}) \text{ at } w \end{array} \right. \right\}$$

and such that $\tilde{\mathcal{G}}(r, \mathcal{W})$ consists of straight line segments and isolated points only. In this section we discuss briefly how this candidate can be verified and we present the detailed argument in Section 4.5.2.

Let $\tilde{\mathcal{P}}(r, \mathcal{W})$ denote the set of payoffs, where $\partial\tilde{\mathcal{B}}(r, \mathcal{W})$ changes from being a solution to (4.4) to being a straight line segment in a continuously differentiable way. With the following steps one can verify whether $\tilde{\mathcal{B}}(r, \mathcal{W}) = \mathcal{B}(r, \mathcal{W})$:

1. Show that all payoffs in $\tilde{\mathcal{P}}(r, \mathcal{W})$ and all corners of $\tilde{\mathcal{B}}(r, \mathcal{W})$ are contained in $\mathcal{B}(r, \mathcal{W})$. Then public randomization and Lemma 4.5.1 imply that $\tilde{\mathcal{B}}(r, \mathcal{W})$ is self-generating up to the arrival of the first event such that at the time of the first event, the jump of the continuation value lands in \mathcal{W} and is directed inward at the boundary. By maximality of $\mathcal{B}(r, \mathcal{W})$ it follows that $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W})$.

2. Show that no payoffs outside of $\tilde{\mathcal{B}}(r, \mathcal{W})$ are decomposable. This means that $\partial\mathcal{B}(r, \mathcal{W}) \setminus \tilde{\mathcal{B}}(r, \mathcal{W})$ is a solution to (4.4), a contradiction to the definition of $\tilde{\mathcal{B}}(r, \mathcal{W})$.

The verification argument for Steps 1 and 2 is contained in more detail in Section 4.5.2. An illustration of how this verification argument is applied in a specific example is contained in Section 4.6.2.

4.5.1 Assumption 4.2.1 and Public Randomization

The following lemma states that under Assumption 4.2.1, changes in incentives provided through rare events can be compensated by a value transfer based on the continuous component of the public signal. As a result, all enforceable action profiles are also enforceable without the observation of rare events.

Lemma 4.5.2. *Suppose that Assumption 4.2.1 is satisfied.*

1. *If $a \in \mathcal{A}$ is enforceable and has the best response property for player i , then a is enforceable by $(\beta, 0)$ with $\beta^i = 0$.*
2. *If $a \in \mathcal{A}$ is enforced by $(T\phi, \delta)$ for some non-coordinate direction T , then for any $\tilde{\delta}$ there exists $\tilde{\phi}$ such that a is also enforced by $(T\tilde{\phi}, \tilde{\delta})$.*

Proof. Let $G^i(a)$ denote the row vector with entries $g^i(\tilde{a}^i, a^{-i}) - g^i(a)$ for every $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$. It is player i 's change in expected flow payoff by deviating from a^i to any other action. Then (β, δ) enforcing a is equivalent to the condition that $G^i(a) + \beta^i M^i(a) + \delta^i \Lambda^i(a) \geq 0$ for both players $i = 1, 2$, where the inequality is understood componentwise. For the first statement, the best-response property for player i implies that a is enforceable by (β, δ)

with β^i and δ^i equal to 0. Since $\delta^{-i}\Lambda^{-i}(a)$ is a linear combination of the row vectors in $\Lambda^{-i}(a)$, Assumption 4.2.1 implies that there exists $\hat{\beta}$ with $\hat{\beta}^{-i}Q^i(a)M^{-i}(a) = \delta^{-i}\Lambda^{-i}(a)$ and $\hat{\beta}^i = 0$. Therefore, $(\beta + \hat{\beta}^{-i}Q^i(a), 0)$ enforces a . For the second statement, Assumption 4.2.1 implies that for any $\tilde{\delta}$, there exists $\hat{\beta}$ such that $\hat{\beta}^i Q^{-i}(a)M^i(a) = (\tilde{\delta}^i - \delta^i)\Lambda^i(a)$ for $i = 1, 2$. It is straightforward to check that $(T\tilde{\phi}, \tilde{\delta})$ enforces a , where

$$\tilde{\phi} = \phi + \frac{1}{T^1}\hat{\beta}^1 Q^2(a) + \frac{1}{T^2}\hat{\beta}^2 Q^1(a). \quad \square$$

This transfer of incentives between continuous and abrupt information is the key ingredient in the characterization of $\mathcal{G}(r, \mathcal{W})$. Suppose that $\mathcal{G}(r, \mathcal{W})$ is C^1 in a neighbourhood of some payoff $w \in \mathcal{G}(r, \mathcal{W})$ that is decomposed by (a, δ_0) with $N_w^\top(g(a) - w) > 0$. Due to Lemma 4.5.2, sufficient incentives can be provided in any direction for any perturbation δ of δ_0 . Thus, for any (v, N) in a neighbourhood of (w, N_w) , there exists (β, δ) restricted-enforcing a , where β can be made arbitrarily small by choosing v and δ close to w and δ_0 . Therefore, there exists a solution \mathcal{C} to (4.4) with arbitrarily large curvature in a neighbourhood of w , hence such a solution intersects $\partial\mathcal{B}(r, \mathcal{W})$ on both sides of w . By Lemma 4.5.1, any payoff in \mathcal{C} can be attained by a suitable solution to (2.8) that remains in $\mathcal{C} \cup \mathcal{B}(r, \mathcal{W})$ until the arrival of the first event, at which point the continuation value jumps to \mathcal{W} . This contradicts maximality of $\mathcal{B}(r, \mathcal{W})$. Therefore, $N_w^\top(g(a) - w) = 0$ for all action profile $a \in \mathcal{A}$ that decompose w , implying that $\mathcal{G}(r, \mathcal{W})$ consists of straight line segments and isolated points only. As a consequence, all corners of $\mathcal{B}(r, \mathcal{W})$ are contained in $\mathcal{V}^N \subseteq \mathcal{B}(r, \mathcal{W})$. See Lemma 4.7.15 and its corollaries for details.




Figure 4.6: The left panel illustrates that if $w \in \mathcal{G}(r, \mathcal{W})$ is decomposed by (a, δ_0) with $N_w^\top(g(a) - w) > 0$, then there exists δ close to δ_0 with $(\beta, \delta) \in \Xi_a(v, N, r, \mathcal{W})$ for (v, N) in a neighbourhood of (w, N_w) . Thus, $\mathcal{B}(r, \mathcal{W})$ could be enlarged by a solution \mathcal{C} to (4.4), a contradiction. The right panel illustrates that the controls a^*, δ^* lead to a strictly inward-pointing drift at $w \in \tilde{\mathcal{P}}(r, \mathcal{W})$. For w to be in $\mathcal{B}(r, \mathcal{W})$, it is thus necessary that a^* is restricted-enforceable in a neighbourhood of w .

Lemma 4.5.2 is also sufficient to show Step 2 above. Let $w \in \tilde{\mathcal{P}}(r, \mathcal{W})$, then by definition of $\tilde{\mathcal{P}}(r, \mathcal{W})$, the boundary of $\tilde{\mathcal{B}}(r, \mathcal{W})$ is a solution to (4.4) on one side of w . Let a^*, ϕ^* and δ^* be the maximizers in (4.4) as v approaches w . One can show that $N_w^\top(g(a^*) + \delta^* \lambda(a^*) - w) > 0$, and hence the continuation value of a strategy profile attaining w may fall strictly into the interior of $\tilde{\mathcal{B}}(r, \mathcal{W})$. For w to be in $\mathcal{B}(r, \mathcal{W})$, it is necessary that a^* is restricted-enforceable in a neighbourhood of w . This is clearly satisfied under Assumption 4.2.1 as a^* is enforceable by $(T_w \phi, 0)$ due to Lemma 4.5.2.

Finally, Lemma 4.5.2 implies the nice corollary that the characterizations of $\mathcal{E}(r)$ and $\mathcal{B}(r, \mathcal{W})$ in Theorem 4.2.1 and Proposition 4.4.2, respectively, do not rely on public randomization if the game satisfies Assumption 4.2.1.

Corollary 4.5.3. *Under Assumptions 4.2.1 and 4.2.2, any payoff in $\mathcal{B}(r, \mathcal{W})$ is achievable without public randomization.*

Proof. We need to show that any payoff in $\mathcal{B}(r, \mathcal{W})$ is attainable by a solution W to (2.8) without public randomization. Lemma 4.5.2 shows that any action

profile is enforceable with $\delta = 0$. Therefore, any payoff in the interior of $\mathcal{B}(r, \mathcal{W})$ can be attained by an enforceable strategy profile without jumps or public randomization until the continuation value reaches the boundary of $\mathcal{B}(r, \mathcal{W})$. On the boundary, it is clear that while such a solution is in extremal payoffs of $\mathcal{B}(r, \mathcal{W})$, it makes no use of public randomization. It thus remains to show that any payoff in the relative interior of a straight line segment $L \subseteq \partial\mathcal{B}(r, \mathcal{W})$ is attainable without public randomization such that an extremal payoff or the interior of $\mathcal{B}(r, \mathcal{W})$ is reached with certainty. This follows from another application of Lemma 4.5.2: Action profile a restricted-enforceable at an end point of L is restricted-enforceable orthogonal to that line segment with $\delta = 0$ by Lemma 4.5.2. Since for $\delta = 0$ the controls are location-independent, a is restricted-enforceable on all of L . Play of the constant strategy profile $A \equiv a$ with $\delta \equiv 0$ thus either enters the interior of $\mathcal{B}(r, \mathcal{W})$ if the drift is strictly inward pointing, or reaches an end point with certainty when the drift is parallel to L . \square

4.5.2 Verification of Candidate

For any $w \in \tilde{\mathcal{P}}(r, \mathcal{W})$, the boundary of $\tilde{\mathcal{B}}(r, \mathcal{W})$ is a non-trivial solution \mathcal{C} to (4.4) on one side of w . For $v \in \mathcal{C} \setminus \{w\}$, let $a^*(v)$, $\phi^*(v)$ and $\delta^*(v)$ denote the maximizers in (4.3) and let $a^*(w)$, $\phi^*(w)$ and $\delta^*(w)$ denote their limit as v approaches w . Observe that $a^*(w)$, $\phi^*(w)$ and $\delta^*(w)$ need not be the maximizers of (4.3) at the payoff w itself, but $(\phi^*(w), \delta^*(w))$ is contained in $\Xi_{a^*(w)}(w, N, r, \mathcal{W})$ for any $N \in \mathcal{N}_{\mathcal{B}(r, \mathcal{W})}(w)$ as we show in the subsequent Lemmas 4.7.6 and 4.7.18. The following lemma expresses conditions on the

optimal incentives at corners and payoffs in $\tilde{\mathcal{P}}(r, \mathcal{W})$ that are sufficient for $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W})$ to hold. Its proof is based on the proof of Proposition 4.4.2, replacing Assumption 4.2.1 where needed. It is thus deferred to Section 4.8.

Lemma 4.5.4. *Suppose that there are finitely many corners and payoffs in $\tilde{\mathcal{P}}(r, \mathcal{W})$ such that*

1. *every $w \in \tilde{\mathcal{P}}(r, \mathcal{W})$ satisfies $\tilde{w} + r\delta^*(w; y) \in \text{int } \mathcal{W}$ for every $y \in Y$,*
2. *every corner w of $\tilde{\mathcal{B}}(r, \mathcal{W})$ is decomposable by (a, δ_0) and there exists $\eta > 0$ sufficiently small such that $w + \eta T \in \tilde{\mathcal{B}}(r, \mathcal{W})$ and $w + r\delta(y) + \eta T \in \mathcal{W}$ for every $y \in Y$, where $T = w - g(a) - \delta_0\lambda(a)$.*

Then $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W})$.

Remark 4.5.1. Note that Condition 2 above is satisfied if $w = g(a) + \delta_0\lambda(a)$ and hence $T = 0$. Often, δ_0 is the minimal value burning necessary to restricted-enforce a in the direction N , which is why $\mathcal{B}(r, \mathcal{W})$ cannot be enlarged at w . When the continuation value W reaches such a payoff, W is temporarily absorbed in w until the arrival of the next event because the drift rate T is zero. See Section 4.6.2 for such an example.

Another sufficient condition is $N^\top(g(a) + \delta\lambda(a) - w) > 0$ for all outward normals N to $\partial\mathcal{B}(r, \mathcal{W})$ at w and $w + r\delta_0(y) \in \text{int } \mathcal{W}$ for every $y \in Y$. Indeed, then T points strictly into the interior of $\tilde{\mathcal{B}}(r, \mathcal{W})$.

For any action profile $a \in \mathcal{A}$, let $I(a)$ denote the set of all controls $\delta \in \mathbb{R}^{m \times I}$ such that $(0, \delta)$ enforces a . The set $I(a)$ is easily computed as solution to a system of linear inequalities. Let Ω_a be the set of all payoffs outside of $\tilde{\mathcal{B}}(r, \mathcal{W})$

that are decomposable by (a, δ) for some $\delta \in I(a)$, that is,

$$\Omega_a := \left\{ w \in \mathcal{W} \setminus \tilde{\mathcal{B}}(r, \mathcal{W}) \left| \begin{array}{l} \exists \delta \in I(a) \text{ and } N \text{ normal to } \text{co}(w, \tilde{\mathcal{B}}(r, \mathcal{W})) \text{ s.t.} \\ N^\top (g(a) + \delta\lambda(a) - w) \geq 0 \text{ and for all } y \in Y \\ w + r\delta(y) \in \mathcal{W} \text{ and } N^\top \delta(y) \leq 0. \end{array} \right. \right\}$$

Suppose that we have verified already that $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W})$. It is possible that $\mathcal{B}(r, \mathcal{W}) \supsetneq \tilde{\mathcal{B}}(r, \mathcal{W})$ only if $\partial\mathcal{B}(r, \mathcal{W})$ contains C^1 line segments in $\mathcal{G}(r, \mathcal{W})$. Clearly, any payoff in $\mathcal{G}(r, \mathcal{W}) \setminus \tilde{\mathcal{B}}(r, \mathcal{W})$ has to be contained in $\bigcup_{a \in \mathcal{A}} \Omega_a$, which implies the following lemma.

Lemma 4.5.5. *Suppose that $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W}) \subseteq \mathcal{W}$. If $\Omega_a = \emptyset$ for every $a \in \mathcal{A}$, then $\tilde{\mathcal{B}}(r, \mathcal{W}) = \mathcal{B}(r, \mathcal{W})$.*

We conclude this section by stating simple sufficient conditions for $\Omega_a = \emptyset$ to hold. Let $\gamma(a) := g(a) + I(a)\lambda(a)$ and $\Delta(a) := \bigcap_{y \in Y} (\mathcal{W} - rI(a)e_y) \cap \mathcal{W}$, where e_y denotes the unit vector in \mathbb{R}^m that corresponds to event y .

Lemma 4.5.6. *Suppose that for any action profile $a \in \mathcal{A}$, either*

1. $g(a) \in \tilde{\mathcal{B}}(r, \mathcal{W})$, or
2. *there exists no $w \in \partial\tilde{\mathcal{B}}(r, \mathcal{W})$ with outward normal N such that the intersections of the upper half-space $\bar{H}(N, w) := \{v \in \mathbb{R}^2 \mid N^\top(v - w) \geq 0\}$ with $\gamma(a)$ and $\Delta(a)$ are non-empty simultaneously.*

Then $\Omega_a = \emptyset$.

Proof. The first condition is obviously sufficient since $N^\top (g(a) + \delta\lambda(a) - w) \geq 0$ for any payoff w outside of $\tilde{\mathcal{B}}(r, \mathcal{W})$ implies that $N^\top \delta(y) > 0$ for at least

one $y \in Y$. For the second condition, suppose that there exists $v \in \Omega_a$, which implies that $v \in \Delta(a)$. By definition of Ω_a , there exist $\delta \in I(a)$ and an outward normal N to $\text{co}(w, \tilde{\mathcal{B}}(r, \mathcal{W}))$ such that $N^\top(g(a) + \delta\lambda(a) - v) \geq 0$. Because $v \notin \tilde{\mathcal{B}}(r, \mathcal{W})$, there exists $w \in \partial\tilde{\mathcal{B}}(r, \mathcal{W})$ with outward normal N such that $N^\top(v - w) \geq 0$ and hence $N^\top(g(a) + \delta\lambda(a) - w) \geq 0$. The intersections of $\overline{H}(w, N)$ with $\gamma(a)$ and $\Delta(a)$ are thus both non-empty, a contradiction. \square

4.6 Computation

In this section we illustrate how to numerically compute $\partial\mathcal{E}(r)$ with two examples where the characterization is new. We focus mainly on the case where Y is non-empty and $\mathcal{E}(r)$ is computed with the algorithm in Lemma 4.4.1 since the continuous monitoring case is described in detail in Section 8 of Sannikov [37]. In Section 4.6.1 we illustrate how the algorithm of Lemma 4.4.1 converges to $\mathcal{E}(r)$ in the climate agreement example of Sections 2.2 and 4.3.1. In Section 4.6.2 we present a partnership example with a one-dimensional public signal both when monitoring is continuous and when, in addition, infrequent events are observed. The characterization is new even in the continuous-monitoring case and $\partial\mathcal{E}(r)$ may have straight line segments that are precluded under the conditions in Sannikov [37]. When information contains both continuous and discontinuous information, Assumption 4.2.1 is not satisfied in this game and we show how to verify the candidate described in Section 4.5.

At payoffs where $\partial\mathcal{B}(r, \mathcal{W})$ is the solution to (4.4), the implementation works similarly to Sannikov [37], that is, the boundary of $\mathcal{B}(r, \mathcal{W})$ is parametrized in terms of the tangential angle θ . Let $w(\theta)$ denote the set of payoffs

in $\mathcal{B}(r, \mathcal{W})$ with normal vector $N(\theta) = (\cos(\theta), \sin(\theta))^\top$. Then $w(\theta)$ is unique where the curvature of $\partial\mathcal{B}(r, \mathcal{W})$ is strictly positive and one can solve

$$\frac{dw(\theta)}{d\theta} = \frac{T(\theta)}{\kappa(\theta)} \quad (4.5)$$

numerically, where $T(\theta) = (-\sin(\theta), \cos(\theta))^\top$ and $\kappa(\theta) = \kappa(w(\theta))$ is given by the optimality equation (4.4). Because $\mathcal{B}(r, \mathcal{W})$ is the largest bounded set with the curvature given by (4.4), we search for an extremal pair of initial values $(\theta, w(\theta))$, for which (4.5) has a closed solution. We can thus search either for the extremal initial angle θ for a fixed starting point $w_0 \in \partial\mathcal{B}(r, \mathcal{W})$, or for the extremal starting point w given a fixed initial angle θ_0 . The former is practical if there exists a static Nash payoff $g(a_e)$ in $\partial\mathcal{V}^*$. Because $\mathcal{V}^N \subseteq \mathcal{B}(r, \mathcal{W}) \subseteq \mathcal{V}^*$, it follows that $g(a_e) \in \partial\mathcal{B}(r, \mathcal{W})$ and hence we know a possible starting point. The latter is useful if the game is symmetric, hence there has to be a point $w \in \partial\mathcal{B}(r, \mathcal{W})$ on the positive diagonal with $\theta = \pi/4$.

4.6.1 Climate agreement

The climate agreement example of Sections 2.2 and 4.3.1 is symmetric, hence we can search for a starting value (w_0, θ_0) on the positive diagonal with $\theta_0 = \pi/4$. Because the game has a static Nash equilibrium at $(0.35, 0.35) \in \partial\mathcal{V}^*$, a continuous solution of the optimality equation must reach $(0.35, 0.35)$ since $(0.35, 0.35) \in \partial\mathcal{B}(r, \mathcal{W}_n)$ for any $n \geq 0$ as argued above. To find the largest of these solutions in the first iteration, we first perform a grid search on the positive diagonal and check whether a solution crosses the line $\{w_2 = 0.35\}$ below or above $w_1 = 0.35$. Because (4.4) is continuous in initial conditions,




Figure 4.7: The left panel shows solutions to the optimality equation with different starting values on the positive diagonal. Because $\mathcal{V}^N \subseteq \mathcal{B}(r, \mathcal{W}_n)$ for every $n \geq 0$, we search for an initial value such that the solution reaches the Nash payoff $(0.35, 0.35)$. The left panel illustrates that for \mathcal{W}_1 , the initial value has to lie in $[0.65, 0.7]$, and a binary search reveals that it equals 0.6593 . The right panel shows an iteration of this procedure, approximating $\mathcal{E}(r)$.

the starting value has to lie between 0.65 and 0.7 as illustrated in the left panel of Figure 4.7. Then we perform a binary search between 0.65 and 0.7 to find the exact starting value of 0.6593 . Solving the optimality equation for starting value 0.6593 , we obtain \mathcal{W}_1 , the largest of the sets depicted in the right panel of Figure 4.7. An iteration of this procedure leads to the decreasing sequence $(\mathcal{W}_n)_{n \geq 0}$, which converges to $\mathcal{E}(r)$. We stop the numerical iteration when we consider the difference in area between \mathcal{W}_n and \mathcal{W}_{n-1} as small enough.

4.6.2 Partnership game

Consider a simple partnership game amongst two players, where each player continuously chooses an effort level from $\mathcal{A}^i = \{0, 1\}$ at every point in time t . Players observe only the total revenue $2X_t$, where $dX_t = \mu(A_t) dt + dZ_t^A$ for a Q^A -Brownian motion Z^A and $\mu(a) = 4a^1 + 4a^2 - a^1a^2$. Suppose that players




Figure 4.8: Because $\mathcal{E}(r)$ is closed, it is contained in the payoff bound \mathcal{M} of Sannikov and Skrzypacz [39]. To increase the speed of the algorithm we choose $\mathcal{W}_0 = \mathcal{M}$ and search for initial values on $K \subseteq \partial\mathcal{M}$ as indicated in the right panel. A binary search between $(0.49, 1.31)$ and $(0.57, 1.53)$ yields that $(0.5529, 1.4742)$ is the largest value (in norm) for which a closed solution to (4.3) exists.

share the revenue equally and are subject to a cost of effort of $5a^i$ so that the expected flow payoff of player i is $g^i(a) = 4(a^1 + a^2) - a^1a^2 - 5a^i$. Because Assumptions 4.2.1 and 4.2.2 are satisfied, $\mathcal{E}(r)$ is closed and hence contained in the payoff bound \mathcal{M} from Sannikov and Skrzypacz [39]; see Figure 4.8. This speeds up the search for the straight line segments significantly.

The static Nash payoff is contained in $\partial\mathcal{M}$ and hence has to be contained in $\partial\mathcal{E}(r)$. However, since action profiles are not pairwise identifiable, $\partial\mathcal{E}(r)$ may contain straight line segments that we have to search for. From Proposition 4.4.2 we know that straight line segments go through $g(a)$ for some action profile a . Thus, $\partial\mathcal{E}(r)$ could have a straight line segment starting right at \mathcal{V}^N . To find the largest possible solution, we first perform a grid search for a starting value on edge K of \mathcal{M} as indicated in the right panel of Figure 4.8, where θ is such that $N(\theta)$ is normal to K . If the grid search is successful, we




Figure 4.9: We search jointly for starting value and initial angle: for each fixed angle (dashed lines), we perform a search over starting value similar to the right panel in Figure 4.8. After each corner, we repeat this search as indicated in the left panel. The right panel shows $\gamma(1, 1)$ and $\Delta(1, 1)$ for the verification with Lemma 4.5.6.

perform a binary search to find the exact starting value. If the grid search is not successful even after refining the grid size, then $\partial\mathcal{E}(r) \cap \partial\mathcal{M} = \mathcal{V}^N$, hence we perform a search over initial angles for fixed starting value $w_0 = (0, 0)$.

Next, suppose that revenue arrives continuously, but accidents may occur at a rate of $\lambda(a) = 21 - 4(a^1 + a^2) - 12a^1a^2$ that cost 0.1 each. The total revenue is given by $2X_t$ where $dX_t = \mu(A_t) dt + dZ_t^A - 0.05 dJ_t^A$ and $\mu(a) = 1.05 + 3.8(a^1 + a^2) - 1.6a^1a^2$ so that the expected flow payoff is the same as in the continuous monitoring example. This information structure violates Assumption 4.2.1 and hence we need to verify that the largest closed solution $\tilde{\mathcal{B}}(r, \mathcal{W}_n)$ to (4.4) coincides with $\mathcal{B}(r, \mathcal{W}_n)$ for any $n \geq 0$ as elaborated in Section 4.5. In this example, we have no indication in which direction a straight line segment of $\partial\tilde{\mathcal{B}}(r, \mathcal{W})$ could go from \mathcal{V}^N , hence we have to search jointly over initial angle and starting value to find $\tilde{\mathcal{B}}(r, \mathcal{W})$ as illustrated in

the left panel of Figure 4.9. We perform a search over the initial angles, say, $\theta_1, \dots, \theta_\ell$, and for each $j = 1, \dots, \ell$, we search for the starting point, given the initial angle θ_j , to find the largest closed solution to (4.3). If the curvature of the optimality equation ever equals 0, as in the point A in the left panel of Figure 4.9, we perform a search on the tangent T_A through A to find the payoff furthest away from A , from which a solution to (4.3) connects to \mathcal{V}^N . This could potentially be a complicated procedure since at every payoff in $\tilde{\mathcal{G}}(r, \mathcal{W})$, the set $\tilde{\mathcal{B}}(r, \mathcal{W})$ could have a corner and one has to search jointly for initial angle and starting value again. In this specific example, the search is concluded when we find a continuous solution connecting to $B \in \tilde{\mathcal{G}}(r, \mathcal{W})$ due to symmetry.

We proceed to verify $\tilde{\mathcal{B}}(r, \mathcal{W}) = \mathcal{B}(r, \mathcal{W})$ as explained in Section 4.5.2. First, we need to check that payoffs A, B and C are contained in $\mathcal{B}(r, \mathcal{W})$ so that $\tilde{\mathcal{B}}(r, \mathcal{W}) \subseteq \mathcal{B}(r, \mathcal{W})$. Indeed, at payoffs A and C , value burning is not necessary to restricted-enforce $(1, 0)$ and $(0, 1)$, respectively. Thus, Condition 1 of Lemma 4.5.4 is satisfied. The minimal value burning necessary to enforce action profile $(1, 1)$ at $B = (1.875, 1.875)$ on the negative diagonal is $\delta_0 = (0.125, 0.125)$. One can check that $(0, \delta_0)$ enforces $(1, 1)$ and that $B = g(1, 1) + \delta_0 \lambda(1, 1)$. Therefore, Condition 2 of Lemma 4.5.4 is satisfied with $T = 0$. A solution to (2.8) with $A \equiv (1, 1)$, $\beta \equiv 0$, $\delta \equiv \delta_0$ and $M \equiv 0$ will thus remain in B until the arrival of an event, at which point it jumps to \mathcal{V}^N and is absorbed there forever.

Next, we check that no payoffs outside of $\tilde{\mathcal{B}}(r, \mathcal{W})$ are decomposable so that $\mathcal{B}(r, \mathcal{W}) \subseteq \tilde{\mathcal{B}}(r, \mathcal{W})$. Because the static Nash payoff is contained in $\tilde{\mathcal{B}}(r, \mathcal{V}^*)$, the static Nash profile $(0, 0)$ satisfies Condition 1 of Lemma 4.5.6. The right

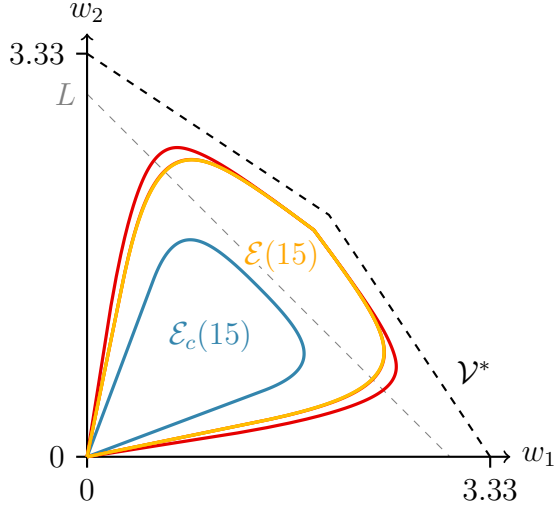


Figure 4.10: Shown is the (fast) convergence of $(\mathcal{W}_n)_{n \geq 0}$ to $\mathcal{E}(15)$ and the comparison to $\mathcal{E}_c(15)$ with continuous monitoring. With continuous monitoring only, action profile $(1, 1)$ cannot be enforced on the negative diagonal and hence $\mathcal{E}_c(r)$ lies below L for any value of the discount rate.

panel of Figure 4.9 shows $\gamma(1, 1)$ and $\Delta(1, 1)$ for the action profile of mutual effort, and it is easy to see that Condition 2 of Lemma 4.5.6 holds. Since $\Delta(0, 1)$ and $\Delta(1, 0)$ are both empty, Condition 2 of Lemma 4.5.6 is trivially satisfied for these two action profiles. Therefore, Lemma 4.5.5 applies and shows that $\tilde{B}(r, \mathcal{V}^*) = \mathcal{B}(r, \mathcal{V}^*)$. The same verification argument shows that $\tilde{B}(r, \mathcal{W}_n) = \mathcal{B}(r, \mathcal{W}_n)$ for any $n \geq 0$ and hence $\mathcal{E}(r)$ can be computed as limit of the algorithm as shown in Figure 4.10.

Figure 4.10 shows once more the drastic effect that abrupt information can have on the equilibrium payoff set. Because the action profile of mutual effort is not enforceable on the negative diagonal by observing the total revenue only, $\mathcal{E}_c(r)$ remains below L for any discount rate. This is consistent with the payoff bound of Sannikov and Skrzypacz [39], which is contained entirely below L in the continuous-monitoring game. With the observation of the accidents,

$(1, 1)$ can be enforced by burning 0.125 units of payoff of each player upon the arrival of an accident. However, this is possible only for discount rates $r \leq 15$. For higher discount rates, the induced jump in the continuation value does not land in \mathcal{V}^* anymore, meaning that the punishment is inconsistent with equilibrium behaviour. Therefore, for $r > 15$, the equilibrium payoff set $\mathcal{E}(r)$ is also bounded below above by L . This shows that $\mathcal{E}(r)$ is not continuous in the discount rate when players observe abrupt information.

Let us discuss what happens when the continuation value of a PPE reaches a corner of $\mathcal{E}(r)$. If it ever reaches a corner $g(a_e) \in \partial\mathcal{E}(r)$ for a static Nash equilibrium $a_e \in \mathcal{A}^N$, then it is absorbed there forever, similarly to Sannikov [37]. Indeed, tangential incentives cannot be provided at corners, hence $\beta = 0$, and because value can only be burnt upon the arrival of an event, $N^\top \delta(y) \leq 0$ for every $y \in Y$ implies that $N^\top (g(a^e) + \delta\lambda(a_e))$ lies in the interior of $\mathcal{E}(r)$, i.e., it would lead to an outward-pointing drift. At corners outside \mathcal{V}^N as in B in this example, the situation looks surprisingly similar. Because $B = g(a) + \delta_0\lambda(a)$ for the minimal amount of value burning δ_0 necessary to enforce a in that direction, incentives $(0, \delta_0)$ are unique: no smaller amount of value burning enforces a and no larger amount of value burning has an inward-pointing (or zero) drift. Thus, the solution W to (2.8) is locally unique with drift rate $w - g(a) - \delta_0\lambda(a) = 0$ and hence W remains in B until the arrival of the next event. Because we consider games of full support public monitoring, the event will occur eventually, after which play becomes dynamic again — or, in this example, the game becomes absorbed in \mathcal{V}^N .

Absorption in the static Nash equilibrium corresponds to the termination of the partnership as no player will ever exert effort afterwards anymore. Sit-

uations, in which incentives can be provided by the abrupt information exclusively correspond to locally static equilibria. In this partnership example, at payoff B the partnership is going at its best. The sum of payoffs is maximized in equilibrium and players trust each other to exert effort without monitoring the continuous revenue. The reward level in this situation is so high that no player wishes to deviate, even though the arrival of an accident will lead to the termination of the partnership (because the jump goes into \mathcal{V}^N with probability one). Nevertheless, this termination of the partnership upon the arrival of an accident is necessary to enforce the trust between players.

4.7 Proof of Proposition 4.4.2

4.7.1 Strategic definition of $\mathcal{B}(r, \mathcal{W})$

A strategic definition of $\mathcal{B}(r, \mathcal{W})$ under suitable conditions that show closedness of $\mathcal{B}(r, \mathcal{W})$ has been provided in Section 4.4. However, because it is not a priori clear whether $\mathcal{B}(r, \mathcal{W})$ is closed or not, in general the condition that the jumps are directed towards the interior on $\partial\mathcal{B}(r, \mathcal{W})$ has to be defined over a limit as the distance to the boundary becomes shrinks to zero.

Definition 4.7.1. For any $\varepsilon > 0$, $r > 0$ and $\mathcal{W} \subseteq \mathbb{R}^2$, call a payoff set $\tilde{\mathcal{W}}$ *locally $(\varepsilon, r, \mathcal{W})$ -admissible* if for every payoff $w \in \tilde{\mathcal{W}}$, there exist a stopping time τ and a solution $(W, A, \beta, \delta, M, Z, (J^y)_{y \in Y})$ to (2.8) with $W_0 = w$ such that for $0 \leq t < \sigma_1 \wedge \tau$, the following conditions are satisfied: $W_t \in \tilde{\mathcal{W}}$, (β_t, δ_t) enforces A_t , $W_t + r\delta_t(y) \in \mathcal{W}$ for every $y \in Y$, $N^\top \delta_t(y) \leq 0$ for every $y \in Y$ and every normal N to $\partial\tilde{\mathcal{W}}$ in $B_\varepsilon(W_t)$.

We say that a payoff set is $(\varepsilon, r, \mathcal{W})$ -admissible if for every payoff w within the set, there exists a solution to (2.8) attaining w satisfying the above properties for $\tau = \infty$. Denote by $\mathcal{B}_\varepsilon(r, \mathcal{W})$ the largest $(\varepsilon, r, \mathcal{W})$ -admissible payoff set. The following lemma and its corollary establish that this set is well defined.

Lemma 4.7.1. *Let $\mathcal{W}_1, \dots, \mathcal{W}_n$ be locally $(\varepsilon, r, \mathcal{W})$ -admissible and let $\overline{\mathcal{W}}$ denote the convex hull of $\mathcal{W}_1, \dots, \mathcal{W}_n$. For any $\eta > 0$ and $k = 1, \dots, n$, define*

$$\mathcal{D}_{k,\eta} := \mathcal{W}_k \cap B_\eta(\partial\mathcal{W}_k \cap \partial\overline{\mathcal{W}}).$$

If $\tau_k = \inf_{w \in \mathcal{D}_{k,\varepsilon/2}} \sup_W \inf\{t \geq 0 \mid W_t \notin \mathcal{D}_{k,\varepsilon}\} > 0$ a.s., where the supremum is taken over all solutions to (2.8) attaining w satisfying properties 1–4 in the definition of local admissibility, then $\overline{\mathcal{W}}$ is $(\varepsilon, r, \mathcal{W})$ -admissible.

Corollary 4.7.2. *If $\mathcal{W}_1, \mathcal{W}_2$ are $(\varepsilon, r, \mathcal{W})$ -admissible, so is the convex hull $\text{co}(\mathcal{W}_1, \mathcal{W}_2)$.*

Proof of Lemma 4.7.1. We show that any payoff in $w \in \overline{\mathcal{W}}$ can be attained by a suitable solution to (2.8). Note that any payoff outside $\mathcal{D}_{\varepsilon/2}^\circ := \bigcup_{k=1}^n \mathcal{D}_{k,\varepsilon/2}$ is attainable by public randomization at time 0 taking values in $\mathcal{D}_{1,\varepsilon/2}, \dots, \mathcal{D}_{n,\varepsilon/2}$. We may thus assume that $w \in \mathcal{D}_{k,\varepsilon/2}$ for some k . By local admissibility, there exists a solution $(W, A, \beta, \delta, M, Z, (J^y)_{y \in Y})$ to (2.8) starting in w such that on $\llbracket 0, \sigma_1 \wedge \tau_k \rrbracket$, W remains in \mathcal{W}_k , (β, δ) enforces A and $\delta(y) \in \mathcal{D}(W) \cap H(N)$ for normal vectors N to $\partial\mathcal{W}_k$ within distance ε of W . This solution also satisfies the inward-jump condition with respect to $\partial\overline{\mathcal{W}}$ outside of the set $\mathcal{H} := \bigcup_{k=1}^n \mathcal{H}_k$, where $\mathcal{H}_k := \mathcal{W}_k \cap B_\varepsilon(\partial\overline{\mathcal{W}}) \setminus \mathcal{D}_{k,\varepsilon}$; see Figure 4.11 for an illustration of \mathcal{H} . The main idea of the proof is thus to concatenate solu-



Figure 4.11: Payoffs in $\mathcal{H} \cup \overline{\mathcal{W}} \setminus (\mathcal{W}_1 \cup \dots \cup \mathcal{W}_n)$ are attained by a public randomization device with values in $\mathcal{D}_{\varepsilon/2}^\circ$.

tions to (2.8) that exist by local admissibility and to use public randomization whenever the solution enters \mathcal{H} . We do this in two steps.

Step 1: We show that if $\tau_k < \infty$ with positive probability, it is possible to find concatenations of solutions to (2.8) that satisfy the inward-jump condition until either an event occurs or the continuation value reaches \mathcal{H} . Indeed, on the set $\{W_{\tau_k} \notin \mathcal{D}_{k,\varepsilon/2}\}$ the payoff W_{τ_k} is attainable by a public randomization device at time τ_k with values in $\mathcal{D}_{\varepsilon/2}^\circ$. We may thus assume that $W_{\tau_k} \in \mathcal{D}_{j,\varepsilon/2}$ for some $j \in \{1, \dots, n\}$ without loss of generality. Therefore, there exists a solution $(\hat{W}, \hat{A}, \hat{\beta}, \hat{\delta}, \hat{M}, \hat{Z}, (\hat{J}^y)_{y \in Y})$ to (2.8) attaining W_{τ_k} such that on $\llbracket 0, \sigma_1 \wedge \tau_j \rrbracket$, $\hat{W} \in \mathcal{W}_j$, $(\hat{\beta}, \hat{\delta})$ enforces \hat{A} , $\hat{W} + r\hat{\delta}(y) \in \mathcal{W}$ for every $y \in Y$ and $N^\top \hat{\delta}(y) \leq 0$ for every $y \in Y$ and every outward normal N to $\partial \mathcal{W}_j$ within $B_\varepsilon(\hat{W})$. Define the concatenations $\tilde{Z}_t := Z_{t \wedge \tau_k} + \hat{Z}_{t - \tau_k} 1_{\{t > \tau_k\}}$ and similarly for \tilde{M} and $(\tilde{J}^y)_{y \in Y}$. Moreover, set $\tilde{W}_t := W_t 1_{\{t \leq \tau_k\}} + \hat{W}_{t - \tau_k} 1_{\{t > \tau_k\}}$ and define \tilde{A} , $\tilde{\beta}$ and $\tilde{\delta}$ analogously to \tilde{W} . Then the tilde-processes are a solution to (2.8) up to a stopping time $\tilde{\tau}$ such that $\tilde{\tau} - \tau_k$ is identically distributed as τ_j and on $\llbracket 0, \tilde{\tau} \rrbracket$, $(\tilde{\beta}, \tilde{\delta})$ enforces \tilde{A} , $\tilde{W} + r\tilde{\delta}(y) \in \mathcal{W}$ for every $y \in Y$ and $N^\top \tilde{\delta}(y) \leq 0$ for every $y \in Y$ and every outward normal N to $\partial \overline{\mathcal{W}}$ within $B_\varepsilon(\tilde{W})$ where $\tilde{W} \notin \mathcal{H}$. Since τ_1, \dots, τ_n are strictly positive, so is $\tau_0 := \min(\tau_1, \dots, \tau_n)$. By repeating this procedure, each step extends the solution on a time interval with length of at least τ_0 . A countable iteration thus extends to infinity with probability 1

by Lemma 3.6.3. We have thus constructed a concatenation of solutions that satisfies the inward-jump condition with respect to $\partial\overline{\mathcal{W}}$ outside of \mathcal{H} .

Step 2: Let $\rho := \inf\{t \geq 0 \mid W_t \in \mathcal{H}\}$. Since \mathcal{H} is bounded away from $\partial\overline{\mathcal{W}}$, W_ρ is attainable by a public randomization device with values in some finite set $\{w_1, \dots, w_K\} \subseteq \mathcal{D}_{\varepsilon/2}^\circ$. For each w_ℓ , denote by $(W^\ell, A^\ell, \beta^\ell, \delta^\ell, M^\ell, Z^\ell, (J^{\ell,y})_{y \in Y})$ the concatenation of solutions to (2.8) obtained through step 1 above. Observe that none of these solutions reach \mathcal{H} again before time τ_0 by assumption. Thus, concatenating independent copies of solutions attaining w_1, \dots, w_n yields a solution on $[0, \infty)$ again by Lemma 3.6.3. \square

Finally, define $\mathcal{B}(r, \mathcal{W}) := \bigcup_{\varepsilon > 0} \mathcal{B}_\varepsilon(r, \mathcal{W})$. Note that $\mathcal{B}_\varepsilon(r, \mathcal{W})$ is monotone in ε because jumps are inward-pointing within distance ε' if they are within distance $\varepsilon > \varepsilon'$. As a consequence, $\mathcal{B}(r, \mathcal{W})$ is convex as the limit of a non-decreasing sequence $\mathcal{B}_{1/n}(r, \mathcal{W})$ of convex sets.

4.7.2 Regularity of the optimality equation

The purpose of this section is to prove that the optimality equation (4.4) is locally Lipschitz continuous at almost every point, so that locally, it admits a unique solution. Together with the subsequent Lemma 4.7.14 this will show that $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ is the unique C^1 solution to the optimality equation. For any fixed $r > 0$, $a \in \mathcal{A}$, and closed and convex $\mathcal{W} \subseteq \mathcal{V}$, consider the optimality equation in the following form:

$$\kappa_a(w, N) = \max_{(\phi, \delta) \in \Xi_a(w, N, r, \mathcal{W})} \frac{2N^\top (g(a) + \delta\lambda(a) - w)}{r\|\phi\|^2}. \quad (4.6)$$

We start by reducing the two-variable optimization problem to a one-variable

optimization by expressing ϕ in terms of δ . For players $i = 1, 2$, define

$$\mathcal{I}_a^i(N, \delta^i) := \{\phi \in \mathbb{R}^{d_c} \mid (T^i \phi, \delta^i) \text{ satisfies (2.11) for player } i\}$$

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

Lemma 4.7.3. *Fix $a \in \mathcal{A}$. Then $(N, \delta) \mapsto \phi(a, N, \delta)$ is locally Lipschitz continuous where $\Phi_a(N, \delta) \neq \emptyset$ and N is different from a coordinate direction.*

In an intermediate step, we show that the set-valued map $(N, \delta) \mapsto \Phi_a(N, \delta)$ is locally Lipschitz continuous for N different from coordinate directions. We refer to Aubin and Frankowska [5] for a detailed overview of set-valued maps and their properties and state here only the most central property.

Definition 4.7.2. A set-valued map $G : x \mapsto G(x)$ is said to be *Lipschitz continuous* if $G(x) \subseteq G(\tilde{x}) + K\|x - \tilde{x}\|B_1(0)$ for some constant K .

Proof of Lemma 4.7.3. For any player $i = 1, 2$ and any $\delta^i \in \mathbb{R}^m$, denote by $\mathcal{I}_a^i(\delta^i) := \{\beta \in \mathbb{R}^{d_c} \mid (\beta, \delta^i) \text{ satisfies (2.11) for player } i\}$ the δ^i -restricted solution set to (2.11) for player i and observe that it is a closed convex polytope. Its hyperfaces have normal vectors $\Delta\mu_{j_i}^i := \mu(a) - \mu(a_{j_i}^i, a^{-i})$, where $a_1^i, \dots, a_{m_i}^i$ is an enumeration of $\mathcal{A}^i \setminus \{a^i\}$. The parameter δ^i determines the location of these hyperfaces. Observe that a change from δ^i to $\tilde{\delta}^i$ shifts face j_i by $(\tilde{\delta}^i - \delta^i)\Delta\lambda_{j_i}^i$, where $\Delta\lambda_{j_i}^i := \lambda(a) - \lambda(a_{j_i}^i, a^{-i})$. The triangle inequality thus implies

$$\mathcal{I}_a^i(\delta) \subseteq \mathcal{I}_a^i(\tilde{\delta}) + B_1(0) \sum_{j_i=1, \dots, m_i} \|\Delta \lambda_{j_i}^i\| \|\tilde{\delta}^i - \delta^i\|,$$

i.e., $\mathcal{I}_a^i(\delta^i)$ is Lipschitz continuous in δ^i . It is clear that $\mathcal{I}_a^i(N, \delta^i) = \frac{1}{T^i} \mathcal{I}_a^i(\delta^i)$ for $i = 1, 2$ is locally Lipschitz continuous in (N, δ^i) for N different from coordinate directions. Since the stretching does not affect the direction of the normal vectors, the normal vectors of $\mathcal{I}_a^i(N, \delta^i)$ are constant, which implies that $(N, \delta) \mapsto \Phi_a(N, \delta) = \mathcal{I}_a^1(N, \delta^1) \cap \mathcal{I}_a^2(N, \delta^2)$ is locally Lipschitz continuous by Lemma 4.A.2. The statement now follows from the following Lemma. \square

Lemma 4.7.4. *Let $f(x, y)$ be a single-valued Lipschitz continuous map and let $G(x)$ be a set-valued (locally) Lipschitz continuous map. Then the restricted maximum $h(x) = \max_{y \in G(x)} f(x, y)$ is (locally) Lipschitz continuous.*

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

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

Lemma 4.7.3 significantly simplifies the constraints in the maximization in (4.6) because we are left with a maximization over δ only. We subsequently characterize the set of all eligible δ , over which the maximization takes place.

For any $\varepsilon \geq 0$ and any set-valued map $w \mapsto \mathcal{D}(w) \subseteq \mathbb{R}^2$, define

$$\Psi_a^\varepsilon(w, N, r, \mathcal{D}) := \left\{ \delta \in \mathbb{R}^{2 \times m} \left| \begin{array}{l} \Phi_a(N, \delta) \neq \emptyset, N^\top (g(a) + \delta\lambda(a) - w) \geq 0, \\ \delta(y) \in \mathcal{D}(w) \cap H(\tilde{N}) \quad \forall \tilde{N} \in B_\varepsilon(N) \quad \forall y \in Y \end{array} \right. \right\}$$

so that (4.6) is equivalent to

$$\kappa_a(w, N) = \max_{\delta \in \Psi_a^\varepsilon(w, N, r, \mathcal{D})} \frac{2N^\top (g(a) + \delta\lambda(a) - w)}{r \|\phi(a, N, \delta)\|^2}. \quad (4.7)$$

for $\varepsilon = 0$ and $\mathcal{D}(w) = (\mathcal{W} - w)/r$. In Section 4.7.3 we will be needing regularity of the optimality equation for different choices of ε and $w \mapsto \mathcal{D}(w)$, which is why we do not limit ourselves to $(\mathcal{W} - w)/r$ in this section. We say a map $w \mapsto \mathcal{D}(w) \subseteq \mathbb{R}^2$ is of *class B* if it is affine, convex- and compact-valued with a uniform bound for $w \in \mathcal{V}$.

Remark 4.7.1. Note here that in the definition of Ψ_a^ε , the dependency on normal vectors close to N differs from the definition of $\mathcal{B}_\varepsilon(r, \mathcal{W})$. For $w \in \mathcal{B}_\varepsilon(r, \mathcal{W})$, we require that $N^\top \delta(y) \leq 0$ for all $N \in \partial \mathcal{B}_\varepsilon(r, \mathcal{W}) \cap B_\varepsilon(w)$, whereas in Ψ_a^ε we choose a definition that depends neither on the boundary of a certain set, nor on the location w . This greatly simplifies establishing regularity properties for Ψ_a^ε . The two definitions are easily related, however, for C^1 parts of the boundary $\mathcal{C} \subseteq \partial \mathcal{B}_\varepsilon(r, \mathcal{W})$. Then there exists $\eta(\mathcal{C}, \varepsilon) > 0$ such that for $w \in \mathcal{C}$, $N \in \mathcal{C} \cap B_\varepsilon(w)$ implies that $N \in B_{\eta(\mathcal{C}, \varepsilon)}(N_w)$.

Lemma 4.7.5. *Let $w \mapsto \mathcal{D}(w)$ be of class B. Then for any $a \in \mathcal{A}$ and $\varepsilon \geq 0$, the map $(w, N) \mapsto \Psi_a^\varepsilon(w, N, r, \mathcal{D})$ is compact- and convex-valued. Moreover, it is locally Lipschitz continuous for N different from coordinate directions and ε such that extremal vectors in $B_\varepsilon(N) \cap S^1$ are not coordinate.*

Proof. Identify $\mathbb{R}^{2 \times m}$ with \mathbb{R}^{2m} by setting $\delta \approx (\delta^1, \delta^2)$. For any subset \mathcal{W} of \mathbb{R}^2 , define $\mathcal{W}^Y := \{(\delta^1, \delta^2) \in \mathbb{R}^{2m} \mid (\delta^1(y), \delta^2(y)) \in \mathcal{W} \ \forall y \in Y\}$. Let $\mathcal{J}_a(N)$ and $\Psi_a(w, N)$ denote the set of all δ , for which $\Phi_a(N, \delta) \neq \emptyset$ and $N^\top(g(a) + \delta\lambda(a) - w) \geq 0$, respectively, are satisfied. We begin by showing that $\mathcal{J}_a(N)$ is closed and convex, hence so is

$$\Psi_a^\varepsilon(w, N, r, \mathcal{D}(w)) = \mathcal{J}_a(N) \cap \Psi_a(w, N) \cap \mathcal{D}(w)^Y \cap H_\varepsilon(N)^Y$$

as intersection of such sets, where $H_\varepsilon(N) = \bigcap_{\tilde{N} \in B_\varepsilon(N)} H(\tilde{N})$. Indeed, let $\delta_1, \delta_2 \in \mathcal{J}_a(N)$. Then there exist ϕ_1, ϕ_2 such that $(\delta_j, T\phi_j)$ for $j = 1, 2$ satisfy (2.11) for every $\tilde{a}^i \in \mathcal{A}^i \setminus \{a^i\}$ and $i = 1, 2$. By linearity of (2.11), so does $(\delta_\kappa, T\phi_\kappa)$ for $\kappa \in [0, 1]$, where we set $\delta_\kappa := \kappa\delta_1 + (1 - \kappa)\delta_2$ and $\phi_\kappa := \kappa\phi_1 + (1 - \kappa)\phi_2$. This shows that $\delta_\kappa \in \mathcal{J}_a(N)$, i.e., $\mathcal{J}_a(N)$ is convex. Let $(\delta_n)_{n \geq 0}$ be a sequence in $\mathcal{J}_a(N)$. Then there exists $(\phi_n)_{n \geq 0}$ such that $(\delta_n, T\phi_n)$ satisfies (2.11). Since the inequalities in (2.11) are not strict, $(\lim_{n \rightarrow \infty} \delta_n, T \lim_{n \rightarrow \infty} \phi_n)$ satisfies (2.11), hence $\lim_{n \rightarrow \infty} \delta_n \in \mathcal{J}_a(N)$ and $\mathcal{J}_a(N)$ is closed. Because $\mathcal{D}(w)^Y$ is bounded, $\Psi_a^\varepsilon(w, N, r, \mathcal{D}(w))$ is compact.

For $\phi \in \mathbb{R}^{d_c}$, introduce the auxiliary sets $\mathcal{J}_a(N, \phi)$ of those $\delta \in \mathbb{R}^{2m}$, for which $(T\phi, \delta)$ enforces a . For $i = 1, 2$, let $a_1^i, \dots, a_{m_i}^i$ be an enumeration of $\mathcal{A}^i \setminus \{a^i\}$ and denote $\Delta\mu_{j_i}^i := \mu(a) - \mu(a_{j_i}^i, a^{-i})$ and $\Delta\lambda_{j_i}^i := \lambda(a) - \lambda(a_{j_i}^i, a^{-i})$ as in the proof of Lemma 4.7.3. Then $\mathcal{J}_a(N, \phi)$ is a closed convex polytope, whose hyperfaces have normal vectors

$$\left(\begin{array}{c} \Delta\lambda_{j_1}^1 \\ 0 \end{array} \right), \quad j_1 = 1, \dots, m_1, \quad \left(\begin{array}{c} 0 \\ \Delta\lambda_{j_2}^2 \end{array} \right), \quad j_2 = 1, \dots, m_2 \quad (4.8)$$

and N only determines the position of these hyperfaces. Thus, similarly as in the proof of Lemma 4.7.3, $N \mapsto \mathcal{J}_a(N, \phi)$ is Lipschitz continuous with a Lipschitz constant that depends only on $\Delta\mu_{j_i}^i$. In particular, the Lipschitz constant of $N \mapsto \mathcal{J}_a(N, \phi)$ is uniformly bounded in ϕ .

Observe that $w \mapsto \Psi_a(w, N)$ is an affine function, and so are the constant maps $w \mapsto H_\varepsilon(N)^Y$ and $w \mapsto \mathcal{J}_a(N)$. Lipschitz continuity in w thus follows from Lemma 4.A.1. For Lipschitz-continuity in N , observe that $H_\varepsilon(N)$ is the intersection of two half-spaces with extremal normal vectors N_-, N_+ in $B_\varepsilon(N) \cap S^1$, where in the case $\varepsilon = 0$, $N_- = N_+ = N$. In particular, $H_\varepsilon(N)^Y$ is a convex polytope with normal vectors $(N_-^1 e_y, N_-^2 e_y)^\top$ and $(N_+^1 e_y, N_+^2 e_y)^\top$, where $e_y \in \mathbb{R}^m$ is the row unit vector in coordinate y . The normal vector $-(N^1 \lambda(a)^\top, N^2 \lambda(a)^\top)^\top$ of $\Psi_a(w, N)$ is a linear combination of all normal vectors of $H_\varepsilon(N)^Y$ for all $N \in S^1$ since $N = k(\varepsilon)(N_+ + N_-)$ for some constant $k(\varepsilon)$. Observe however, that $-(N^1 \lambda(a)^\top, N^2 \lambda(a)^\top)^\top$ is not a linear combination of a proper subset of normal vectors of $H_\varepsilon(N)^Y$ because $\lambda(y|a) > 0$. Therefore, the ranks of the matrices formed by any combination of normal vectors of $H_\varepsilon(N)^Y$ and $\Psi_a(w, N)$ are constant for all $N \in S^1$ as required by Lemma 4.A.2. While $H_\varepsilon(N)^Y$ and $\Psi_a(w, N)$ may fail to be Lipschitz continuous because they are unbounded, their intersection with a constant and bounded polytope is. Let \bar{W} be such a polytope containing $\mathcal{D}(w)^Y$ such that none of its normal vectors are arbitrarily close to being linearly dependent to any $2m - 1$ normal vectors of $H_\varepsilon(N)^Y$, $\Psi_a(w, N)$ or any of the vectors in (4.8). Then $H_\varepsilon(N)^Y \cap \bar{W}$ and $\Psi_a(w, N) \cap \bar{W}$ satisfy the conditions in Lemma 4.A.2 and hence $N \mapsto H_\varepsilon(N)^Y \cap \Psi_a(w, N) \cap \bar{W}$ is Lipschitz continuous. For any $\phi \in \mathbb{R}^{d_c}$, $\mathcal{J}_a(N, \phi)$ is a polytope with normal vectors in (4.8). Let N be differ-

ent from a coordinate direction and suppose that ε is such that neither N_- or N_+ are coordinate directions. Let $\mathcal{X}(N)$ be any subset of normal vectors to $H_\varepsilon(N)^Y$, $\Psi_a(w, N)$ and $\mathcal{J}_a(N, \phi)$. If there exists a linear combination amongst the vectors in $\mathcal{X}(N)$, then there exists a linear combination also in $\mathcal{X}(\tilde{N})$ for \tilde{N} arbitrarily close to N by multiplying the coefficients by \tilde{N}^i/N^i or \tilde{N}_-/N_-^i and \tilde{N}_+/N_+^i , respectively. Therefore, Lemma 4.A.2 applies and shows that $N \mapsto H_\varepsilon(N)^Y \cap \Psi_a(w, N) \cap \mathcal{J}_a(N, \phi) \cap \bar{\mathcal{W}}$ is locally Lipschitz continuous in N except for coordinate directions. Since $\mathcal{D}(w)^Y \subseteq \bar{\mathcal{W}}$ and the intersection of a Lipschitz continuous map with a convex set is Lipschitz continuous, it follows that for any $\phi \in \mathbb{R}^{d_c}$, $N \mapsto H_\varepsilon(N)^Y \cap \Psi_a(w, N) \cap \mathcal{J}_a(N, \phi) \cap \mathcal{D}(w)^Y$ is Lipschitz continuous. Local Lipschitz continuity of $N \mapsto \Psi_a^\varepsilon(w, N, r, \mathcal{D})$ now follows from the fact that the arbitrary union of Lipschitz continuous maps with uniformly bounded Lipschitz constants is Lipschitz again. \square

So far we have shown that (4.7) is locally Lipschitz continuous for almost every $N \in S^1$, where $\phi(a, N, \delta)$ is well defined and bounded away from 0. Define

$$E_a^\varepsilon(r, \mathcal{D}) := \{(w, N) \in \mathbb{R}^2 \times S^1 \mid \Psi_a^\varepsilon(w, N, r, \mathcal{D}) \neq \emptyset\}$$

$$\Gamma_a^\varepsilon(r, \mathcal{D}) := \{(w, N) \in \mathbb{R}^2 \times S^1 \mid \exists \delta \in \Psi_a^\varepsilon(w, N, r, \mathcal{D}) \text{ with } \phi(a, N, \delta) = 0\}$$

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

Lemma 4.7.6. *Let $\varepsilon \geq 0$ and let \mathcal{D} be of class B. If a sequence $(w_n, N_n)_{n \geq 0}$ converges to $(w, N) \notin \mathcal{P}$ such that $\Psi_a^\varepsilon(w_n, N_n, r, \mathcal{D}) \neq \emptyset$ for all $n \geq 0$, then $\Psi_a^\varepsilon(w, N, r, \mathcal{D}) \neq \emptyset$.*

Proof. Let $\delta_n \in \Psi_a^\varepsilon(w_n, N_n, r, \mathcal{D})$. Because $\mathcal{D}(w_n)$ is uniformly bounded, so is δ_n , hence $(\delta_n)_{n \geq 0}$ converges along a subsequence $(n_k)_{k \geq 0}$ to some finite limit δ with $N^\top(g(a) + \delta\lambda(a) - w) \geq 0$ and $\tilde{N}^\top \delta(y) \leq 0$ for every $\tilde{N} \in B_\varepsilon(N) \cap S^1$ and $y \in Y$. Since \mathcal{D} is closed-valued and Lipschitz continuous, $\delta(y) \in \mathcal{D}(w)$ for every $y \in Y$. It remains to show that $\Phi_a(N, \delta) \neq \emptyset$. If the converse is true, closedness of $\mathcal{I}_a^i(N, \delta^i)$ for $i = 1, 2$ implies that $\mathcal{I}_a^1(N, \delta^1)$ and $\mathcal{I}_a^2(N, \delta^2)$ are strictly separated. By continuity, $\mathcal{I}_a^1(N_{n_k}, \delta_{n_k}^1)$ and $\mathcal{I}_a^2(N_{n_k}, \delta_{n_k}^2)$ are separated as well for k sufficiently large, a contradiction. \square

Corollary 4.7.7. *For any $a \in \mathcal{A}$ and $\varepsilon \geq 0$, $E_a^\varepsilon(r, \mathcal{D}) \cup \mathcal{P}$ and $\Gamma_a^\varepsilon(r, \mathcal{D})$ are closed. Therefore, so is $\Gamma^\varepsilon(r, \mathcal{D})$.*

Proof. Indeed, $\Gamma_a^\varepsilon(r, \mathcal{D})$ is closed since $0 \in \Phi_a(N, \delta)$ for some $N \in S^1$ if and only if $0 \in \Phi_a(N, \delta)$ for all $N \in S^1$. \square

Proposition 4.7.8. *For $\varepsilon \geq 0$ and $w \mapsto \mathcal{D}(w)$ of class B,*

$$\kappa(w, N) = \max_{a \in \mathcal{A}} \max_{\delta \in \Psi_a^\varepsilon(w, N, r, \mathcal{D})} \frac{2N^\top(g(a) + \delta\lambda(a) - w)}{r\|\phi(a, N, \delta)\|^2} \quad (4.9)$$

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

Proof. We first show local Lipschitz continuity of κ_a in (4.7) for fixed $a \in \mathcal{A}$. Suppose first that $(w, N) \in E_a^\varepsilon(r, \mathcal{D}) \setminus (\Gamma_a^\varepsilon(r, \mathcal{D}) \cup \mathcal{P})$, i.e., N is not a coordinate direction. Since $\Gamma_a^\varepsilon(r, \mathcal{D})$ is closed by Corollary 4.7.7, there exists an open neighbourhood U of (w, N) bounded away from $\Gamma_a^\varepsilon(r, \mathcal{D}) \cup \mathcal{P}$. Therefore, $\inf_{N, \delta} \|\phi(a, N, \delta)\| \geq c$ and hence the function that is maximized in the right-

hand side of (4.7) is Lipschitz continuous on U by Lemma 4.7.3. It follows that κ_a is Lipschitz continuous by Lemmas 4.7.4 and 4.7.5. Because (4.9) is the maximum over finitely many functions κ_a , it is Lipschitz continuous except where (w, N) leaves the domain of the maximal function κ_a . \square

When we refer to a solution to (4.9), we will always mention explicitly with respect to what ε and which map \mathcal{D} (4.9) is being solved.

Lemma 4.7.9. *Let $a \notin \mathcal{A}^N$ have the unique best response property for player i with $\kappa_a(w, N) > 0$ for $(w, N) \in \mathcal{P}$. Then for every $\varepsilon \in (0, \kappa_a(w, N))$ there exists a neighbourhood \mathcal{U} of (w, N) such that $\kappa_a(\tilde{w}, \tilde{N}) \geq \varepsilon$ for any $(\tilde{w}, \tilde{N}) \in \mathcal{U}$.*

Proof. Let $N = \pm e_i$ and let β be the vector with minimal length such that $(\beta, 0)$ enforces a . Observe that such a vector exists by Lemma 4.5.2. Moreover, $\beta \neq 0$ because $a \notin \mathcal{A}^N$. Since $0 \in \text{int } \mathcal{I}_a^i(0)$, it follows that $(\beta/\tilde{T}^{-i}, 0)$ restricted-enforces a at (\tilde{w}, \tilde{N}) in a neighbourhood of (w, N) . Let K be the Lipschitz constant of

$$\tilde{\kappa}_a(\tilde{w}, \tilde{N}) := \frac{2\tilde{N}^\top (g(a) - \tilde{w})}{r\|\beta\|^2} (\tilde{T}^{-i})^2.$$

Because $\varepsilon \leq \tilde{\kappa}_a(\tilde{w}, \tilde{N}) \leq \kappa_a(\tilde{w}, \tilde{N})$ holds for every $(\tilde{w}, \tilde{N}) \in B_{\varepsilon/K}(w, N)$, the neighbourhood $\mathcal{U} = B_{\varepsilon/K}(w, N)$ satisfies the desired properties. \square

4.7.3 Characterization of $\partial\mathcal{B}(r, \mathcal{W})$

We start by showing that $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ is given by the optimality equation. Because the continuous part of the signal is what creates the curvature, these steps are similar in ideas to Sannikov [37]. Some technical bounds on the

provision of incentives and proximity of solutions to (4.9) for different choices of \mathcal{D} are deferred to Section 4.A.2. We start by showing the following slight generalization of Lemma 4.5.1.

Lemma 4.7.10. *For $\varepsilon \geq 0$ and \mathcal{D} of class B, let \mathcal{C} be a C^1 solution to*

$$\kappa(w) = \frac{2N_w^\top (g(a^*(w)) + \delta^*(w)\lambda(a^*(w)) - w)}{r \|\phi(a^*(w), N_w, \delta^*(w))\|^2} \quad (4.10)$$

oriented by $w \mapsto N_w$, where a^, δ^* are such that for every w in the relative interior of \mathcal{C} , the expression on the right-hand side of (4.10) is strictly positive and $\delta^*(w) \in \Psi_{a^*(w)}^\varepsilon(w, N_w, r, \mathcal{D})$. Then the solution W to (2.8) with $A = a^*(W)$, $\delta = \delta^*(W_-)$, $\beta = \phi(A, N_W, \delta)T_W$ and $M \equiv 0$ remains on \mathcal{C} until an endpoint of \mathcal{C} is reached or an 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 open ball around w with radius η . On $B_\eta(w)$, $\partial\mathcal{W}$ admits a local parametrization f in the direction N_w . For any $v \in B_\varepsilon(w)$, define the orthogonal projection $\hat{v} = T_w^\top v$ onto the tangent T_w . Denote by $\pi(v) = (\hat{v}, f(\hat{v}))$ the projection of $v \in B_\eta(w)$ onto $\partial\mathcal{W}$ in the direction N_w .

Let $(W, A, \beta, \delta, Z, (J^y)_{y \in Y}, M)$ be a weak solution to (2.8) with $W_0 = w$ and $M \equiv 0$ such that on $\llbracket 0, \tau \rrbracket$, $A = a^*(\pi(W))$, $\beta = T\phi$, $\delta(y) = \delta^*(\pi(W_-), y)$ for every $y \in Y$, where $\tau := \sigma_1 \wedge \inf\{t \geq 0 \mid W_t \notin B_\eta(w)\}$, σ_1 is the first time, any of the processes $(J^y)_{y \in Y}$ jump and we abbreviated $N = N_{\pi(W)}$, $T = T_{\pi(W)}$ and $\phi = \phi(A, N, \delta)$. Then the solution satisfies (b) in Lemma 2.4.1 up to time τ . Indeed, $\delta \in \Psi_A^\varepsilon(\pi(W), N, r, \mathcal{D})$ a.e. by construction. Since the maximizer of a measurable function is measurable and π is measurable, A, β

and δ are predictable. Moreover, because δ^* is bounded due to the condition that $W + r\delta^*(y) \in \mathcal{W}$ for every $y \in Y$ and ϕ is a Lipschitz-continuous function of δ^* , they are both square-integrable.

We measure the distance of W to \mathcal{C} by $D_t = N_w^\top W_t - f(\hat{W}_t)$ as illustrated in Figure 3.7. 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 the Meyer-Itô formula (see Theorem 19.5 in Kallenberg [24]) that

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

Note that $N^\top T = 0$ and $\Delta J^y \equiv 0$ for any $y \in Y$ on $\llbracket 0, \sigma_1 \rrbracket$. Using (4.10) and the fact that $N_w^\top N_t = T_w^\top T_t = \ell_t^{-1}$, we obtain that on $\llbracket 0, \tau \rrbracket$,

$$\begin{aligned} dD_t &= r\ell_t N_t^\top (W_t - g(A_t) - \delta_t \lambda(A_t)) dt + \frac{r^2}{2} \kappa(\pi(W_t)) \ell_t^3 |T_w^\top T_t|^2 |\phi_t|^2 dt \\ &= rD_t dt, \end{aligned}$$

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 $D_t = D_0 e^{rt}$, which is identically zero because $D_0 = 0$. On $\{\tau < \sigma_1\}$ we can repeat this procedure and concatenate the solutions to obtain a solution to (2.8) that remains on \mathcal{C} until either an accident occurs or an endpoint of \mathcal{C} is reached. Let ρ denote the hitting time of an endpoint of \mathcal{C} . Then $D_0 = 0$ on $\llbracket 0, \rho \wedge \sigma_1 \rrbracket$ implies that $\pi(W) = W$ and hence $\delta \in \Psi_A^\varepsilon(W, N_W, r, \mathcal{D})$. \square

Corollary 4.7.11. *For $\varepsilon \geq 0$ and \mathcal{D} of class B, let \mathcal{C} be a C^1 solution to (4.9) with positive curvature throughout. Then any payoff in the relative interior of \mathcal{C} is attainable by a solution to (2.8) with $\delta \in \Psi_A^\varepsilon(W, N_W, r, \mathcal{D})$ such that W remains on \mathcal{C} until either an endpoint of \mathcal{C} is reached or an event occurs.*

Proof. For any $w \in \mathcal{C}$, let $a^*(w)$ and $\delta^*(w)$ denote the maximizers in (4.9). Since \mathcal{C} is assumed to have positive curvature, the maximization in (4.9) is not taken over empty sets. By Corollary 4.7.7, the maximizers are attained. \square

The following two lemmas establish that locally, $\partial\mathcal{B}(r, \mathcal{W})$ coincides with a solution to (4.4) at almost every point outside $\mathcal{G}(r, \mathcal{W})$. Lemma 4.7.12 states that it is impossible for a solution to (4.4) to cut through $\mathcal{B}(r, \mathcal{W})$.

Lemma 4.7.12. *Let $\varepsilon > 0$ and let $w \in \partial\mathcal{B}_\varepsilon(r, \mathcal{W})$ with outward normal N' . Define the projection $\pi : B_{\varepsilon/2}(w) \rightarrow \partial\mathcal{B}_\varepsilon(r, \mathcal{W})$ onto $\partial\mathcal{B}_\varepsilon(r, \mathcal{W})$ in the direction of N' and set*

$$\mathcal{D}(w) := \{\delta \in \mathbb{R}^2 \mid \exists \kappa \in [0, 1] \text{ with } \kappa w + (1 - \kappa)\pi(w) + r\delta \in \mathcal{W}\}. \quad (4.11)$$

It is impossible for a C^1 solution \mathcal{C} to (4.9) with $(0, \mathcal{D})$ oriented by $v \mapsto N_v$ to simultaneously satisfy

1. $\mathcal{C} \cap \mathcal{B}_\varepsilon(r, \mathcal{W}) \subseteq B_{\varepsilon/2}(w)$,
2. $\inf_{v \in \mathcal{C}} N_v^\top N' > 0$,
3. $\mathcal{N}_\mathcal{C} \cap (\Gamma^0(r, \mathcal{D}) \cup \mathcal{P}) = \emptyset$,
4. for any $a \in \mathcal{A}$, $\mathcal{N}_\mathcal{C} \cap \partial E_a^0(r, \mathcal{D}) = \emptyset$,
5. there exists $v_0 \in \mathcal{C}$ such that $v_0 + \eta N' \in \mathcal{B}_\varepsilon(r, \mathcal{W})$ for some $\eta > 0$.

Proof. Suppose towards a contradiction that there exists such a curve \mathcal{C} . Observe that \mathcal{D} is of class B, hence it follows from Conditions 3 and 4 as well as Proposition 4.7.8 that \mathcal{C} is C^2 at almost every point. By Condition 2, there exists a local parametrization f of \mathcal{C} in the direction N' . Define the orthogonal projection $\hat{v} = T'^\top v$ onto the tangent for any $v \in B_{\varepsilon/2}(w)$, where T' is the counterclockwise rotation of N' by 90° . Denote by $\hat{\pi}(v) = (\hat{v}, f(\hat{v}))$ the projection of $v \in \mathcal{B}_{\varepsilon/2}(w)$ onto \mathcal{C} in the direction N' .

By definition of $\mathcal{B}_\varepsilon(r, \mathcal{W})$, there exists a solution $(W, A, \beta, \delta, M, Z, (J^y)_{y \in Y})$ to (2.8) with $W_0 = v_0 + \eta N'$ such that on the set $\llbracket 0, \sigma_1 \rrbracket$, (β, δ) enforces A , $W + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ and $N^\top \delta(y) \leq 0$ for every normal vector N to $\partial \mathcal{B}_\varepsilon(r, \mathcal{W})$ at $v \in \partial \mathcal{B}_\varepsilon(r, \mathcal{W}) \cap B_\varepsilon(W)$ and every $y \in Y$. Define the stopping time $\tau_1 := \inf\{t \geq 0 \mid W_t \notin B_{\varepsilon/2}(w)\}$. Condition 1 together with convexity implies that \mathcal{C} intersects $\partial \mathcal{B}_\varepsilon(r, \mathcal{W})$ at two points v_L, v_R . Since any two points in $B_{\varepsilon/2}(w)$ are within distance ε of each other, it follows that $N^\top \delta(y) \leq 0$ on $\llbracket 0, \tau_1 \rrbracket$ for every normal vector N to $\partial \mathcal{B}_\varepsilon(r, \mathcal{W})$ between v_L and v_R . Since \mathcal{C} cuts through $\mathcal{B}_\varepsilon(r, \mathcal{W})$ by Condition 5, convexity implies that the set of normal vectors to \mathcal{C} between v_L and v_R is a subset of the normal vectors to $\mathcal{B}_\varepsilon(r, \mathcal{W})$ between these two points. Therefore, $N^\top \delta(y) \leq 0$ also for any normal vector N to \mathcal{C} between v_L and v_R on $\llbracket 0, \tau_1 \rrbracket$.

Suppose first that $\mathcal{N}_\mathcal{C} \subseteq E_a^0(r, \mathcal{D})$ for some $a \in \mathcal{A}$, i.e., \mathcal{C} is a non-trivial solution to (4.9). Let $N_t := N_{\hat{\pi}(W_t)}$ and $T_t := T_{\hat{\pi}(W_t)}$ and observe that these projections are well defined on $\llbracket 0, \tau_1 \rrbracket$. We measure the distance of W to \mathcal{C} by $D_t = N'^\top W_t - f(\hat{W}_t)$. Denote $\ell_t := 1/T_t^\top T'$ and $\gamma_t := \ell_t N_t^\top T'$ for the sake of brevity and observe that $\bar{\gamma} := \sup_{w \in \mathcal{C}} N_w^\top T' / T_w^\top T' < \infty$ by Condition 5. Then,

similarly as in Footnote 3 of Hashimoto [21], it follows from Itô's formula that

$$D_t \geq D_0 + \int_0^t \zeta_s \, ds + \int_0^t \xi_s (dZ_s - \mu(A_s) \, ds) + \sum_{y \in Y} \int_0^t \rho_s(y) \, dJ_s^y + \tilde{M}_t,$$

where

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

$\xi_t = r\ell_t N_t^\top \beta_t$, $\rho_t(y) = r\ell_{t-} N_{t-}^\top \delta_t(y)$ and $\tilde{M}_t = \int_0^t r\ell_{t-} N_{t-}^\top \, dM_t$. Define the stopping time $\tau_2 := \inf\{t \geq 0 \mid D_t \leq 0\}$ and observe that $\tau_2 \leq \tau_1$ a.s. because Condition 2 implies that $v_L + \eta N' \notin \mathcal{B}_\varepsilon(r, \mathcal{W})$ for any $\eta > 0$ and similarly for v_R . We will show that there exists an equivalent probability measure R such that the drift rate of D_t is bounded from below by rD_t . Then D_t becomes arbitrarily large with positive R -probability, and hence positive Q^A -probability. Because it may take arbitrarily long until an accident arrives, this leads to a contradiction since \mathcal{V} is bounded.

Let Ξ_1 denote the set where $N^\top(\hat{\pi}(W) - g(A) - \delta\lambda(A)) \geq 0$. On Ξ_1 , $\xi_t \geq rD_t$, hence there is no need to change the probability measure. It follows from Condition 2 that $\beta \neq 0$ on Ξ_1^c . Let $\Xi_2 \subseteq \Xi_1^c$ be the set where $\mathcal{N}_{\hat{c}} \subseteq E_A^0(r, \mathcal{D})$, i.e., $\Psi_A^0(\hat{\pi}(W), N, r, \mathcal{D}) \neq \emptyset$. Set

$$\hat{\delta} \in \arg \min_{\Psi_A^0(\hat{\pi}(W), N, r, \mathcal{D})} \|\hat{\delta}^1 - \delta^1\| + \|\hat{\delta}^2 - \delta^2\|,$$

then (4.9) implies that

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

Denote $\Lambda := \max_{a \in \mathcal{A}} \sum_{y \in Y} \lambda(y|a)$ and observe that $N^\top(g(A) + \hat{\delta}\lambda(A) - \hat{\pi}(W))$ is uniformly bounded above by the constant $K_1 := \text{diam } \mathcal{V} + \sup(\mathcal{W} - \mathcal{V})\Lambda < \infty$. The condition that $W + r\delta(y) \in \mathcal{W}$ implies that $\delta(y) \in \mathcal{D}(W)$ on $\llbracket 0, \tau_2 \rrbracket$ for every $y \in Y$. Since $N^\top\delta(y) \leq 0$ holds by the choice of ε , Lemma 4.A.3 asserts the existence of constants $K_2, \bar{\Psi}$ such that

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

On the set $\Xi_1^c \cap \Xi_2^c$, condition 3 implies that $\mathcal{N}_{\hat{c}}$ is bounded away from $E_A^0(r, \mathcal{D}) \cup \mathcal{P}$ by virtue of Corollary 4.7.7. Lemma 4.A.4 thus implies that $\|N^\top\beta\| \geq K_4$ for some constant K_4 and hence

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

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

$$\frac{dL_t}{L_t} = \psi_t dZ_t + \sum_{y \in Y} \left(\frac{1}{\lambda(y|A_{t-})} - 1 \right) (dJ_t^y - \lambda(y|A_{t-}) dt),$$

where

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

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

$$D_t \geq D_0 + \int_0^t r D_s ds + \int_0^t \xi_s^\top dZ'_s + \tilde{M}_t + \sum_{y \in Y} \int_0^t \rho_s(y) dJ_s^y. \quad (4.12)$$

Since W is bounded, $\int_0^\cdot \xi_s^\top dZ_s$ is a $BMO(Q_T^A)$ -martingale. Hence, $\int_0^\cdot \xi_s^\top dZ'_s$ is a $BMO(R)$ -martingale by Theorem 3.6 in Kazamaki [28]. Define the stopping time $\tau_3 := \inf\{t \geq 0 \mid D_t \leq D_0(1 + rt)/2\} \leq \tau_2 \wedge T$. It follows from (4.12) that

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

where $F_t = \int_0^t \xi_s dZ'_s + \tilde{M}_t$ is an R -martingale starting at 0. Define the R -martingale $G_t := e^{|Y|t} 1_{\{t < \sigma_1\}}$ and observe that G is orthogonal to F . Because $\tau_3 \leq T$ a.s.,

$$\begin{aligned} 0 &\geq \mathbb{E}_R \left[\left(D_{\tau_3} - \frac{D_0}{2}(1 + r\tau_3) \right) 1_{\{T < \sigma_1\}} \right] \geq \mathbb{E}_R \left[\frac{D_0}{2} 1_{\{T < \sigma_1\}} + F_{\tau_3} 1_{\{T < \sigma_1\}} \right] \\ &= \frac{D_0}{2} R(T < \sigma_1) + e^{-mT} \mathbb{E}_R[F_{\tau_3} G_T] > 0, \end{aligned}$$

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

Suppose now that $\mathcal{N}_C \subseteq \bigcap_{a \in \mathcal{A}} E_a^0(r, \mathcal{D})^c$, i.e., \mathcal{C} is a straight line segment. Let D denote the distance of W to \mathcal{C} in the direction of the normal vector N_C




Figure 4.12: Construction of a curve \mathcal{C}' that cuts through $\mathcal{B}(r, \mathcal{W})$ and hence also through $\mathcal{B}_\varepsilon(r, \mathcal{W})$ for sufficiently small $\varepsilon > 0$.

of \mathcal{C} . Condition 2 makes it possible to apply Lemma 4.A.4, hence any (β, δ) enforcing A it follows that $\|N^\top \beta\| \geq K$ for some constant K . Similarly as before, the drift of D_t is thus bounded from below by $rD_t - K_1/Kr\ell_t\|N^\top \beta_t\|$. Therefore, there exists an equivalent probability measure under which D grows arbitrarily large with positive probability, a contradiction. \square

Lemma 4.7.12 shows that a solution \mathcal{C} to (4.4) cannot escape $\text{cl } \mathcal{B}(r, \mathcal{W})$. Indeed, if it did, there would exist \mathcal{C}' close to \mathcal{C} as indicated in Figure 4.12 that cuts through $\mathcal{B}_\varepsilon(r, \mathcal{W})$ for small ε . As motivated in Section 4.1, \mathcal{C} cannot fall into the interior of $\mathcal{B}(r, \mathcal{W})$ either, hence $\partial \mathcal{B}(r, \mathcal{W})$ is given by (4.4).

Lemma 4.7.13. *Fix $w \in \mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ with outward normal N , where (4.4) is locally Lipschitz continuous. Then $\partial \mathcal{B}(r, \mathcal{W})$ coincides with a solution to (4.4) in a neighbourhood of (w, N) .*

Proof. We first show that a solution to (4.9) with $\varepsilon = 0$ and $\mathcal{D}(w)$ given as in (4.11) coincides with $\partial \mathcal{B}(r, \mathcal{W})$, which implies that also a solution to (4.4) stays on $\partial \mathcal{B}(r, \mathcal{W})$. In a sufficiently small neighbourhood of (w, N) , (4.9) admits a unique C^2 solution that is continuous in initial values. Let \mathcal{C} be solution with initial value (w, N) and suppose towards a contradiction that

\mathcal{C} escapes $\text{cl } \mathcal{B}(r, \mathcal{W})$ in a neighbourhood of w . Then we can change initial conditions slightly to obtain a curve \mathcal{C}' that cuts through $\mathcal{B}(r, \mathcal{W})$, and hence also through $\mathcal{B}_\varepsilon(r, \mathcal{W})$ for ε sufficiently close. Specifically:

- If $\partial\mathcal{B}(r, \mathcal{W})$ is not C^1 at w , we obtain \mathcal{C}' as a solution to (4.4) with initial conditions $(w - \eta N, N)$ for $\eta > 0$ sufficiently small.
- If $\partial\mathcal{B}(r, \mathcal{W})$ is C^1 at w , we obtain \mathcal{C}' for initial conditions (w, N') , where N' is a slight rotation of N as illustrated in the left panel of Figure 4.12.

Because the set where (4.9) fails to be locally Lipschitz continuous is closed by Corollary 4.7.7 and Proposition 4.7.8, a small enough perturbation satisfies that $\mathcal{N}_{\mathcal{C}'} \cap \Gamma_a^0(r, \mathcal{D}) = \emptyset$ for every $a \in \mathcal{A}$ and either $\mathcal{N}_{\mathcal{C}'} \subseteq E_a^0(r, \mathcal{D})$ or $\mathcal{N}_{\mathcal{C}'} \cap (E_a^0(r, \mathcal{D}) \cup \mathcal{P}) = \emptyset$ for any $a \in \mathcal{A}$, that is, \mathcal{C}' satisfies conditions 3–5 of Lemma 4.7.12. By choosing ε and η or N' suitably, we can get conditions 1 and 2 to hold as well, and hence \mathcal{C}' is impossible due to Lemma 4.7.12. We conclude that $\partial\mathcal{B}(r, \mathcal{W})$ is C^1 where (4.9) is locally Lipschitz continuous and that a solution to (4.9) cannot escape $\text{cl } \mathcal{B}(r, \mathcal{W})$.

Suppose towards a contradiction that \mathcal{C} falls into the interior of $\mathcal{B}(r, \mathcal{W})$ in a neighbourhood of (w, N) , that is, there exists $v \in \mathcal{C} \cap \text{int } \mathcal{B}(r, \mathcal{W})$ arbitrarily close to w . By convexity of $\mathcal{B}(r, \mathcal{W})$, this is not possible if \mathcal{C} is a trivial solution to (4.9), 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 Figure 4.13. Let v be close enough to w such that (4.9) with $(0, \mathcal{D})$ is Lipschitz continuous on a neighbourhood of $\mathcal{N}_{\mathcal{C}} := \{(\tilde{w}, N_{\tilde{w}}) \mid \tilde{w} \in \mathcal{C} \text{ between } w \text{ and } v\}$. Let $\delta > 0$ such that the closed ball $B_\delta(v)$ is contained in the interior of $\mathcal{B}(r, \mathcal{W})$. Then for $\varepsilon > 0$ small enough, $B_\delta(v) \subseteq \text{int } \mathcal{B}_\varepsilon(r, \mathcal{W})$ and it follows from Re-




Figure 4.13: If \mathcal{C} falls into the interior of $\mathcal{B}(r, \mathcal{W})$, then there exists a solution \mathcal{C}'_ε to (4.6) with initial conditions (w, N') and a slightly reduced set of available controls such that \mathcal{C}'_ε escapes $\mathcal{B}(r, \mathcal{W})$. For small ε and N' close to N , there exists an enforceable strategy profile attaining $w' \notin \mathcal{B}(r, \mathcal{W})$, whose continuation value reaches $\mathcal{B}_\varepsilon(r, \mathcal{W})$ with certainty. This leads to a contradiction.

mark 4.7.1 that there exists $\eta(\varepsilon)$ such that for any $\tilde{w} \in \mathcal{C}$, $N_{\tilde{v}} \in B_{\eta(\varepsilon)}(N_{\tilde{w}})$ for any $\tilde{v} \in \mathcal{C} \cap B_\varepsilon(\tilde{w})$. Note that $\eta(\varepsilon)$ can be made arbitrarily small by choosing small ε , hence for ε sufficiently small, (4.9) with $(\eta(\varepsilon), \mathcal{D})$ is Lipschitz continuous on an open neighbourhood of $\mathcal{N}_\mathcal{C}$. For $\zeta > 0$ to be chosen later, let $\mathcal{W}_\zeta := \{w \in \mathcal{W} \mid d(w, \partial\mathcal{W}) \leq \zeta\}$, where $d(w, \partial\mathcal{W})$ denotes the minimal distance of d from $\partial\mathcal{W}$. Set

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

where π is the projection onto $\partial\mathcal{B}(r, \mathcal{W})$ in the direction N . Observe that for ζ sufficiently small, (4.9) with $(\eta(\varepsilon), \mathcal{D}_\zeta)$ is Lipschitz continuous in a neighbourhood of $\mathcal{N}_\mathcal{C}$, hence it admits a unique solution \mathcal{C}_ε . Choose now ε and ζ small enough such that Lemma 4.A.5 asserts the existence of $v' \in \mathcal{C}_\varepsilon \cap B_\delta(v)$.

Because \mathcal{C}_ε is continuous in initial conditions, a solution \mathcal{C}'_ε to (4.9) with $(\eta(\varepsilon), \mathcal{D}_\zeta)$ for a slight rotation N' of N reaches a neighbourhood of v' in $\mathcal{B}_\varepsilon(r, \mathcal{W})$. As illustrated in Figure 4.13, \mathcal{C}'_ε will escape $\text{cl } \mathcal{B}(r, \mathcal{W})$ to the right of w and enter $\mathcal{B}_\varepsilon(r, \mathcal{W})$ to the left of w . Thus, for N' close enough to N , there

exist $v_L, v_R \in \mathcal{C}'_\varepsilon \cap \mathcal{B}_\varepsilon(r, \mathcal{W})$, such that $\|\tilde{w} - \pi(\tilde{w})\| \leq \zeta$ for all $\tilde{w} \in \mathcal{C}'_\varepsilon$. By Corollary 4.7.11, for any $w' \in \mathcal{C}'_\varepsilon$ there exists a solution to (2.8) with $W_0 = w'$ such that $\delta \in \Psi_A^{\eta(\varepsilon)}(W, N, r, \mathcal{D}_\zeta)$ on $\llbracket 0, \sigma_1 \rrbracket$ and $W \in \mathcal{C}'_\varepsilon$ until it reaches an end point of \mathcal{C}'_ε or an event occurs. Let $\tau := \inf\{t \geq 0 \mid t \in \{v_L, v_R\}\}$ and observe that $W_\tau \in \mathcal{B}_\varepsilon(r, \mathcal{W})$ on $\{\tau < \sigma_1\}$. The condition that $\delta(y) \in \mathcal{D}_\zeta(W)$ a.e. for every $y \in Y$ implies that $x + r\delta_t(y) \in \mathcal{W}_\zeta$ for some x between W_t and $\pi(W_t)$. On $\llbracket 0, \tau \wedge \sigma_1 \rrbracket$ it holds that $\|W_t - x\| \leq \zeta$, and hence $\delta \in \Psi_A^{\eta(\varepsilon)}(W, N_W, r, \mathcal{D})$. Because $W_\tau \in \mathcal{B}_\varepsilon(r, \mathcal{W})$ on $\{\tau < \sigma_1\}$, by definition of $\mathcal{B}_\varepsilon(r, \mathcal{W})$ there exists a solution $(\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta})$ with $\tilde{W}_0 = W_\tau$ such that on $\llbracket 0, \sigma_1 \rrbracket$, $(\tilde{\beta}, \tilde{\delta})$ enforces $\tilde{A}, \tilde{\delta}(y) \in \mathcal{D}(\tilde{W})$ and $N^\top \tilde{\delta}(y) \leq 0$ for every normal vector N at $\partial\mathcal{B}_\varepsilon(r, \mathcal{W})$ sufficiently close to \tilde{W} . Therefore, a concatenation of (W, A, β, δ) with $\tilde{W}, \tilde{A}, \tilde{\beta}, \tilde{\delta}$ satisfies the same properties, which shows that $\text{co}\mathcal{C}'_\varepsilon \cup \mathcal{B}_\varepsilon(r, \mathcal{W}) \subseteq \mathcal{B}_\varepsilon(r, \mathcal{W})$ by maximality of $\mathcal{B}_\varepsilon(r, \mathcal{W})$. This is a contradiction to $w' \notin \text{cl}\mathcal{B}(r, \mathcal{W})$. \square

Finally, because (4.4) is Lipschitz continuous almost everywhere, we need to show that $\partial\mathcal{B}(r, \mathcal{W})$ is C^1 to grant uniqueness of the solution. By convexity, $\mathcal{B}(r, \mathcal{W})$ cannot have inward corners, and it follows from another escaping argument that it cannot have outward corners outside of $\mathcal{G}(r, \mathcal{W})$ either.

Lemma 4.7.14. *$\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ is C^1 where (4.4) fails to be Lipschitz continuous. Moreover, outside of \mathcal{P} , the set of all points in $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$, where (4.4) fails to be Lipschitz continuous, has relative measure 0.*

Proof. We first check that $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ is C^1 . Because of Lemma 4.7.13, it is enough to verify this property at payoffs w in $\partial\mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ where (4.4) fails to be locally Lipschitz continuous. Observe that $w \notin \mathcal{G}(r, \mathcal{W})$ implies that $(w, N) \notin \Gamma^0(r, \mathcal{D})$ for any outward normal $N \in \mathcal{N}_w(\mathcal{B}(r, \mathcal{W}))$, where

$\mathcal{D}(w) = (\mathcal{W} - w)/r$. Since $\Gamma^0(r, \mathcal{D})$ is closed by Corollary 4.7.7, it follows that $(v, N) \notin \Gamma^0(r, \mathcal{D})$ for any $N \in \mathcal{N}_w(\mathcal{B}(r, \mathcal{W}))$ and any v in a sufficiently small open neighbourhood U of w . Because the boundary of a Borel set has measure 0, it follows that for an arbitrarily small neighbourhood U of w , there exists (v, N) in the set $U \cap \text{int } \mathcal{B}(r, \mathcal{W}) \times \text{int } \mathcal{N}_w(\mathcal{B}(r, \mathcal{W})) \setminus \{\pm e_i\}$ of positive measure, for which (4.4) is locally Lipschitz continuous.

For sufficiently small $\varepsilon > 0$, let $(w_\varepsilon, N') \in \mathcal{N}_{\mathcal{B}_\varepsilon(r, \mathcal{W})}$ be such a pair. Because the points where (4.4) fails to be Lipschitz continuous is closed, (4.4) is also Lipschitz continuous in a small neighbourhood U_ε of (w_ε, N') . Let \mathcal{D}_ε be defined as in (4.11) with respect to N' . By the definition of \mathcal{D}_ε it follows that $(v, N) \notin \bigcup_{a \in \mathcal{A}} \partial E_a^0(r, \mathcal{D}_\varepsilon) \cup \Gamma^0(r, \mathcal{D}_\varepsilon) \cup \mathcal{P}$ for any $(v, N) \in U_\varepsilon$. Choose such a pair (v, N) close enough to (w_ε, N') and let \mathcal{C} be a solution to (4.9) with $(0, \mathcal{D}_\varepsilon)$ starting in (v, N) . We have shown that conditions 3 and 4 of Lemma (4.7.12) are satisfied on U_ε for any $\varepsilon > 0$. Since $\mathcal{B}(r, \mathcal{W})$ has a corner at w , the curvature of $\partial \mathcal{B}_\varepsilon(r, \mathcal{W})$ becomes arbitrarily large, hence for ε small enough we can get \mathcal{C} to cut through $\text{int } \mathcal{B}_\varepsilon(r, \mathcal{W})$, i.e., we can get conditions 1, 2 and 5 to hold for (v, N) sufficiently close to (w_ε, N') . Such a curve \mathcal{C} is impossible by Lemma 4.7.12, hence the first statement follows

For the second statement, suppose that there exists $\mathcal{C} \subseteq \partial \mathcal{B}(r, \mathcal{W}) \setminus \mathcal{G}(r, \mathcal{W})$ of positive length. By shortening the line segment we may assume that $\mathcal{N}_\mathcal{C} \subseteq \mathcal{P}$ or $\mathcal{N}_{\text{int } \mathcal{C}} \cap \mathcal{P} = \emptyset$. Suppose towards a contradiction that $\mathcal{N}_{\text{int } \mathcal{C}} \cap \mathcal{P} = \emptyset$. Then Proposition 4.7.8 shows that (w, N) enters and leaves $E_a^0(r, \mathcal{D})$ of the maximizing action profile a at almost every $(w, N) \in \mathcal{N}_\mathcal{C}$. Because \mathcal{A} is finite we may assume that this is the same action profile. This implies that $\mathcal{N}_\mathcal{C} \subseteq \partial E_a^0(r, \mathcal{D})$ and hence $\mathcal{N}_{\text{int } \mathcal{C}} \subseteq E_a^0(r, \mathcal{D})$ by Corollary 4.7.7, a contradiction. \square

The characterization of $\partial\mathcal{B}(r, \mathcal{W})$ is concluded by showing that any C^1 segment in $\mathcal{G}(r, \mathcal{W})$ must lie on a straight line through $g(a)$ for some $a \in \mathcal{A}$.

Lemma 4.7.15. *Let \mathcal{W} have non-empty interior let $w \in \mathcal{G}(r, \mathcal{W})$ with outward normal N . Then $(w, N) \in \Gamma_a^0(r, \mathcal{D})$ implies that $N^\top(g(a) - w) = 0$.*

Proof. It is sufficient to verify the statement for extremal normal vectors as this implies that $g(a) = w$ if $\mathcal{B}(r, \mathcal{W})$ has a corner at w . Let N' be an extremal normal vector at w and let $a \in \mathcal{A}$ with $(w, N') \in \Gamma_a^0(r, \mathcal{D})$. Then there exists δ such that (a, δ) decomposes w . Suppose first that $\mathcal{B}(r, \mathcal{W})$ has empty interior. By convexity of $\mathcal{B}(r, \mathcal{W})$ it follows that $\mathcal{B}(r, \mathcal{W})$ is a straight line segment and hence $-N'$ is an outward normal vector as well. Since the jumps are directed inwards, this implies that $N'^\top \delta(y) = 0$ for every $y \in Y$ and hence $N'^\top(g(a) - w) = 0$. Suppose therefore that $\mathcal{B}(r, \mathcal{W})$ has non-empty interior. Because \mathcal{W} has non-empty interior by assumption, $H(N') \cap \mathcal{D}(w)$ has non-empty interior as well. Therefore, for any (a, δ) decomposing w there exists a small perturbation $\tilde{\delta}$ of δ with $\tilde{\delta}(y) \in \text{int}(H(N') \cap \mathcal{D}(w))$ for every $y \in Y$.

Suppose first that N' is not coordinate and that w is decomposed by (a, δ_0) with $N'^\top(g(a) + \delta_0 \lambda(a) - w) > 0$. Since $\mathcal{B}(r, \mathcal{W})$ has non-empty interior, there exists $d \in \mathbb{R}^{2 \times m}$ with $\delta(\eta; y) := \delta_0(y) + \eta d(y) \in \text{int}(\mathcal{D}(w) \cap H(N'))$ for every $y \in Y$ and $\eta > 0$ sufficiently small. Thus, for small $\eta > 0$, there exists a neighbourhood $\mathcal{U}(\eta)$ of (w, N') such that for every $(v, N) \in \mathcal{U}(\eta)$ it holds that $\delta(\eta) \in \text{int}(\mathcal{D}(v) \cap H(N))$. By making $\mathcal{U}(\eta)$ small enough such that $\mathcal{U}(\eta) \cap \mathcal{P} = \emptyset$, Lemma 4.5.2 implies the existence of $\phi(a, N, \delta(\eta)) \in \mathbb{R}^{d_c}$ such that $(\phi(a, N, \delta(\eta)), \delta(\eta))$ restricted-enforces a in the direction N at v for every $(v, N) \in \mathcal{U}(\eta)$. Observe that $\|\phi(a, N, \delta(\eta))\|$ can be made arbitrarily small

by choosing η small. Since $\mathcal{B}(r, \mathcal{W})$ is convex, $\partial\mathcal{B}(r, \mathcal{W})$ is one-sided C^1 at w , hence it is possible to choose η sufficiently small such that a solution \mathcal{C} to

$$\kappa(v) = \frac{2N_v^\top (g(a) + \delta(\eta)\lambda(a) - v)}{r \|\phi(a, N_v, \delta(\eta))\|^2} \quad (4.13)$$

with initial state (\tilde{w}, N') for \tilde{w} close enough to w intersects $\partial\mathcal{B}(r, \mathcal{W})$ on both sides of \tilde{w} with $\mathcal{N}_{\mathcal{C}} \subseteq \text{int}\mathcal{U}(\eta)$. Denote $\mathcal{B}'_\varepsilon := \text{co}(\mathcal{C} \cup \mathcal{B}_\varepsilon(r, \mathcal{W}))$ and fix $\varepsilon > 0$ small enough such that $\mathcal{N}_{\mathcal{C} \cap B_\varepsilon(\partial\mathcal{B}'_\varepsilon)} \cap \mathcal{N}_{\mathcal{B}_\varepsilon(r, \mathcal{W})}$ is contained in the interior of $\mathcal{U}(\eta)$. We will now show that \mathcal{B}'_ε is $(\varepsilon, r, \mathcal{W})$ -admissible, which contradicts maximality of $\mathcal{B}_\varepsilon(r, \mathcal{W})$. Indeed, by Lemma 4.7.10 any payoff on \mathcal{C} is attainable by a solution to (2.8) that is continuous on $\llbracket 0, \sigma_1 \rrbracket$ until an endpoint is reached. Therefore, the time required for that solution to reach $\mathcal{C} \cap \partial B_\varepsilon(\partial\mathcal{B}'_\varepsilon)$ from any point in $\mathcal{C} \cap B_{\varepsilon/2}(\partial\mathcal{B}'_\varepsilon)$ is bounded from below by the times needed to get there from $\mathcal{C} \cap \partial B_{\varepsilon/2}(\partial\mathcal{B}'_\varepsilon)$, which is strictly positive. Therefore, \mathcal{C} is locally $(\varepsilon, r, \mathcal{W})$ -admissible and together with $\mathcal{B}_\varepsilon(r, \mathcal{W})$ satisfies the conditions in Lemma 4.7.1. This lemma thus implies that \mathcal{B}'_ε is $(\varepsilon, r, \mathcal{W})$ -admissible as well.

Observe that a contradiction can be obtained in the same way if w is decomposed by (a, δ_0) with $N'^\top (g(a) + \delta_0\lambda(a) - w) = 0$ and $N'^\top \delta_0(y) < 0$ for at least one $y \in Y$. Indeed, this implies that $N'^\top (g(a) - w) > 0$ and hence for a suitable choice of d it follows that $N'^\top (g(a) + \delta(\eta)\lambda(a) - w) > 0$ and $\delta(\eta) \in \text{int}\mathcal{D}(w) \cap H(N')$ for all $\eta > 0$ small enough. If, instead, $N'^\top \delta(y) = 0$ for every $y \in Y$ and every δ decomposing w , then $N'^\top (g(a) - w) = 0$.

Finally, let $N' = \pm e_i$ be a coordinate direction. By Assumption 4.2.2 there exists an action profile $\tilde{a} \in \mathcal{A}$ with the unique best response property for player i such that $N'^\top (g(\tilde{a}) - w) \geq 0$. Moreover, due to Lemma 4.5.2,

\tilde{a} is restricted-enforceable in the direction N' by $(\beta, 0)$. Suppose towards a contradiction that $N'^\top(g(\tilde{a}) - w) \neq 0$. Then it follows from the unique best response-property that $(\tilde{\phi}(\tilde{a}, N_v, 0), 0)$ restricted-enforces a in the direction N_v sufficiently close to N' , where $\tilde{\phi}(\tilde{a}, N_v, 0) := \beta/T_v^{-i}$. Moreover, for this choice of $\tilde{\phi}$, (4.13) is locally Lipschitz continuous in a neighbourhood of coordinate directions. Therefore, a solution \mathcal{C} to (4.13) with initial conditions (\tilde{w}, N') for \tilde{w} sufficiently close to w intersects $\partial\mathcal{B}(r, \mathcal{W})$ on both sides of \tilde{w} . The proof is completed by showing that $\text{co } \mathcal{C} \cup \mathcal{B}_\varepsilon(r, \mathcal{W})$ is $(\varepsilon, r, \mathcal{W})$ -admissible in the same way as before, thereby contradicting maximality of $\mathcal{B}_\varepsilon(r, \mathcal{W})$. \square

Corollary 4.7.16. *Any C^1 segment in $\mathcal{G}(r, \mathcal{W})$ is a straight line segment whose infinite continuation goes through $g(a)$ for some $a \in \mathcal{A}$.*

Corollary 4.7.17. *All corners of $\mathcal{B}(r, \mathcal{W})$ are contained in \mathcal{V}^N .*

4.7.4 Closedness of $\mathcal{B}(r, \mathcal{W})$

We first show that the maximization in (4.9) is taken over a non-empty set of controls at all extremal payoffs of $\mathcal{B}(r, \mathcal{W})$. This follows from Corollary 4.7.7 for non-coordinate payoffs, and from the following two lemmas for coordinate payoffs. Denote $\mathcal{D}(w) = (\mathcal{W} - w)/r$ throughout this section.

Lemma 4.7.18. *Let $(w_n, N_n)_{n \geq 0} \subseteq \mathcal{N}_{\mathcal{B}(r, \mathcal{W})}$ converge to $(w, N) \in \mathcal{P}$ with $N_n \neq N$ for every $n \geq 0$, and denote by a the maximizer of $\kappa(w_n, N_n)$ in (4.9) for n sufficiently large. Then $\Psi_a^0(w, N, r, \mathcal{D}) \neq \emptyset$.*

Proof. Assume without loss of generality that $N = \pm e_1$. If $(w, N) \in \Gamma_a^0(r, \mathcal{D})$, then, by definition of $\Gamma_a^0(r, \mathcal{D})$, there exists $(0, \delta)$ restricted-enforcing a , which

establishes the claim. Suppose now that $(w, N) \notin \Gamma_a^0(r, \mathcal{D})$. Since $\Gamma_a^0(r, \mathcal{D})$ is closed by Corollary 4.7.7, the same is true for (w', N') in a neighbourhood of (w, N) . It follows from Lemma 4.7.13 that $\partial\mathcal{B}(r, \mathcal{W})$ is a solution to (4.4) in a neighbourhood of (w, N) . Let $v \in \partial\mathcal{B}(r, \mathcal{W})$ converge to w . Then $N_v \rightarrow N$ as otherwise w would be corner of $\mathcal{B}(r, \mathcal{W})$ and hence $(w, N) \in \Gamma_a^0(r, \mathcal{D})$. Because Ψ_a^0 is compact-valued, it follows that the arg max correspondence is upper hemicontinuous and there exists a continuous choice of maximizer $v \mapsto \delta_v$.

Suppose towards a contradiction that $\Psi_a^0(w, N, r, \mathcal{D}) = \emptyset$. This implies that $0 \notin \mathcal{I}_a^1(\delta^1)$ for every continuous choice $v \mapsto \delta_v$ with limit $\delta := \lim_{v \rightarrow w} \delta_v$. Since $\mathcal{I}_a^1(\delta^1)$ is closed and \mathcal{I}_a^1 is continuous, it follows that $0 \notin \mathcal{I}_a^1(\delta_v^1)$ for v close enough to w . Therefore, the length of any vector in $\frac{1}{T_v^1} \mathcal{I}_a^1(\delta_v^1) \cap \frac{1}{T_v^2} \mathcal{I}_a^2(\delta_v^2)$ converges to ∞ as $v \rightarrow w$ and hence $\kappa_a(v, N_v) \rightarrow 0$. Observe that this implies $\kappa(w, N) = 0$ as otherwise there would exist an action profile \tilde{a} with the unique best response property for player 1 and $N^\top(g(\tilde{a}) - w) > 0$ by Assumption 4.2.2. Therefore, due to Lemma 4.7.9, $\kappa_{\tilde{a}}(v, N_v) > \kappa_a(v, N_v)$ for (v, N_v) in a neighbourhood of (w, N) , contradicting the fact that a is the maximizer.

Since $\frac{1}{T_v^2} \mathcal{I}_a^2(\delta_v^2)$ has finitely many hyperfaces, $\phi(a, N_v, \delta_v)$ remains on the same hyperface for v sufficiently close to w . Therefore, $\|\phi(a, N_v, \delta_v)\|$ behaves like c/T_v^1 , in particular it is convex for v close enough to w . This implies that $1/\|\phi_v\|^2$ is convex as well and, because it is positive and monotonically decreasing for v close to w , Lipschitz continuous. Therefore, κ_a and κ are Lipschitz continuous by Lemma 4.7.4. Solutions to κ are thus unique. However, the solution starting in (w, N) is a straight line, a contradiction to $N_n \neq N$. \square

Lemma 4.7.19. *The endpoints of any coordinate straight line segment in $\partial\mathcal{B}(r, \mathcal{W})$ are in $\mathcal{G}(r, \mathcal{W})$.*

Proof. Suppose that $\partial\mathcal{B}(r, \mathcal{W})$ has a coordinate straight line segment L with end point w outside of $\mathcal{G}(r, \mathcal{W})$. Because $\mathcal{G}(r, \mathcal{W})$ is closed, there exists a subsegment $L' \subseteq L$ of positive length entirely outside of $\mathcal{G}(r, \mathcal{W})$, such that $\partial\mathcal{B}(r, \mathcal{W})$ on the other side of w is a solution \mathcal{C} to (4.4) by Lemma 4.7.13. Lemma 4.7.18 implies that the maximizer a^* of (4.4) as $v \in \mathcal{C}$ approaches w is restricted-enforceable in the direction N . We show that a^* has the unique best response property.

Any action profile a with the property that $(w, N), (w', N) \in E_a^0(r, \mathcal{D})$ for some other $w' \in L'$ also satisfies $(\kappa w + (1-\kappa)w', N) \in E_a^0(r, \mathcal{D})$ for any $\kappa \in [0, 1]$ by convexity of \mathcal{W} and linearity of all other constraints. Since \mathcal{A} is finite it follows that (4.4) is locally Lipschitz continuous on the relative interior of L' by making L' smaller if necessary. Therefore, $\kappa(\hat{w}, N) = 0$ for any \hat{w} in the relative interior of L' by Lemma 4.7.13 and, in particular, $N^\top(g(a) - w) = 0$ for any action profile a with the best response property. Let i be the player for whom L is coordinate and let \hat{a}^i be the best response to $a^{*, -i}$. Assumption 4.2.2 implies that $\hat{a} = (\hat{a}^i, a^{*, -i})$ satisfies $(w, N) \in E_{\hat{a}}^0(r, \mathcal{D})$. Therefore, by Lemma 4.5.2, \hat{a} is also restricted-enforced by $(\beta, 0)$. This shows that $(w', N) \in E_{\hat{a}}^0(r, \mathcal{D})$ also for any other $w' \in L'$ and hence $N^\top(g(\hat{a}) - w) = 0$ by the above argument. Unless $a^* = \hat{a}$, the unique best response property implies $N^\top(g(a^*) - w) < 0$, a contradiction to the fact that a^* is the maximizer in (4.4) as v approaches w . Thus, a^* has the unique best response property.

For $v \in \mathcal{C}$, let δ_v denote the maximizing control in κ_{a^*} at (v, N_v) . Observe that as $v \in \mathcal{C}$ approaches w , N_v approaches N due to Lemma 4.7.14. Therefore, $N_v^\top(g(a^*) - v) \rightarrow 0$ as well and hence $\delta_v^1 \rightarrow 0$. By the unique best response property it follows that $0 \in \text{int } \mathcal{I}_{a^*}^1(0)$ and hence $0 \in \text{int } \mathcal{I}_{a^*}^1(\delta_v^1)$ for v sufficiently

close to w . Therefore, H_{a^*} is equal to

$$\tilde{\kappa}_{a^*}(v, N_v) := \max_{\delta \in \Psi_{a^*}^2(v, N_v, r, \mathcal{D})} \frac{2N_v^\top (g(a^*) + \delta\lambda(a^*) - v)}{r\|\phi(a^*, N, \delta)\|^2}$$

in a sufficiently small neighbourhood of (w, N) , where $\Psi_{a^*}^2(v, N_v, r, \mathcal{D})$ is the set of all δ such that $N_v^\top (g(a^*) + \delta\lambda(a^*) - v) \geq 0$, for every $y \in Y$ it holds that $N_v^\top \delta(y) \leq 0$ and $\delta(y) \in \mathcal{D}(v)$ as well as $\mathcal{I}_{a^*}^2(\delta^2) \neq \emptyset$. It will follow from Lemma 4.7.4 that $\tilde{\kappa}_{a^*}$ is Lipschitz continuous in a neighbourhood of (w, N) once we show that $(v, N_v) \mapsto \Psi_{a^*}^2(v, N_v, r, \mathcal{D})$ is. Lipschitz continuity in v follows from Lemma 4.A.1. For Lipschitz continuity in N_v , write

$$\Psi_{a^*}^2(v, N_v, r, \mathcal{D}) = H_0(N_v)^Y \cap \mathcal{D}(v)^Y \cap \Psi_{a^*}(v, N_v) \cap \mathcal{J}_{a^*},$$

where $\mathcal{J}_{a^*} := \{\delta \mid \mathcal{I}_{a^*}^2(\delta^2) \neq \emptyset\}$ and the other sets are defined as in the proof of Lemma 4.7.5. With the same argument as in the proof of Lemma 4.7.5, it follows that $N_v \mapsto H_0(N_v)^Y \cap \Psi_{a^*}(v, N_v) \cap \bar{\mathcal{W}}$ is Lipschitz continuous for a suitable polytope $\bar{\mathcal{W}}$ containing $\mathcal{D}(v)^Y$. Since both $\mathcal{D}(v)^Y$ and \mathcal{J}_{a^*} are constant and convex, and the intersection with constant and convex sets remains Lipschitz continuous, it follows that $N_v \mapsto \Psi_{a^*}^2(v, N_v, r, \mathcal{D})$ is indeed Lipschitz continuous. This shows that (4.4) is Lipschitz continuous on $L \cup \mathcal{C}$ and hence a solution is unique. This is a contradiction because a solution to (4.4) with initial value (w, N) is a straight line. \square

Corollary 4.7.17 implies that all corners of $\mathcal{B}(r, \mathcal{W})$ are contained in $\mathcal{B}(r, \mathcal{W})$. If the boundary between two corners has strictly positive curvature throughout, then it is contained in $\mathcal{B}(r, \mathcal{W})$ by Corollary 4.7.11. It remains to verify points where the boundary changes from a curve to a straight line segment.

Lemma 4.7.20. *Let $w \in \partial\mathcal{B}(r, \mathcal{W})$ be a point where $\partial\mathcal{B}(r, \mathcal{W})$ changes from a curved solution \mathcal{C} to (4.4) to a straight line segment L in a differentiable way. Then w is in $\mathcal{B}(r, \mathcal{W})$ if the other end points of \mathcal{C} and L are.*

Proof. Let $\varepsilon > 0$ be small enough such that the maximizer a^* in \mathcal{C} does not change in $B_\varepsilon(w) \cap (\mathcal{C} \setminus \{w\})$ and such that $B_\varepsilon(w) \cap \partial\mathcal{B}(r, \mathcal{W})$ admits a parametrization f in the direction N_w . Denote by T_w the tangent vector in the direction of \mathcal{C} and let $\hat{v} := T_w^\top v$ denote the projection onto the tangent. Let $\pi(v) = (\hat{v}, f(\hat{v}))$ denote the projection onto $\partial\mathcal{B}(r, \mathcal{W})$ in the direction N_w . For $v \in B_\varepsilon(w) \cap \mathcal{C}$, let $\delta^*(v)$ denote the maximizer of (4.7) and observe that $\lim_{v \rightarrow w} \delta^*(w) \in \Psi_{a^*}(w, N_w, r, \mathcal{D})$ due to Lemmas 4.7.6 and 4.7.18.

If $w \in \mathcal{G}(r, \mathcal{W})$, then there exists $v \in \mathcal{C}$ arbitrarily close to w with $\|\phi(a^*, N_v, \delta^*(v))\|$ arbitrarily small. Thus a contradiction can be obtained by an enlargement procedure analogous to the proof of Lemma 4.7.15. If $w \notin \mathcal{G}(r, \mathcal{W})$, Lemma 4.7.19 implies that L is a non-coordinate straight line segment. Moreover, $\mathcal{C} \cap B_\varepsilon(w)$ is a solution to H_{a^*} in (4.7), which is Lipschitz continuous on $B_\varepsilon(w) \cap \mathcal{C}$. Solutions are thus unique and hence $N_w^\top (g(a^*) - w) > 0$ as otherwise the solution starting in (w, N_w) would be a straight line. Let $(W, A, \beta, \delta, Z, (J^y)_{y \in Y}, M)$ be a weak solution to (2.8) with initial condition $W_0 = w$ such that $M \equiv 0$ and for all $t \geq 0$, $A_t = a^*$, $\beta_t = T_t \phi_t$ and

$$\delta_t(y) = \begin{cases} \delta^*(\pi(W_{t-}), y) - \frac{1}{r} D_t N_w, & \hat{W}_t \geq 0, \\ \delta^*(w, y) - \frac{1}{r} (W_{t-} - w), & \hat{W}_t < 0, \end{cases} \quad (4.14)$$

for every $y \in Y$ on $\llbracket 0, \sigma_1 \wedge \tau \rrbracket$, where $N_t = N_{\pi(W_t)}$, $T_t = T_{\pi(W_t)}$, $\phi_t = \phi(a^*, N_t, \delta_t)$, $D_t := N_w^\top W_t - f(\hat{W}_t)$, $\tau := \inf\{t \geq 0 \mid W_t \notin B_\varepsilon(w)\}$ and σ_1 is the first time any of the processes $(J^y)_{y \in Y}$ jump. Note that ϕ is well defined by Lemma 4.5.2.

Let $\rho := \inf\{t \geq 0 \mid D_t > 0\}$. We will show that $\rho \geq \tau \wedge \sigma_1$ a.s. Indeed, on $\llbracket 0, \tau \wedge \sigma_1 \rrbracket$, W and D are continuous, hence D has to reach 0 before crossing it. Since $\delta = \delta^*$ on $\{D = 0\} \cap \{\hat{W} \geq 0\}$, it follows in the same way as in the proof of Lemma 4.7.10 that D remains 0 after reaching it until either $W \notin B_\varepsilon(w)$ or $\hat{W} < 0$. On $\{\hat{W} < 0\}$, the drift is strictly inward pointing and the volatility tangential, hence D is strictly decreasing. It follows that $D_t \leq 0$ on $\llbracket 0, \tau \wedge \sigma_1 \rrbracket$ and hence $\rho \geq \tau \wedge \sigma_1$. In particular W remains in $\mathcal{B}(r, \mathcal{W})$. On $\{\sigma_1 < \tau\}$, it holds by construction that $W_{\sigma_1} \in \mathcal{W}$ and the jumps are parallel to the boundary if W is on the boundary. On the set $\{\tau < \sigma_1\}$, it follows that W_τ is either in $\mathcal{C} \cap B_\varepsilon(w)$ or in the interior of $\mathcal{B}(r, \mathcal{W})$ since D is strictly decreasing on $\{\hat{W} < 0\}$. If $W_\tau \in \text{int } \mathcal{B}(r, \mathcal{W})$, concatenate the solution with a solution to (2.8) attaining W_τ that remains in $\mathcal{B}(r, \mathcal{W})$ until time σ_1 and jumps into \mathcal{W} . If $W_\tau \in \mathcal{C} \cap B_\varepsilon(w)$, Lemma 4.7.10 implies that there exists a solution to (2.8) that remains on \mathcal{C} until either an end point is reached or an event occurs. Since the other end point is in $\mathcal{B}(r, \mathcal{W})$ by assumption, repeating the same procedure yields $w \in \mathcal{B}(r, \mathcal{W})$. \square

Lemma 4.7.21. $\mathcal{B}(r, \mathcal{W})$ is closed.

Proof. By public randomization, a straight line segment is contained in $\mathcal{B}(r, \mathcal{W})$ if both of its end points are contained in $\mathcal{B}(r, \mathcal{W})$. Similarly, Lemma 4.7.10 implies that curved parts of $\partial\mathcal{B}(r, \mathcal{W})$ are contained in $\mathcal{B}(r, \mathcal{W})$ if and only if its end points are in $\mathcal{B}(r, \mathcal{W})$. These end points are either corners of $\mathcal{B}(r, \mathcal{W})$ or points where a curved solution to (4.4) turns into a straight line segment. Lemma 4.7.20 implies that the closure of $\mathcal{B}(r, \mathcal{W})$ is $(0, r, \mathcal{W})$ -admissible, hence it is contained in $\mathcal{B}(r, \mathcal{W})$ by maximality of $\mathcal{B}(r, \mathcal{W})$. \square

Proof of Proposition 4.4.2. Lemmas 4.7.13 and 4.7.14 imply that outside of $\mathcal{G}(r, \mathcal{W})$, the boundary $\partial\mathcal{B}(r, \mathcal{W})$ is a C^1 solution to (4.4). It follows from Corollary 4.7.16 that $\mathcal{G}(r, \mathcal{W})$ has the desired properties and Corollary 4.7.17 shows that all corners are contained in the set of static Nash payoffs. Finally, $\mathcal{B}(r, \mathcal{W})$ is closed by Lemma 4.7.21. \square

4.8 Proof of Lemma 4.5.4

Proof of Lemma 4.5.4. Let w_1, \dots, w_n be an enumeration of all payoffs in $\tilde{\mathcal{P}}(r, \mathcal{W})$ and all corners of $\tilde{\mathcal{B}}(r, \mathcal{W})$. We start by showing that there exists $\varepsilon > 0$ such that for any $k = 1, \dots, n$, there exist a solution to (2.8) starting in w_k satisfying the necessary properties while it is within distance ε of w_k . Indeed, suppose first that w_k is a corner of $\tilde{\mathcal{B}}(r, \mathcal{W})$. By assumption, w_k is decomposable by (a, δ_0) such that $w_k + \eta T \in \tilde{\mathcal{B}}(r, \mathcal{W})$ and $w_k + r\delta(y) + \eta T \in \mathcal{W}$ for every $y \in Y$ and $\eta > 0$ sufficiently small. Consider a solution to (2.8) with $A \equiv a$, $\delta \equiv \delta_0$, $\beta \equiv 0$ and $M \equiv 0$. The continuation value takes the explicit form $W_t = w_k + (e^{rt} - 1)T$ and hence remains on the straight line through w_k in the direction T . By assumption, $W \in \tilde{\mathcal{B}}(r, \mathcal{W})$ and $W + r\delta(y) \in \mathcal{W}$ for every $y \in Y$ on $\Xi := \llbracket 0, \sigma_1 \rrbracket \cap \{W \in B_\varepsilon(w_k)\}$ for ε sufficiently small. Moreover, (β, δ) enforces A and $N^\top \delta(y) \leq 0$ for every $y \in Y$ by decomposability.

If $w_k \in \tilde{\mathcal{P}}(r, \mathcal{W})$ instead, then the solution to (2.8) is constructed similarly as in the proof of Lemma 4.7.20, with the exception that the definition of $\delta_t(y)$ in (4.14) is replaced by

$$\delta_t(y) = \begin{cases} \delta^*(\pi(W_{t-}), y), & \hat{W}_t \geq 0 \\ \delta^*(\pi(w), y), & \hat{W}_t < 0 \end{cases},$$

where we adopt the same notation as in the proof of Lemma 4.7.20. The condition that $w + r\delta^*(w; y) \in \text{int } \mathcal{W}$ implies that $W + r\delta(y) \in \mathcal{W}$ on Ξ for $\varepsilon > 0$ small enough. The other conditions are satisfied by construction.

Fix $k = 1, \dots, n$ and a solution to (2.8) starting in w_k . Define the stopping time $\tau := \inf\{t \geq 0 \mid W_t \notin \mathcal{B}_\varepsilon(w_k)\}$ and observe that $W_\tau \in \partial B_\varepsilon(w_k) \cap \tilde{\mathcal{B}}(r, \mathcal{W})$ by continuity of W on Ξ . Because there are finitely many corners and payoffs in $\tilde{\mathcal{P}}(r, \mathcal{W})$ by assumption, any payoff in $\bigcup_{k=1}^n \partial B_\varepsilon(w_k) \cap \tilde{\mathcal{B}}(r, \mathcal{W})$ can be attained by a public randomization device taking finitely many values v_1, \dots, v_K in $\partial \tilde{\mathcal{B}}(r, \mathcal{W})$, such that any v_ℓ is either an element of $\{w_1, \dots, w_n\}$ or $\partial \tilde{\mathcal{B}}(r, \mathcal{W})$ is locally a solution to (2.8) around v_ℓ . Therefore, any v_ℓ can be attained by a solution W^ℓ to (2.8) that is continuous on $\llbracket 0, \sigma_1 \rrbracket$ until time

$$\tau_\ell := \inf \left\{ t \geq 0 \mid W^\ell \in \bigcup_{k=1}^n \partial B_\varepsilon(w_k) \cap \text{int } \tilde{\mathcal{B}}(r, \mathcal{W}) \right\},$$

at which point it can be attained by a public randomization device again with values in v_1, \dots, v_m . Note that an iteration of this procedure will extend to σ_1 with certainty because $\tilde{\tau} := \min_{\ell=1, \dots, m} \tau_\ell > 0$ a.s. and the countable sum of identical and independent copies of $\tilde{\tau}$ extends to ∞ by Lemma 3.6.3. For any w_k , we have thus constructed a solution to (2.8) with the necessary properties on $\llbracket 0, \sigma_1 \rrbracket$. By public randomization and Corollary 4.7.11, this holds true for any payoff in $\tilde{\mathcal{B}}(r, \mathcal{W})$ and hence $\tilde{\mathcal{B}}(r, \mathcal{W})$ is admissible. \square

4.A Appendix: Auxiliary results

4.A.1 Lipschitz continuity of set-valued maps

While the arbitrary union of Lipschitz continuous maps $(F_i)_{i \in I}$ is Lipschitz continuous if the Lipschitz constants K_i are uniformly bounded, the intersection of two Lipschitz continuous maps may fail to be Lipschitz continuous in general. In this appendix, we show that for two special cases that are relevant in our setting, the intersection is indeed Lipschitz continuous.

Lemma 4.A.1. *The intersection of two convex-valued affine maps is Lipschitz continuous.*

Proof. Let F and G be two convex-valued affine maps. If $F \cap G$ is continuous, then it is Lipschitz continuous by affinity of F and G . Suppose towards a contradiction that $F \cap G$ fails to be continuous at x_0 , that is, there exists $v \in F(x_0) \cap G(x_0)$ such that $B_\varepsilon(v) \cap F(x) \cap G(x) = \emptyset$ for $\varepsilon > 0$ arbitrarily small and $x \in \text{supp } F \cap G$ arbitrarily close to x_0 . Since F and G are affine, this is possible only if $N_F = -N_G$, where N_F and N_G denote the normal vectors to $\partial F(x_0)$ and $\partial G(x_0)$, respectively, at v . Convexity implies that $F(x) \cap G(x) = \emptyset$ for x arbitrarily close to x_0 , contradicting the fact that $x \in \text{supp } F \cap G$. \square

Lemma 4.A.2. *Let F and G be Lipschitz continuous with bounded support taking values in closed convex polytopes. Let $\pi_i^F(x)$, $i \in I_F$ and $\pi_i^G(x)$, $i \in I_G$ denote the outward normal vectors to their hyperfaces, respectively. If for any $J_F \subseteq I_F$, $J_G \subseteq I_G$, the matrix $[(\pi_j^F(x))_{j \in J_F}, (\pi_j^G(x))_{j \in J_G}]$ has constant column rank in a neighbourhood of x , then $F \cap G$ is locally Lipschitz continuous at x . If the ranks are constant on $\text{supp } F \cap G$, then $F \cap G$ is Lipschitz continuous.*




Figure 4.14: Level sets $H(z)$ of $\partial(F(x) \cap G(x))$ containing points p_z with maximal distance from $p_0 \in \partial(F(x) \cap G(x))$. Clearly, $\|p_z - p_0\| = z\|p_1 - p_0\|$.

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

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

However, the right-hand side is larger than $F(x) \cap G(x) + \|\tilde{x} - x\|B_K(0)$. Let $H(z) := \partial(F(x) + B_z(0)) \cap \partial(G(x) + B_z(0))$ be the level sets of $\partial(F(x) \cap G(x))$. Let p_1 denote a point in $H(1)$ with maximal distance from $\partial(F(x) \cap G(x))$ and let p_0 be a point in $\partial(F(x) \cap G(x))$ with minimal distance from p_1 as illustrated in Figure 4.14. Let $\{\pi_1, \dots, \pi_n\}$ be a minimal subset of normal vectors to the hyperfaces of $F(x) \cap G(x)$ that intersect at p_0 such that p_1 is the unique point in $H(1)$, which is related to p_0 by

$$\pi_j^\top(p_1 - p_0) = 1 \quad \text{for } j = 1, \dots, n \quad \text{and} \quad p_1 - p_0 \in \underset{j=1, \dots, n}{\text{span}} \pi_j. \quad (4.15)$$

By linearity of (4.15), it follows that $p_{K\|\tilde{x}-x\|} := p_0 + K\|\tilde{x} - x\|(p_1 - p_0)$ is a point in $H(K\|\tilde{x} - x\|)$ with maximal distance of $F(x) \cap G(x)$. Its distance from p_0 equals $K\|p_1 - p_0\|\|\tilde{x} - x\|$. The statement thus follows once we show that $\|p_1(x) - p_0(x)\|$ is uniformly bounded in x .

By minimality of $\{\pi_1, \dots, \pi_n\}$, the vectors π_1, \dots, π_n are linearly independent. Thus by assumption, $\pi_1(\tilde{x}), \dots, \pi_n(\tilde{x})$ are linearly independent also for \tilde{x} in a neighbourhood of x . Since F and G are continuous, the norm of the solution is continuous, hence by making the neighbourhood smaller and compact, its maximum is bounded. Because F and G have finitely many hyperfaces, the finite maximum over all possible combinations of normal vectors $\pi_1, \dots, \pi_{\bar{n}}$ yields a bound for $\|p_1 - p_0\|$ on a sufficiently small neighbourhood of x . Finally, if the rank is constant on $\text{supp } F \cap G$, then $\|p_1 - p_0\|$ is uniformly bounded since $\text{supp } F \cap G$ is compact because F and G are closed-valued. \square

4.A.2 Bounds on incentives and solutions to (4.9)

Lemma 4.A.3. *Let $w \mapsto \mathcal{D}(w)$ be of class B and let \mathcal{C} be a C^1 solution to (4.9) with $(0, \mathcal{D})$ oriented by $w \mapsto N_w$ with endpoints v_L, v_R such that $\mathcal{N}_{\mathcal{C}} \subseteq E_a^0(r, \mathcal{D})$ and $\mathcal{N}_{\mathcal{C}} \cap (\Gamma_a^0(r, \mathcal{D}) \cup \mathcal{P}) = \emptyset$ for some $a \in \mathcal{A}$. Then there exists $K > 0$ such that for any $w \in \mathcal{C}$, $(T_w \phi + N_w \chi, \delta)$ enforcing a with $N_w^\top (g(a) + \delta \lambda(a) - w) \geq 0$ and $\delta(y) \in \mathcal{D}(w)$ and $N_w^\top \delta(y) \leq 0$ for all $y \in Y$ implies that for any $\alpha \geq 0$,*

$$K \|\chi\| \geq \|\hat{\delta}^1 - \delta^1\| + \|\hat{\delta}^2 - \delta^2\|, \quad \frac{2K + 2\alpha}{\bar{\Psi}} \|\chi\| \geq 1 - \frac{(\|\phi\| - \alpha \|\chi\|)^2}{\|\phi(a, N, \hat{\delta})\|^2}, \quad (4.16)$$

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

Proof. Fix a, w, δ . We start by extending $\phi(a, N, \delta)$ to δ with $\Phi_a(N, \delta) = \emptyset$ in a Lipschitz continuous way. For $i = 1, 2$, let $\mathcal{I}_a^i(N, \delta^i)$ be defined as in Lemma 4.7.3. Because $(T\phi + N\chi, \delta)$ enforces a by assumption, it follows that

$\mathcal{I}_a^i(N, \delta^i) \neq \emptyset$. Any pair (ϕ', χ') with $(T\phi' + N\chi', \delta)$ enforcing a has to satisfy $\phi' + N^i/T^i\chi' \in \mathcal{I}_a^i(N, \delta)$ for $i = 1, 2$, and hence

$$\left(\frac{N^1}{T^1} + \frac{N^2}{T^2}\right) \|\chi'\| \geq d(\mathcal{I}_a^1(N, \delta^1), \mathcal{I}_a^2(N, \delta^2)),$$

where $d(\mathcal{I}_a^1(N, \delta^1), \mathcal{I}_a^2(N, \delta^2))$ denotes the minimal distance between the two sets. Because $\mathcal{I}_a^i(N, \delta)$ are closed, the minimal distance is attained between two points $p_i \in \mathcal{I}_a^i(N, \delta^i)$. If the minimal distance is attained for more than one such pair (p_1, p_2) , let (p_1, p_2) be the pair that minimizes the norm of

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

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

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

That is, $(\phi(a, N, \delta), \chi(a, N, \delta))$ is the pair (ϕ', χ') with $(T\phi' + N\chi', \delta)$ enforcing a that minimizes first $\|\chi'\|$ and then $\|\phi'\|$; see also Figure 4.15. Observe that this definition is indeed an extension since $\Phi_a(N, \delta) \neq \emptyset$ implies $\chi(a, N, \delta) = 0$. For $\kappa \in [0, 1]$, let $\delta_\kappa := \kappa\hat{\delta} + (1-\kappa)\delta$ and define $d(\kappa) := d(\mathcal{I}_a^1(N, \delta_\kappa^1), \mathcal{I}_a^2(N, \delta_\kappa^2))$. We will show that $d(\kappa)$ is piecewise linear and strictly decreasing in κ . Indeed, because a change in δ_κ shifts the hyperfaces of $\mathcal{I}_a^i(N, \delta_\kappa^i)$ in a linear way, it is clear that d is piecewise linear. Suppose towards a contradiction that $d(\kappa)$ is increasing on an interval. This means that one of the hyperfaces




Figure 4.15: The left panel illustrates the position of $\phi(a, N, \delta)$, p_1 and p_2 relative to $\mathcal{I}_a^1(N, \delta^1)$ and $\mathcal{I}_a^2(N, \delta^2)$. It also shows that $d_i \leq N^i/T^i\|\chi\|$ has to hold. The right panel shows that $\theta_1 + \theta_2 = \gamma$ and hence $\theta_j \geq \gamma/2$.

of $\mathcal{I}_a^i(N, \delta_\kappa^i)$ is moving away from $\mathcal{I}_a^{-i}(N, \delta_\kappa^{-i})$ at a faster rate than all the hyperfaces of $\mathcal{I}_a^{-i}(N, \delta_\kappa^{-i})$ are catching up. Because δ_κ is linear, this implies that the hyperfaces of $\mathcal{I}_a^{-i}(N, \delta_\kappa^{-i})$ can never catch up and hence $d(1) > 0$, a contradiction. Finally, $d(\kappa)$ is strictly decreasing by minimality of $\hat{\delta}$. Let $K_1(\delta) = \min|dd(\kappa)/d\kappa| > 0$ and observe that $K_1(\delta)$ is continuous in δ . Indeed, $\Psi_a^0(w, N_w, r, \mathcal{D})$ is convex by Lemma 4.7.5, hence $\hat{\delta}$ is continuous in δ as the minimizer of a convex function over a convex set. Therefore, $d(\kappa)$ is continuous in δ as well and so is $K_1(\delta)$. Let \mathcal{J}_a denote the set of all δ , where there exists β such that (β, δ) enforces a . Observe that $\mathcal{D}(\mathcal{C})$ is compact as the image of a compact set under an affine function and that $\bigcup_{w \in \mathcal{C}} H(N_w) = H(N_{v_L}) \cap H(N_{v_R})$ due to positive curvature of \mathcal{C} . Therefore, $\mathcal{J}_a \cap \bigcup_{w \in \mathcal{C}} (\mathcal{D}(w) \cap H(N_w)) = \mathcal{J}_a \cap \mathcal{D}(\mathcal{C}) \cap H(N_{v_L}) \cap H(N_{v_R})$ is compact and the minimum K_1 of $K_1(\delta)$ over all eligible δ is attained, hence positive. Thus,

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

Since $\|\chi(a, N, \delta)\| \leq \|\chi\|$ by definition, it follows that there exists a constant K_2 such that $\|\hat{\delta}^1 - \delta^1\| + \|\hat{\delta}^2 - \delta^2\| \leq K_2\|\chi\|$. The second inequality follows once we show that

$$\|\phi(a, N, \hat{\delta})\| - \|\phi\| \leq K_3\|\hat{\delta} - \delta\| + K_3\|\chi\| \quad (4.17)$$

for some constant K_3 , where we denote $\|\hat{\delta} - \delta\| := \|\hat{\delta}^1 - \delta^1\| + \|\hat{\delta}^2 - \delta^2\|$ for the sake of brevity. Indeed, the right-hand side of (4.17) is bounded by $K_3(K_2 + 1)\|\chi\|$ due to the already established inequality. Thus,

$$\begin{aligned} 1 - \frac{\|\phi\| - \alpha\|\chi\|}{\|\phi(a, N, \hat{\delta})\|} &\leq \frac{\|\phi(a, N, \hat{\delta})\| - \|\phi\| + \alpha\|\chi\|}{\bar{\Psi}} \\ &\leq \frac{K_3(K_2 + 1) + \alpha}{\bar{\Psi}}\|\chi\|. \end{aligned} \quad (4.18)$$

Observe that $\bar{\Psi}$ is positive because \mathcal{N}_C is bounded away from $\Gamma(r, \mathcal{D}) \cup \mathcal{P}$ by closedness of both sets. The second inequality in (4.16) then follows from (4.18) in conjunction with the elementary inequality $1 - x \geq \frac{1}{2}(1 - x^2)$. It remains to show (4.17).

Suppose first that $\|\phi\| \geq \|\phi(a, N, \delta)\|$. Then Lipschitz continuity implies that $\|\phi\| \geq \|\phi(a, N, \hat{\delta})\| - K\|\hat{\delta} - \delta\|$, which readily implies (4.17). Suppose therefore that $\|\phi\| < \|\phi(a, N, \delta)\|$. Let d_i denote the distance of ϕ from $\mathcal{I}_a^i(N, \delta)$ for $i = 1, 2$ and observe that $d_i \leq N^i/T^i\|\chi\|$ as illustrated in the left panel of Figure 4.15. Define the auxiliary sets

$$\tilde{\mathcal{I}}_a^i(N, \delta^i) := \mathcal{I}_a^i(N, \delta^i) - \frac{N^i}{T^i}\chi(a, N, \delta)$$




Figure 4.16: Because q_2 is the projection of ϕ onto $\partial\tilde{\mathcal{I}}_a^2(N, \delta^2)$, the angle at q_2 in the triangle shown to the left is at least 90° . Therefore, $\tilde{d}_2 \geq b_2 = \ell_2 \tan(\alpha_2)$ and thus the triangle inequality implies $\|\phi(a, N, \delta) - \phi\| \leq \tilde{d}_2 + \ell_2 \leq \tilde{d}_2(1 + \frac{1}{\tan(\alpha_2)})$. The right panel illustrates that $\alpha_j \geq \theta_j \geq \gamma/2$.

so that $\phi(a, N, \delta) \in \tilde{\mathcal{I}}_a^1(N, \delta^1) \cap \tilde{\mathcal{I}}_a^2(N, \delta^2)$ as shown in the right panel of Figure 4.15. Let \tilde{d}_i denote the distance of ϕ from $\tilde{\mathcal{I}}_a^i(N, \delta)$ and observe that $\tilde{d}_i \leq d_i$. Let q_i for $i = 1, 2$ denote the point in $\partial\tilde{\mathcal{I}}_a^i(N, \delta)$ closest to ϕ and let ϕ' be the projection of ϕ onto the plane through $\phi(a, N, \delta)$, q_1 and q_2 . Let $j \in \{1, 2\}$ be the index i for which the angle θ_i between $\phi(a, N, \delta) - \phi'$ and $\phi(a, N, \delta) - q_i$ is maximal. Then $\theta_j \geq \gamma/2$, where γ is the angle between $\phi(a, N, \delta) - q_1$ and $\phi(a, N, \delta) - q_2$. Let α_i be the angles between $\phi(a, N, \delta) - \phi$ and $\phi(a, N, \delta) - q_i$ and observe that $\alpha_i \geq \theta_i$. Then

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

as illustrated in Figure 4.16. Observe that it is impossible for γ to be 0 by the definition of $\phi(a, N, \delta)$. Since changes in N and δ do not change the direction of the hyperplanes bounding $\tilde{\mathcal{I}}_a^i(N, \delta)$, a uniform lower bound $\underline{\gamma}$ for γ is given by taking the minimum over all strictly positive angles between the finitely many hyperfaces of $\partial\tilde{\mathcal{I}}_a^1(N, \delta)$ and $\partial\tilde{\mathcal{I}}_a^2(N, \delta)$. Thus, $\|\phi(a, N, \delta) - \phi\| \leq K_4 \|\chi\|$ for

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

Finally, (4.17) follows from the triangle inequality

$$\begin{aligned} \|\phi(a, N, \hat{\delta}) - \phi\| &\leq \|\phi(a, N, \hat{\delta}) - \phi(a, N, \delta)\| + \|\phi(a, N, \delta) - \phi\| \\ &\leq K_5 \|\hat{\delta} - \delta\| + K_4 \|\chi\|, \end{aligned}$$

where K_5 is the Lipschitz constant of ϕ . □

Lemma 4.A.4. *Let $w \mapsto \mathcal{D}$ be of class B and let \mathcal{C} be a C^1 solution to (4.6) oriented by $w \mapsto N_w$ such that $\mathcal{N}_{\mathcal{C}} \cap (\Gamma_a^\varepsilon(r, \mathcal{D}) \cup E_a^\varepsilon(r, \mathcal{D}) \cup \mathcal{P}) = \emptyset$ for some $a \in \mathcal{A}$. Then there exists $K > 0$ such that for any $w \in \mathcal{C}$, $(T_w\phi + N_w\chi, \delta)$ enforcing a implies $K \leq \|\chi\|$.*

Proof. Let \mathcal{J}_a denote the set of all δ , for which there exists β such that (β, δ) enforces a . Fix $w \in \mathcal{C}$ and $\delta \in \mathcal{J}_a \cap \mathcal{D}(w)$ and suppose that $(T_w\phi + N_w\chi, \delta)$ enforces a . From $\mathcal{N}_{\mathcal{C}} \cap \Gamma(r, \mathcal{D}) = \emptyset$ it follows that $T_w\phi + N_w\chi \neq 0$. Moreover, $\Psi_a(w, N_w, r, \mathcal{D}) = \emptyset$ since $\mathcal{N}_{\mathcal{W}} \cap E_a(r, \mathcal{D}) = \emptyset$. This implies that $\mathcal{I}_a^1(N_w, \delta) \cap \mathcal{I}_a^2(N_w, \delta) = \emptyset$, hence $\mathcal{I}_a^1(N_w, \delta)$ and $\mathcal{I}_a^2(N_w, \delta)$ are strictly separated by closedness. Let $d(N_w, \delta) > 0$ denote the distance of the two sets. Because N_w is bounded away from coordinate directions, the map $(N_w, \delta) \mapsto \mathcal{I}_a^i(N_w, \delta)$ is continuous for $i = 1, 2$. Because $w \mapsto \mathcal{D}(w)$ is affine, $\mathcal{D}(\mathcal{C})$ is compact, hence so is $\mathcal{D}(\mathcal{C}) \cap \mathcal{J}_a$. Therefore, the minimum of $d(N_w, \delta)$ over the compact set $\mathcal{C} \times \mathcal{D}(\mathcal{C}) \cap \mathcal{J}_a$ is attained, in particular positive. □

The last result of this thesis is a Gronwall-type lemma, giving a measure of closeness of solutions to the optimality equation (4.9) when the target sets for the jumps are close to each other.

Lemma 4.A.5. *Let $\varepsilon_1 \leq \varepsilon_2$ and let $\mathcal{D}_1(w)$ and $\mathcal{D}_2(w)$ be maps of class B such that there exists $\varepsilon > 0$ such that for every $w \in \mathcal{V}$ and every $N \in S^1$,*

$$\mathcal{D}_2(w) \cap H_{\varepsilon_2}(N) \subseteq \mathcal{D}_1(w) \cap H_{\varepsilon_1}(N) \subseteq \mathcal{D}_2(w) \cap H_{\varepsilon_2}(N) + B_\varepsilon(0).$$

Let (w, N) such that (4.9) with $(\varepsilon_1, \mathcal{D}_1(w))$ and $(\varepsilon_2, \mathcal{D}_2(w))$ is Lipschitz continuous in a neighbourhood U of (w, N) . Let \mathcal{C}_1 and \mathcal{C}_2 be two solutions to (4.9) with $(\varepsilon_1, \mathcal{D}_1(w))$ and $(\varepsilon_2, \mathcal{D}_2(w))$, respectively, with initial value (w, N) such that $\mathcal{N}_{\mathcal{C}_1}, \mathcal{N}_{\mathcal{C}_2} \subseteq U$. Then there exist constants K_1, K_2, K_3 such that for any $v \in \mathcal{C}_1$, there exists $v' \in \mathcal{C}_2$ with $\|v - v'\| \leq K_1\varepsilon(\|v - w\|^2 + K_2e^{K_3\|v-w\|})$.

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

$$f''(x) = F(x, f(x), f'(x)), \quad h''(x) = H(x, h(x), h'(x)) \quad (4.19)$$

with $f(0) = h(0) = 0$ and $f'(0) = h'(0) = 0$ for Lipschitz continuous F and H with Lipschitz constants K_F and K_H , respectively. By Lemma 4.7.3, the right-hand side of (4.9) is Lipschitz in δ for $(v, N_v) \in U$ with Lipschitz constant K . Since $\mathcal{D}_2(w) \cap H_{\varepsilon_2}(N) \subseteq \mathcal{D}_1(w) \cap H_{\varepsilon_1}(N) \subseteq \mathcal{D}_2(w) \cap H_{\varepsilon_2}(N) + B_\varepsilon(0)$, it follows that $0 \leq F(x, d, v) - H(x, d, v) \leq K\sqrt{m}\varepsilon$, hence integrating (4.19) yields

$$\begin{aligned} f'(x) - h'(x) &= \int_0^x \left(F(t, f(t), f'(t)) - H(t, f(t), f'(t)) \right. \\ &\quad \left. + H(t, f(t), f'(t)) - H(t, h(t), h'(t)) \right) dt \\ &\leq K\sqrt{m}\varepsilon x + K_H \int_0^x (|f(t) - h(t)| + |f'(t) - h'(t)|) dt. \end{aligned}$$

Since $H(x, d, v) \leq F(x, d, v)$ in a neighbourhood of 0, we may assume that $f'(x) > h'(x)$ and $f(x) > h(x)$ by choosing U small enough. Therefore, $f - h$ satisfies the conditions of Theorem 1.8.1 in Pachpatte [35], which implies that

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

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

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

For any $v = (x, h(x))$, let $v' = (x, f(x))$, hence the result follows from the inequality $x \leq \|v - v'\|$. □

Chapter 5

Conclusion

This thesis treats continuous-time repeated games with imperfect public information in a fairly general setting that includes both continuous and abrupt information through the arrival of infrequent events. At the heart of this thesis is a rigorous mathematical foundation on which we build the theory of these games: we develop many continuous-time analogues to well-known concepts of discrete time and show that the continuous-time techniques often lead to clean, quantitative results. Building on this framework, we establish several versions of folk theorems in Chapter 3, finding sufficient conditions for players to attain asymptotic efficiency as they get increasingly patient. This is possible in a surprisingly simple class of strategies that are well suited for implementation. In Chapter 4, we find an exact description of the equilibrium payoff set $\mathcal{E}(r)$ in two-player games, and provide an algorithm similar to that known from discrete time, with which $\mathcal{E}(r)$ can be computed. In contrast to its discrete-time counterpart, however, the payoff set at each step of the algorithm is explicitly characterized through an ordinary differential equation.

In Chapter 2, we provide the rigorous foundation necessary for the main results of the other chapters. We establish and discuss many continuous-time analogues to concepts that are well-known from discrete time: We provide a tractable framework for modelling mixed and behaviour strategies in continuous time and show that the two notions are realization equivalent. This framework is also suitable to model private sources of information, which may be useful for future research on continuous-time games with private information. We use it to show that best responses to public strategy profiles always exist in public strategies and we discuss uniqueness of outcomes. We extend the incentive compatibility conditions of Sannikov [37] and Sannikov and Skrzypacz [39] to multiplayer games in mixed strategies and provide the continuous-time analogue to the one-shot deviation principle and the important notion of self-generating payoff sets.

In addition to establishing the continuous-time analogues to the above concepts, the continuous-time setting gives rise to explicit relations that do not exist in discrete time between the informativeness of the public signal, players' patience, exchangeability of game primitives and the use of public randomization. Low discount rates are associated with patient players because they value future payoffs more. In continuous time, this is related very visually to the execution speed of strategies: an equilibrium profile can be transformed to a lower discount rate (more patient players) by reducing the execution speed of the strategy profile and using public randomization suitably. The use of public randomization in this case is related to the increased informativeness of the public signal: When players become more patient and strategy profiles are executed at a slower speed, a longer period of observation becomes available

at the same cost. Therefore, players can better estimate the underlying drift rate, which increases the relative informativeness of the public signal. To balance this in equilibrium, players create noise artificially by using public randomization. This procedure is, in fact, a special case of the more general concept that continuous-time game primitives $(r, \mu, \sigma, \lambda)$ can be exchanged for each other and that public randomization can be used to continuously move across game primitives $(r, \mu, \sigma, \lambda)$ with an increased informativeness of the public signal relative to players' patience.

In Chapter 3, we develop the techniques of Fudenberg, Levine and Maskin [16] for a continuous-time setting and establish several versions of folk theorems in pure and behaviour strategies. This turns out to be more challenging in a continuous-time setting than in discrete time: Because incentives have to be provided continuously, rather than only at discrete time points, abrupt information cannot be used on payoff sets arbitrarily close to \mathcal{V}^* . Moreover, the notion of decomposability on tangent hyperplanes has to be strengthened to uniform decomposability so that action profiles can be enforced on tangent hyperplanes arbitrarily close to being coordinate.

An interesting contribution is the observation that payoffs on a uniformly decomposable payoff set are attainable by concatenations of locally strong solutions to the stochastic differential equation characterizing the continuation value. A strong solution differs from a weak solution by the fact that the source of uncertainty is fixed, and locally, constant strategy profiles are enforced. Therefore, the resulting equilibrium profiles switch action profiles only finitely many times on finite time intervals, which is very desirable for implementation of these strategies. This shows that despite the possibility to react

infinitesimally fast, players do not need to do so to attain asymptotic efficiency. This is a partial rebuttal to the concern that continuous-time models lead to strategies that cannot be implemented because of unbounded oscillation between action profiles as it is the case in Sannikov [37], for example.

This construction of locally constant strategy profiles could be significant for future research on the connection between discrete- and continuous-time games. Indeed, on a uniformly decomposable payoff set \mathcal{W} , equilibrium profiles need only be adapted at independent copies of a certain stopping time τ . This suggests that payoffs within \mathcal{W} can also be attained in equilibrium of a discrete-time game, where the lengths of the time intervals between periods are random (identical copies of τ), rather than fixed. By considering smooth inner approximations $(\mathcal{W}_n)_{n \geq 0}$ of $\mathcal{E}(r)$, one could construct a sequence of discrete-time repeated games such that the set of equilibria in the discrete-time game converges to the set of continuous-time equilibria.

In Chapter 4, we vastly generalize the class of games, for which an exact description of $\mathcal{E}(r)$ is known via a differential equation describing its boundary. Not only do we extend the type of information from Sannikov [37] to an information structure that includes both continuous and abrupt information, but we also lower the necessary assumptions on game primitives so that the characterization is new even when the public signal is continuous, but one-dimensional. This includes the important games of a Cournot duopoly in a single homogeneous good and partnership games, where only the total revenue of the partnership is observed.

Our result shows the drastically different impacts that continuous and abrupt information have on equilibrium payoffs. This is due to the funda-

mentally different ways in which the two types of information are used to provide incentives on $\partial\mathcal{E}(r)$: The continuous information can be used only to transfer value tangentially amongst players, but the amount of value that can be transferred is unbounded. In contrast, the abrupt information can be used to transfer value as well as to destroy value, but only bounded amounts of value can be transferred or destroyed in this way. We also show that this amount is decreasing in the discount rate and hence, fewer incentives can be provided through the observation of infrequent events for impatient players.

Despite the burning of value, the addition of abrupt information may increase efficiency in equilibrium by quite a large amount because of a tradeoff with the tangential incentives from the continuous component of the public signal. This additional efficiency can be very significant in some games, even when the observed events are of purely informational nature and do not themselves affect players' payoffs as we show with examples in Sections 4.3.1 and 4.6.2. Because of the quantitative nature of our result, it is possible quantify the value of observing infrequent events given the continuous information, which may lead to interesting future research in mechanism design. A mechanism designer may, for example, compare the value of observing certain events to the cost of reporting them. If the gained efficiency outweighs the cost of reporting, requiring participants to report these events will increase overall efficiency. It may also lead to future research questions in information disclosure and industrial organization, where firms choose to either release their information continuously or abruptly to maximize their objectives.

In contrast to Sannikov [37], the characterization of $\mathcal{E}(r)$ is not explicit but rather a fixed-point characterization because the local description of $\partial\mathcal{E}(r)$ is

an ordinary differential equation that depends on its entire solution $\mathcal{E}(r)$. By restricting the jumps of the continuation value to land in some fixed payoff set \mathcal{W} , we obtain the largest payoff set $\mathcal{B}(r, \mathcal{W})$ that is self-generating up to the arrival of the first rare event. This is a continuous-time analogue to the standard set operator \mathcal{B} in Abreu, Pearce and Stacchetti [2], again suggesting that continuous-time games may better be approximated by discrete-time games with random period lengths rather than fixed ones. Similarly to [2], a successive application of \mathcal{B} to \mathcal{V}^* converges to $\mathcal{E}(r)$. Contrary to its discrete counterpart, however, $\mathcal{B}(r, \mathcal{W})$ can be explicitly characterized through an ordinary differential equation, leading to an efficient computation of $\mathcal{E}(r)$.

Our proof also reveals interesting implications of abrupt information on equilibrium strategies and not just their associated payoffs. In the setting with abrupt information, $\mathcal{E}(r)$ may have corners outside of \mathcal{V}^N as illustrated in Section 4.6.2. When the continuation value reaches such a corner, it is absorbed there until the arrival of the next rare event. Such an event will occur eventually because we consider games of full support public monitoring, after which play becomes dynamic again. This shows that in certain situations, the arrival of infrequent events may have such a high threat or reward level that the observation of the continuous information becomes irrelevant to provide incentives. Payoffs that correspond to these situations can thus be thought of as locally static Nash payoffs.

In conclusion, this thesis makes several important contributions to the theory of repeated games. It introduces the arrival of informative but infrequent events to continuous-time games with imperfect public monitoring. Capitalizing on the mathematical tractability of continuous-time techniques, we prove

several folk theorems and explicit results. The field of continuous-time repeated games is quite young and it seems promising that more results can be obtained using continuous-time techniques that help us deepen our understanding of repeated interactions. In addition to its contributions to the theory of repeated games, this thesis develops techniques that may be useful in the connection between discrete- and continuous-time games, mechanism design and information disclosure.

Bibliography

- [1] Dilip Abreu, David Pearce and Ennio Stacchetti: Optimal cartel equilibria with imperfect monitoring, *Journal of Economic Theory*, **39** (1986), 251–269
- [2] Dilip Abreu, David Pearce and Ennio Stacchetti: Toward a theory of discounted repeated games with imperfect monitoring, *Econometrica*, **58** (1990), 1041–1063
- [3] Dilip Abreu and Yuliy Sannikov: An algorithm for two-player repeated games with perfect monitoring, *Theoretical Economics*, **9** (2014), 313–338
- [4] Robert M. Anderson: Quick-response equilibrium. Working Papers in Economic Theory and Econometrics # IP-323, Center for Research in Management, University of California, Berkeley (1984)
- [5] Jean-Pierre Aubin and Hélène Frankowska: *Set-Valued Analysis*, Springer, Berlin, 1990
- [6] Benjamin Bernard and Christoph Frei: The folk theorem with imperfect public information in continuous time, *Theoretical Economics*, forthcoming
- [7] Patrick Billingsley: *Probability and Measure*, 2nd ed., Wiley, 1986
- [8] Patrick Bolton and Christopher Harris: Strategic experimentation, *Econometrica*, **67** (1999), 349–374
- [9] Mark Cronshaw and David G. Luenberger: Strongly symmetric subgame perfect equilibrium in infinitely repeated games with perfect monitoring and discounting, *Games and Economic Behaviour*, **6** (1994), 220–237
- [10] Brendan Daley and Brett Green: Waiting for news in the market for lemons, *Econometrica*, **80** (2012), 1433–1504
- [11] Claude Dellacherie and Paul-André Meyer: *Probabilities and Potential, B: Theory of Martingales*, North-Holland Pub. Co, 1982
- [12] Eduardo Faingold and Yuliy Sannikov: Reputation in continuous-time games, *Econometrica*, **79** (2011), 773–876

- [13] Drew Fudenberg, David M. Kreps and Eric Maskin: Repeated games with long-run and short-run players, *Review of Economic Studies*, **57** (1990), 555–573
- [14] Drew Fudenberg and David K. Levine: Efficiency and observability with long-run and short-run players, *Journal of Economic Theory*, **62** (1994), 103–135
- [15] Drew Fudenberg and David K. Levine: Continuous time limits of repeated games with imperfect public monitoring, *Review of Economic Dynamics*, **10** (2007), 173–192
- [16] Drew Fudenberg, David K. Levine and Eric Maskin: The folk theorem with imperfect public information, *Econometrica*, **62** (1994), 997–1039
- [17] Drew Fudenberg and Eric Maskin: The folk theorem in repeated games with discounting or with incomplete information, *Econometrica*, **54** (1986), 553–554
- [18] Drew Fudenberg and Jean Tirole: *Game Theory*, MIT Press, Cambridge, 1991
- [19] Edward J. Green and Robert H. Porter: Noncooperative collusion under imperfect price information, *Econometrica*, **52** (1994), 87–100
- [20] John C. Harsanyi: Games with incomplete information played by ‘Bayesian’ players, *Management Science*, **14** (1967–68), 159–182, 320–334, 486–502
- [21] Tadashi Hashimoto: Corrigendum to ‘Games with imperfectly observable actions in continuous time’, *Econometrica*, **78** (2010), 1155–1159
- [22] Kenneth L. Judd: The law of large numbers with a continuum of IID random variables, *Journal of Economic Theory*, **35** (1985), 19–25
- [23] Kenneth L. Judd, Sevin Yeltekin and James Conklin: Computing supergame equilibria, *Econometrica*, **71** (2003), 1239–1254
- [24] Olav Kallenberg: *Foundations of Modern Probability*, Springer, New York, 1997
- [25] Michihiro Kandori: The use of information in repeated games with imperfect monitoring, *Review of Economic Studies*, **59** (1992), 581–593
- [26] Michihiro Kandori and Ichiro Obara: Efficiency in repeated games revisited: the role of private strategies, *Econometrica*, **74** (2006), 499–519
- [27] Ioannis Karatzas and Steven Shreve: *Methods of Mathematical Finance*, Applications of Mathematics 39, Springer, New York, 1998
- [28] Norihiko Kazamaki: *Continuous Exponential Martingales and BMO*, Lecture Notes in Mathematics 1579, Springer, Berlin, 1994

- [29] Godfrey Keller and Sven Rady: Strategic experimentation with Poisson bandits, *Theoretical Economics*, **5** (2010), 275–311
- [30] Harold W. Kuhn: Extensive games and the problem of information, *Contributions to the Theory of Games*, **2** (1953), 193–216
- [31] George J. Mailath and Larry Samuelson: *Repeated Games and Reputations*, Oxford University Press, 2006
- [32] James R. Munkres: *Topology*, 2nd ed., Prentice Hall, 2000
- [33] John F. Nash: Equilibrium points in n-person games, *Proceedings of the National Academy of Sciences*, **36** (1950), 48–49
- [34] Bernt Øksendal: *Stochastic Differential Equations*, 5th ed., Springer, New York, 1998
- [35] Baburao G. Pachpatte: *Inequalities for Differential and Integral Equations*, San Diego: Academic Press, 1998
- [36] Philip Protter: *Stochastic Integration and Differential Equations*, 2nd ed., Stochastic Modelling and Applied Probability 21, Springer, New York, 2005
- [37] Yuliy Sannikov: Games with imperfectly observable actions in continuous time, *Econometrica*, **75** (2007), 1285–1329
- [38] Yuliy Sannikov and Andrzej Skrzypacz: Impossibility of collusion under imperfect monitoring with flexible production, *American Economic Review*, **97** (2007), 1794–1823
- [39] Yuliy Sannikov and Andrzej Skrzypacz: The role of information in repeated games with frequent actions, *Econometrica*, **78** (2010), 847–882
- [40] Reinhard Selten: Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragerträglichkeit, *Zeitschrift für die Gesamte Staatswissenschaft*, **121** (1965), 301–324, 667–689
- [41] Leo K. Simon and Maxwell B. Stinchcombe: Extensive form games in continuous time: pure strategies, *Econometrica*, **57** (1989), 1171–1214
- [42] Mathias Staudigl and Jan-Henrik Steg: On repeated games with imperfect public monitoring: from discrete to continuous time, Preprint, 2014, available at SSRN: <http://ssrn.com/abstract=2497608>