

UNIVERSITY COLLEGE LONDON

DOCTORAL THESIS

**Monitoring and Heterogeneity in
Dynamic Games**

Yves GUÉRON

*A thesis submitted in fulfilment of the requirements
for the degree of Doctor of philosophy*

in the

Department of Economics

August 20, 2013

Declaration of Authorship

I, Yves GUÉRON, declare that this thesis titled “Heterogeneity and Monitoring in Dynamic Games”, and the work presented are my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Signed: _____

Date: _____

Abstract

In this thesis we study the impact of monitoring and heterogeneity on the set of equilibria of dynamic games.

In Chapter 1 we show how heterogeneity in time preferences can help create new intertemporal incentives. Proving the folk theorem in a game with three or more players usually requires imposing restrictions on the dimensionality of the stage-game payoffs. Considering a class of games in which those restrictions do not hold, we show how to recover a folk theorem by allowing time preferences to vary across players.

In Chapters 2 and 3 we show how a small degree of imperfection in the monitoring technology can have large effects on the set of equilibria of dynamic games. We study a dynamic voluntary contribution game with irreversibility and a game with an asymptotically finite horizon. In both settings, when monitoring is perfect, players can cooperate and obtain payoffs in the repeated game that are strictly greater than the payoffs from the unique inefficient stage-game equilibrium. We show however that introducing an arbitrarily small amount of noise in the monitoring technology can cause a complete breakdown in cooperation.

Finally in Chapter 4 we investigate how information is transmitted in a revision game with one-sided incomplete information. Players aim to coordinate on an action which depends on an unknown state of the world and players can only revise their actions stochastically during a preparation stage, at the end of which the prepared action profile is implemented. Miscoordination arises from the possibility of no longer receiving revision opportunities until the deadline. We show that close to the deadline no information is transmitted and that far from the deadline the uninformed player prefers to be miscoordinated.

Acknowledgements

I would like to thank my supervisor, Martin Cripps, for his unwearied support over the last five years. He always managed to uplift my spirit at each of my numerous visits to his office and will continue to be a great source of inspiration. I was privileged to have the opportunity to work with him. I would also like to thank my second supervisor, Philippe Jehiel, who always took a keen interest in my work.

I thank Robert Evans and Antonio Guarino for accepting to be my examiners and for having the patience to wait until the completion of this manuscript. I am extremely grateful to David Myatt and Syngjoo Choi for being available as examiners on such short notice and for the very enjoyable viva.

I wish to thank Ran Spiegler, who entrusted me in reviewing his book on bounded rationality and industrial organization. This taught me how important ideas could be illustrated through simple models and I can only hope to write such inspiring research and express it so clearly. I would also like to express my gratitude to the theory group at UCL, who always provided a challenging and stimulating environment from which I greatly benefited: V. Bhaskar, Jan Eeckhout, Steffen Huck, Suehyun Kwon, Guy Laroque, Nikita Roketskiy and Georg Weizsacker.

I am also grateful to Marco Scarsini and Frédéric Koesler who gave me the opportunity to present my work in front of distinguished audiences at the Workshop on Stochastic Methods in Game Theory, Erice, and at the Paris School of Economics. Presenting my work to such audiences has without a doubt helped me improve it.

I probably would not have done a Ph.D. in Economics if it were not for Jean Lainé, who introduced me to microeconomic theory while at the ENSAI and helped me find my way from statistics to economics. After graduating from the the LSE, he said to me as I was leaving his office: “Congratulations, now you really are an economist.” While it is still unclear whether this is true or not, I will always strive to honour the faith he had in me. I am also

grateful to Alain Trognon, who has been a fantastic teacher and also greatly contributed to my transition from statistics to economics. I also wish to thank Gérard Kremer, who first gave me the opportunity to apply my knowledge of statistics and economics in a professional environment and took good care of me at the French Central Bank.

I wish to thank the many PhD students for stimulating interactions: Alex, Andreas, Andres, Antoine, Brendon, Boris, Christian, Dan, Eeva, Florian, Italo, Jan, Jaime, Kari, Lena, Lucia, Luigi, Manzur, Marieke, Michael ($\times 3$), Nicolas ($\times 2$), Renata, Roberta and Sami. I thank Andrea for providing a much needed balance throughout my Ph.D., Marc for the many stimulating discussions over lunches, and Xavier for always answering my math questions on Google chat. I am particularly thankful to Caroline and Thibaut, with whom the first chapter of this thesis was written. It was a great experience writing a paper together and having our first publication, and I hope more collaboration will follow.

I also wish to thank Nicolas and Yohan, who helped me define my long-term goals as a student of ENSAI and still continue to do so to this day.

I thank Sophie, who inspires me to become a better person every day. Finally, I am eternally grateful to my parents and my brother, who always provided me with a safe, loving and stimulating home environment.

Contents

Abstract	2
Acknowledgements	3
Introduction	7
1 Repeated Games with One-Dimensional Payoffs and Different Discount factors	10
1.1 Introduction	10
1.1.1 An Example	12
1.1.2 Notation	13
1.1.3 Main Result and Outline of the Proof	14
1.2 Lowest Equilibrium Payoffs	16
1.2.1 Strategy Profiles and Incentive Compatibility Constraints	16
1.2.2 Proof of Theorem 1.1	18
1.3 Conclusion	28
1.3.1 Two intermediate players have the same discount factor	28
1.3.2 Generalization	31
2 Failure of Gradualism under Imperfect Monitoring	33
2.1 Introduction	33
2.2 Public monitoring	38
2.2.1 The stage game	39
2.2.2 The dynamic game and monitoring structure	40
2.2.3 Private and public histories	43
2.2.4 Stochastic process of public signals	44
2.2.5 Notation	45
2.3 Main Results under Perfect Monitoring	45
2.4 Failure of Cooperation	46

2.4.1	Outline of the Proof	46
2.4.2	Preliminary Results	48
2.4.3	Proof of Theorem 2.1	50
2.5	All-or-Nothing Monitoring	56
2.6	Conclusion	60
3	Repeated Games with Asymptotically Finite Horizon and Im-	
	perfect Monitoring	62
3.1	Introduction	62
3.2	Perfect monitoring	64
3.2.1	The stage game	64
3.2.2	The dynamic structure	65
3.2.3	Cooperation in repeated games with asymptotically fi-	
	nite horizons and perfect monitoring	66
3.3	Public monitoring	67
3.4	Breakdown of cooperation	68
3.4.1	Preliminary results	69
3.4.2	Proof of theorem 3.2	71
3.5	Conclusion	72
3.5.1	Tapering off	73
3.5.2	Private monitoring	74
3.6	Proof of Lemma 3.3	74
4	Revision Games with One-Sided Incomplete Information	77
4.1	Introduction	77
4.2	Setting	79
4.2.1	The Bayesian game	79
4.2.2	Timing and revision opportunities	79
4.2.3	Belief updating of Player 2	80
4.2.4	Histories	81
4.2.5	Strategies	82
4.2.6	Equilibrium	82
4.3	Informative equilibria	82
4.3.1	No signalling close to the deadline	83
4.3.2	The uninformed player can prefer miscoordination	87
4.4	Conclusion	94
	Bibliography	98

Introduction

Repeated interactions are pervasive in economic exchanges: the same firms compete against each other over time; consumers often use the same service providers, go to the same shops, or eat in the same restaurants; regulators set up long-term rules to control financial institutions or natural monopolies; policy makers seek re-election from the same group of voters. Dynamic games model such situations in which agents interact repeatedly over time.

Repeated interactions give rise to incentives that are not found in one-shot interactions and can help players cooperate and coordinate. For example, consider the Prisoner's Dilemma presented in Figure 1. When played once,¹ both players have strict incentives to play D (defect), yielding the inefficient outcome $(1,1)$. When repeated an infinite number of times and if players are sufficiently patient the following "grim-trigger" strategy yields the efficient outcome: start by playing C and play C until one of the players play D ; then play D forever. Indeed, the average payoff from this strategy is 2, while a deviation will lead to an instantaneous gain of 1 followed by a loss of 1 in each future period. Such a deviation is not profitable if players are sufficiently patient.²

	C	D
C	2,2	0,3
D	3,0	1,1

Figure 1: The Prisoner's Dilemma

¹Or any finite number of times

²More specifically if $\delta \geq 1/2$.

In repeated games, the same strategic interaction is repeated over time. The folk theorem then states that if players are patient enough, any feasible and strictly individually rational payoff can be obtained as the payoff of a subgame-perfect equilibrium. In particular, intertemporal incentives allow players to cooperate and reach efficient outcomes where it might not be possible in one-shot interactions, as illustrated previously with the Prisoner's Dilemma.

In Chapter 1 we study the role of heterogeneity in the folk theorem. In particular, in games of three or more players, the folk theorem fails when players have perfectly aligned preferences. This is because the individual rewards and punishments that are needed to provide incentives to each player are no longer available given that all players share the same payoffs.³ The folk theorem therefore usually requires the assumption that the set of stage-game payoffs has non-empty interior, or full-dimensionality. In this first chapter we consider a game with three or more players in which preferences are perfectly aligned, so that the interior of stage-game payoffs is empty. By introducing heterogeneity in the discount rates, we are able to re-create individual intertemporal incentives and recover the folk theorem.⁴

While in Chapter 1 we consider games in which players perfectly observe each other's actions, in many economic situations this is not the case and players only observe noisy signals of each other's actions. For example two competing firms which choose how much to produce might only observe the equilibrium market price, which can be a random function of the total quantity produced. The folk theorem still holds in repeated games with imperfect monitoring, provided players observe a public signal with sufficient statistical information. In Chapters 2 and 3 we show that this is not necessarily the case in dynamic games and illustrate how the introduction of an arbitrarily small amount of noise generates important discontinuities in the equilibrium payoff

³This does not matter in two-player games as players have the ability to "mutually min-max" each other. In our repeated Prisoner's Dilemma example this is what happened.

⁴This chapter is based on joint work with Thibaut Lamadon and Caroline Thomas, both Ph.D. students at UCL at the time of the research. All other chapters are sole-authored.

set. Indeed, we show that cooperation is no longer feasible.

In dynamic games, players do not necessarily face the same strategic interaction over time. For example, two firms might repeatedly interact in a market with a declining demand; or players' actions might be constrained by previous choices, if players for example make some irreversible investments. In Chapter 2 we consider a dynamic contribution game in which past contributions are non-refundable. This irreversibility makes the dynamic game non-stationary. In Chapter 3 we consider a repeated game in which the probability of further interactions occurring, although always positive, declines over time, again making the environment non-stationary. Both stage-games have a continuum of actions and a unique and inefficient stage-game Nash equilibrium.

When players perfectly observe each other's actions, cooperation is feasible and player can obtain payoffs strictly higher than those of the unique stage-game Nash equilibrium. We show that introducing an arbitrarily small amount of smooth noise in the monitoring causes a breakdown in cooperation and that players are reduced to playing the unique stage-game Nash equilibrium in each period.

Repeated interactions sometimes allow for players to coordinate in order to reach efficient outcomes. In Chapter 4 we study a revision game with asymmetric information and look at how information is transmitted by the more informed player. Players aim to coordinate on an action which depends on a state of the world which is known only by one of the players. Players can only revise their actions stochastically during a preparation stage, at the end of which the prepared action profile is implemented. Miscoordination arises from the possibility of no longer receiving revision opportunities until the deadline. We show that close to the deadline no information is transmitted, and that far away from the deadline, the uninformed player prefers to be miscoordinated.

Chapter 1

Repeated Games with One-Dimensional Payoffs and Different Discount factors

This chapter is based on joint work with Thibaut Lamadon and Caroline Thomas.

1.1 Introduction

For the folk theorem to hold with more than two players, it is necessary to have the ability to threaten any single player with a low payoff, while also offering rewards to the punishing players. In assuming full dimensionality of the convex hull of the set of feasible stage-game payoffs, Fudenberg and Maskin (1986) guarantee that those individual punishments and rewards exist. Abreu et al. (1994) show that the weaker NEU condition (“nonequivalent utilities”), whereby no two players have identical preferences in the stage-game, is sufficient for the folk theorem to hold.

When the NEU condition fails, players that have equivalent utilities can no longer be individually punished in equilibrium. Wen (1994) introduces the notion of *effective minmax* payoff, which takes into account the fact that

when a player is being minmaxed, another player with equivalent utility might unilaterally deviate and best respond. The effective minmax payoff of a player cannot be lower than his individual minmax payoff (when NEU is satisfied, they coincide), and Wen shows that when NEU fails it is the effective minmax that constitutes the lower bound on subgame-perfect equilibrium payoffs. He establishes the following folk theorem: when players are sufficiently patient, any feasible payoff vector can be supported as a subgame-perfect equilibrium, provided it dominates the effective minmax payoff vector. We show that this can be relaxed by allowing for unequal discounting.

As pointed out by Lehrer and Pauzner (1999), when players have different discount factors, the set of feasible payoffs in a two-player repeated game is typically larger and of higher dimensionality than the set of feasible stage-game payoffs.¹ In a particular three-player game in which two players have equivalent utilities, Chen (2008) illustrates how with unequal discounting payoffs below the effective minmax may indeed be achieved in equilibrium for one of the players.

In this chapter, we explore the notion that unequal discounting restores the ability to punish players individually in an n -player game where all players have equivalent utilities. Our result is stronger than Chen's as we show that all players can be held down to their individual minmax payoff in equilibrium. Moreover we argue that our result holds for all possible violations of NEU.

¹Mailath and Samuelson (2006, Remark 2.1.4) present a simple example to show how the set of feasible payoffs can increase when allowing for different discount factors. Consider the game of battle of the sexes depicted in Figure 1.1 and assume that players have different discount factors, $\delta_1 > \delta_2$. Consider an outcome in which (B, R) is played for T periods while (T, L) is played in subsequent periods. That is, first the less patient player is favored while the more patient player is rewarded subsequently. The payoffs to player 1 and 2 from this outcome are $(1 - \delta_1^T) + 3\delta_1^T$ and $3(1 - \delta_2^T) + \delta_2^T$, which is outside the convex hull of the set $\{(3, 1), (0, 0), (1, 3)\}$ because $\delta_1 > \delta_2$.

	L	R
T	$3, 1$	$0, 0$
B	$0, 0$	$1, 3$

Figure 1.1: Battle of the sexes

We find that a small difference in the discount factors suffices to hold a player to his individual minmax for a certain number of periods while still being able to reward the punishing players. For discount factors sufficiently close to one, any strictly individually rational payoff, including those dominated by the effective minmax payoff, can be obtained as the outcome of a subgame-perfect equilibrium with public correlation, restoring the validity of the folk theorem.

Although our result is stated for games where all players have equivalent utilities, we conjecture that it extends to weaker violations of NEU, as long as any two players with equivalent utilities have different discount factor. The intuition behind this conjecture is that following Abreu et al. (1994) we could design specific punishments for each group of players with equivalent utilities and use the difference in discount factors within each group to enforce those specific punishments.

1.1.1 An Example

	<i>L</i>	<i>R</i>		<i>L</i>	<i>R</i>
<i>T</i>	1,1,1	0,0,0	<i>T</i>	0,0,0	0,0,0
<i>B</i>	0,0,0	0,0,0	<i>B</i>	0,0,0	1,1,1
	<i>C</i>			<i>D</i>	

Figure 1.2: A stage game with one-dimensional payoffs

Consider the stage-game in Figure 1.2, where Player 1 chooses rows, Player 2 chooses columns and Player 3 chooses matrices. This stage-game is infinitely repeated and the players evaluate payoff streams according to the discounting criterion. When the players share a common discount factor $\delta < 1$, Fudenberg and Maskin (1986, Example 3) show that any subgame-perfect equilibrium yields a payoff of at least $1/4$ (the effective minmax) to each player, whereas the individual minmax payoff of each player is zero.² The low dimensionality of the set of stage-game payoffs weakens the punishment that can be imposed on a player as another player with equivalent utility can deviate and best respond. The

²For example, when Player 1 plays *T* and Player 2 plays *R*, Player 3 gets a payoff of 0 whether he plays *C* or *D*.

inability to achieve subgame-perfect equilibrium payoffs in $(0, 1/4)$ means that the “standard” folk theorem fails in this case.³

We show however that if all three players have different discount factors, there exists a subgame-perfect equilibrium in which the payoff to each player is arbitrarily close to zero, the individual minmax, provided that the discount factors are sufficiently close to one. Any payoff in the interval $(0, 1/4)$ can then be achieved in equilibrium, restoring the validity of the folk theorem in the context of this game.

1.1.2 Notation

We consider an n -player repeated game, where all players have equivalent utilities. We normalize payoffs to be in $\{0, 1\}$ and let each player’s individual minmax payoff be zero.⁴ We use public correlation to convexify the payoff set, although we argue later that this assumption can be dispensed with. Players have different discount factors, and are ordered according to their patience level: $0 < \delta_1 < \dots < \delta_{n-1} < \delta_n < 1$.⁵ We use an exponential representation of discount factors: $\forall i, \delta_i := e^{-\Delta\rho_i}$, where $\Delta > 0$ could represent the length of time between two repetitions of the stage game. As $\Delta \rightarrow 0$, all discount factors tend to one. The ρ ’s are strictly ordered: $0 < \rho_n < \dots < \rho_2 < \rho_1$. We assume that the stage game has a (mixed) Nash equilibrium which yields a payoff $Q < 1$ to all players.⁶

We summarize our assumptions about the game and introduce a notation for the lowest subgame-perfect equilibrium payoff of a player i in the following definitions:

³One may not be too concerned about our inability to achieve low payoffs. However if the game of Figure 1.2 is part of a more general game then our inability to reach low payoffs (that is, to punish players) might reduce the scope for cooperation in the more general game.

⁴We only use two payoffs as we only need to consider the minmax payoff and the maximum possible payoff.

⁵Note that the result no longer holds if several but not all players have the same discount factor. We address this point in Section 1.3.1.

⁶For example in the game of Figure 1.2, the mixture $\left\{ (1/2, 1/2), (1/2, 1/2), (1/2, 1/2) \right\}$ is a Nash equilibrium that yields a payoff of $1/4$.

Definition 1.1. Let $\Gamma(\Delta)$ be the set of n -player infinitely repeated games such that:

- A1. The set of stage-game payoffs is one-dimensional and all players receive the same payoff in $\{0, 1\}$.
- A2. The stage game has a mixed-strategy Nash equilibrium which yields a payoff of $Q < 1$ to all players.
- A3. Each player's pure action individual minmax payoff is zero.
- A4. Players evaluate payoff streams according to the discounting criterion, and discount factors are strictly ordered: $0 < \delta_1 < \dots < \delta_n < 1$, where $\delta_i := e^{-\Delta\rho_i}$.

Note that the stage game of Figure 1.2 satisfies assumptions A1 to A3 of Definition 1.1.

Definition 1.2. We denote by a_i the lowest subgame-perfect equilibrium payoff of Player i in a game $G_\Delta \in \Gamma(\Delta)$.

For given discount factors, the existence of the $(a_i)_{i=1,\dots,n}$ is ensured by the compactness of the set of subgame-perfect equilibrium payoffs (see Fudenberg and Levine (1983, Lemma 4.2)).

1.1.3 Main Result and Outline of the Proof

Our main result, Theorem 1.1, states that for games in $\Gamma(\Delta)$, the lowest subgame-perfect equilibrium payoff of each player goes to zero (the common individual minmax payoff) as discount factors tend to one:

Theorem 1.1. *Consider an n -player infinitely repeated game $G_\Delta \in \Gamma(\Delta)$. Then $a_i \in O(\Delta)$ for all i .⁷*

⁷That is, $\exists M \geq 0$ and $\Delta^* > 0$ such that $a_i \leq M \cdot \Delta$ for $\Delta \leq \Delta^*$.

Theorem 1.1 states that for discount factors sufficiently close to one (that is for Δ sufficiently close to zero), the lowest subgame-perfect equilibrium payoff of each player i , a_i , is arbitrarily close to zero. We do not provide a full characterization of the set of subgame-perfect equilibrium payoffs but note that any feasible and strictly individually rational payoff is a subgame-perfect equilibrium payoff. In recent work, Sugaya (2010) characterises the set of perfect and public equilibrium payoffs in games with imperfect public monitoring when players have different discount factors, under a full-dimensionality assumption.

To prove Theorem 1.1, we first show that when stage-game payoffs are identical, the lowest subgame-perfect equilibrium payoffs are ordered according to the discount factors (Lemma 1.1). A player's lowest subgame-perfect equilibrium payoff cannot be below that of another player who is less patient. We then show that the lowest subgame-perfect equilibrium payoffs of the two most patient players (Player $n - 1$ and Player n) are arbitrarily close to each other when discount factors tend to one (Lemma 1.2). This is done by explicitly constructing a subgame-perfect equilibrium of the repeated game.

In a similar way, we then construct a set of subgame-perfect equilibria (one for each player $i \in \{2, \dots, n - 1\}$) (Lemma 1.3) and use those to bound the distance between the lowest subgame-perfect equilibrium payoffs of players i and $i - 1$ (Lemma 1.4). We then show by induction that the lowest subgame-perfect equilibrium payoffs of *any* two players are arbitrarily close to each other as discount factors tend to one (Lemma 1.5). Finally we show that Player 1's lowest subgame-perfect equilibrium payoff can be made arbitrarily close to zero as discount factors tend to one (Lemma 1.6). We are then able to conclude and prove Theorem 1.1.

Note that the assumption of strictly different discount factors cannot be dispensed with. In particular our result does not hold when some but not all player share a common discount factor. In a similar fashion to Fudenberg and Maskin (1986, Example 3), we construct a four-player example where the

stage game satisfies assumptions A1 to A3 but where the two “intermediate” players share a common discount factor. That is we have $\delta_1 < \delta_2 = \delta_3 < \delta_4$. This example is presented in Section 1.3.1.

1.2 Lowest Equilibrium Payoffs

1.2.1 Strategy Profiles and Incentive Compatibility Constraints

To prove Theorem 1.1, we explicitly construct several subgame-perfect equilibria of the repeated game. To do so, we consider strategy profiles that give a constant expected stage-game payoff between zero and one (using public correlation) to all players for a given number of periods, and then stage-game payoffs of one forever:

Definition 1.3. Let $\sigma(\mu, \tau, i)$ be the strategy profile such that:

- (i) For τ periods, in each stage-game, players use a public correlating device to generate an expected payoff of μ . When the public correlating device generates a payoff of zero, players minmax Player i .
- (ii) In all subsequent periods $t > \tau$, players play an action profile yielding a stage-game payoff of 1 to each player.
- (iii) During the first τ periods, deviations by Player i are ignored. After that, if Player i deviates from the equilibrium path, players play a subgame-perfect equilibrium which gives Player i his lowest possible payoff, a_i .
- (iv) If a deviation by Player $j \neq i$ occurs at any time, players then play a subgame-perfect equilibrium which gives Player j his lowest possible payoff, a_j .

Assuming that the correlating device generates a payoff of zero at $t = 0$, a

player $j \neq i$ will not have an incentive to deviate from $\sigma(\mu, \tau, i)$ if:^{8,9}

$$(1 - \delta_j) + \delta_j a_j \leq \delta_j \left((1 - \delta_j^{\tau-1}) \mu + \delta_j^{\tau-1} \right), \quad (1.1)$$

which can be rewritten as

$$\delta_j^\tau \geq \frac{1 - \delta_j + \delta_j a_j - \delta_j \mu}{1 - \mu}. \quad (1.2)$$

To prove Theorem 1.1, we show that there exists a “low” μ and a large τ such that for Δ sufficiently close to zero, the strategy profile $\sigma(\mu, \tau, i)$ is subgame perfect, that is, we show that (1.2) is satisfied for any $j \neq i$. To do so, we identify the player with the tightest incentive compatibility constraint as j_i^* and find the largest τ such that (1.2) is satisfied for Player j_i^* (Lemma 1.3). Notice that Player j_i^* is not necessarily the player with the lowest discount factor. By a “low” μ we mean that μ must be close to a_{i-1} . To this end, we choose a stage-game payoff μ_i that is slightly above a_{i-1} :

Definition 1.4. For all $i \in \{1, \dots, n\}$, let μ_i be such that:¹⁰

$$\mu_i = \begin{cases} a_{i-1} + \frac{1-\delta_1}{\delta_1} & \text{if } 2 \leq i \leq n, \\ 0 & \text{if } i = 1. \end{cases}$$

To illustrate, consider a player i with intermediate patience, such that $1 < i < n$. The strategy profile $\sigma(\mu, \tau, i)$ does not give him an opportunity to deviate, as he is being minmaxed when payoffs of zero are generated. For this reason, that strategy profile can be thought of as the other players colluding

⁸First note that zero is the lowest possible stage game payoff and so if it is enforceable all other payoffs will be. Second the strategy starts by giving zeros and ones and then rewards the players with ones forever, so the tightest incentive compatibility constraint will be when $t = 0$ as for $t > 0$ players are closer to getting ones for ever.

⁹The left-hand side of (1.1) is the payoff to Player j if he deviates: he get an instantaneous payoff of 1 followed by a repeated game payoff of a_j . If Player j follows the strategy he gets a payoff of zero today, followed by $\tau - 1$ periods during which he gets an expected payoff of μ , after which he receives a payoff of one in each period.

¹⁰Note that for all i and for Δ sufficiently close to zero, $\mu_i \leq 1$. Indeed, $\mu_i \leq Q + \frac{1-\delta_1}{\delta_1} \rightarrow_{\Delta \rightarrow 0} Q < 1$.

against player i . Lowering the payoff to player i from that strategy profile may conflict with making it incentive compatible both for players that are more and less patient than him. Players less patient than i must get a payoff sufficiently higher than their lowest SPE payoff, and players more patient than i must be promised payoffs of 1 soon enough to make them accept an early stream of low payoffs. We show that these constraints can be reconciled with keeping player i 's payoff very close to the lowest equilibrium payoff of the player just less patient than him.

1.2.2 Proof of Theorem 1.1

In a first step towards Theorem 1.1 we now show that the lowest subgame-perfect equilibrium payoffs are ordered according to the discount factors (Lemma 1.1), and that Player n 's lowest subgame-perfect equilibrium payoff is arbitrarily close to Player $n - 1$'s for Δ close enough to zero (Lemma 1.2).

Lemma 1.1. $\forall i \in \{2, \dots, n\}, a_{i-1} \leq a_i$.

Proof. The main idea is to find a stream of payoffs $(z_t)_{t=0, \dots, \infty}$ in $[0, 1]^{\mathbb{N}}$ that minimizes Player i 's average discounted payoff, given Player $i - 1$ is guaranteed his lowest subgame-perfect equilibrium payoff at each stage. By definition, the resulting average discounted payoff for Player i cannot be greater than a_i . We show that the constraints imposed by Player $i - 1$'s lowest subgame-perfect equilibrium payoff must all be binding and that $z_t = a_{i-1}, \forall t \geq 0$.

Formally, we solve the following minimization problem:

$$\min_{(z_t)_{t=0, \dots, \infty} \in [0, 1]^{\mathbb{N}}} (1 - \delta_i) \sum_{t=0}^{\infty} \delta_i^t z_t \quad (1.3)$$

subject to

$$(1 - \delta_{i-1}) \sum_{t=s}^{\infty} \delta_{i-1}^{t-s} z_t \geq a_{i-1}, \quad \forall s \geq 0 \quad (1.4)$$

We show by induction that all constraints in (1.4) will be binding, which implies that $z_s = a_{i-1}, \forall s \geq 0$. Our induction hypothesis is that the con-

straints in (1.4) must bind for $s = 0, \dots, \tau$ and therefore, that the minimization problem (1.3) subject to the constraints (1.4) can be rewritten as:

$$\min_{(z_t)_{t=\tau, \dots, \infty} \in [0,1]^{\mathbb{N}}} \lambda_{\tau-1}(a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \left(\sum_{t=\tau+1}^{\infty} \delta_i^{\tau} (\delta_i^{t-\tau} - \delta_{i-1}^{t-\tau}) z_t \right) \quad (1.5)$$

subject to

$$(1 - \delta_{i-1}) \sum_{t=s}^{\infty} \delta_{i-1}^{t-s} z_t \geq a_{i-1}, \quad \forall s \geq \tau + 1 \quad (1.6)$$

where the function λ_{τ} is recursively defined by

$$\lambda_0(a_{i-1}, \delta_{i-1}, \delta_i) = (1 - \delta_i) \frac{a_{i-1}}{1 - \delta_{i-1}}$$

and

$$\lambda_{\tau}(a_{i-1}, \delta_{i-1}, \delta_i) = \lambda_{\tau-1}(a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \delta_i^{\tau} + (\delta_i - \delta_{i-1}) \frac{a_{i-1}}{1 - \delta_{i-1}}.$$

Initialization: $\tau = 0$ The first constraint is the only constraint featuring z_0 and can be rewritten as $z_0 \geq \frac{a_{i-1}}{1 - \delta_{i-1}} - \sum_{t=1}^{\infty} \delta_{i-1}^t z_t$. Moreover, z_0 enters with a positive coefficient in the objective function, therefore, the first constraint must be binding. The constraint is then used to eliminate z_0 from the objective function: the minimization problem (1.3) subject to (1.4) can therefore be written in the following way:

$$\min_{(z_t)_{t=1, \dots, \infty} \in [0,1]^{\mathbb{N}}} (1 - \delta_i) \left(\frac{a_{i-1}}{1 - \delta_{i-1}} + \sum_{t=1}^{\infty} (\delta_i^t - \delta_{i-1}^t) z_t \right)$$

subject to

$$(1 - \delta_{i-1}) \sum_{t=s}^{\infty} \delta_{i-1}^{t-s} z_t \geq a_{i-1}, \quad \forall s \geq 1$$

This verifies (1.5) and (1.6).

Induction We assume that our minimization problem can be rewritten as (1.5) subject to (1.6) for some $\tau > 1$. Because $\delta_i > \delta_{i-1}$, $z_{\tau+1}$ enters with a positive coefficient in the objective function and $z_{\tau+1}$ only appears in the

constraint $z_{\tau+1} \geq \frac{a_{i-1}}{1-\delta_{i-1}} - \sum_{t=\tau+2}^{\infty} \delta_{i-1}^{t-(\tau+1)} z_t$, this constraint will be binding and the objective function can be rewritten by substituting for $z_{\tau+1}$ as follows:

$$\begin{aligned}
& \lambda_{\tau-1} (a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \left(\sum_{t=\tau+1}^{\infty} \delta_i^{\tau} \left(\delta_i^{t-\tau} - \delta_{i-1}^{t-\tau} \right) z_t \right) \\
&= \lambda_{\tau-1} (a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \left(\delta_i^{\tau} (\delta_i - \delta_{i-1}) \left(\frac{a_{i-1}}{1 - \delta_{i-1}} - \sum_{t=\tau+2}^{\infty} \delta_{i-1}^{t-(\tau+1)} z_t \right) \right) \\
&\quad + (1 - \delta_i) \sum_{t=\tau+2}^{\infty} \delta_i^{\tau} \left(\delta_i^{t-\tau} - \delta_{i-1}^{t-\tau} \right) z_t \\
&= \lambda_{\tau} (a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \sum_{t=\tau+2}^{\infty} \left(\delta_i^{\tau} \left(\delta_i^{t-\tau} - \delta_{i-1}^{t-\tau} \right) - \delta_i^{\tau} (\delta_i - \delta_{i-1}) \delta_{i-1}^{t-(\tau+1)} \right) z_t \\
&= \lambda_{\tau} (a_{i-1}, \delta_{i-1}, \delta_i) + (1 - \delta_i) \left(\sum_{t=\tau+2}^{\infty} \delta_i^{\tau+1} \left(\delta_i^{t-(\tau+1)} - \delta_{i-1}^{t-(\tau+1)} \right) z_t \right),
\end{aligned}$$

where the first equality is obtained by substituting for $z_{\tau+1}$ and the other equalities are obtained by grouping the terms in z_t ($t \geq \tau + 2$) together. Thus (1.5) and (1.6) hold for $\tau + 1$ also.

This concludes the proof by induction and so all constraints in (1.4) must bind: $(1 - \delta_{i-1}) \sum_{t=s}^{\infty} \delta_{i-1}^{t-s} z_t = a_{i-1}$, $\forall s \geq 0$. We now show that this implies that $z_s = a_{i-1}$, $\forall s \geq 0$. Consider the constraint for some $s \geq 0$:

$$\begin{aligned}
a_{i-1} &= (1 - \delta_{i-1}) \sum_{t=s}^{\infty} \delta_{i-1}^{t-s} z_t \\
&= (1 - \delta_{i-1}) \left\{ z_s + \delta_{i-1} \sum_{t=s+1}^{\infty} \delta_{i-1}^{t-(s+1)} z_t \right\} \\
&= (1 - \delta_{i-1}) \left\{ z_s + \frac{\delta_{i-1}}{1 - \delta_{i-1}} a_{i-1} \right\},
\end{aligned}$$

where the last inequality holds because the constraint is binding for $s + 1$. This implies that $z_s = a_{i-1}$, $\forall s \geq 0$.

Given the constraints imposed on stage-game payoffs by player $i - 1$'s lower subgame-perfect equilibrium bound, the lowest average discounted payoff which can be given to player i is a_{i-1} . We therefore have $a_{i-1} \leq a_i$. \square

Lemma 1.2. $|a_n - a_{n-1}| \in O(\Delta)$.

Proof. Consider the strategy profile $\sigma(\mu_n, \infty, n)$, where $\mu_n = a_{n-1} + \frac{1-\delta_1}{\delta_1}$. We are going to show that this constitutes a subgame-perfect equilibrium.

First, note that in a period in which the public correlating device generates a payoff of one, no player has a one-shot profitable deviation. Secondly, because Player n is being minmaxed in a period in which the public correlating device generates a payoff of zero, he doesn't have a profitable one-shot deviation. Thirdly, because punishment phases consist of subgame-perfect equilibrium strategies, no player has a profitable one-shot deviation during one of those. Thus, to verify that $\sigma(\mu_n, \infty, n)$ is subgame perfect, we only need to check that players $i \leq n-1$ do not have profitable one-shot deviations when the public correlating device generates a payoff of zero.

A deviation from Player $i \leq n-1$ leads at most to a one-off gain of one followed by a payoff of a_i forever. Therefore, there is no one-shot profitable deviation if $(1-\delta_i) + \delta_i a_i \leq \delta_i \left(a_{n-1} + \frac{1-\delta_1}{\delta_1} \right)$, where the right-hand-side is the repeated game payoff to Player i if the public correlation device indicates a zero payoff action profile in that period. This inequality is always satisfied for $i \leq n-1$ as $a_i \leq a_{n-1}$ (Lemma 1.1) and as $\frac{1-\delta_i}{\delta_i} \leq \frac{1-\delta_1}{\delta_1}$.

By definition of a_n , and by Lemma 1.1, we have that $a_{n-1} \leq a_n \leq a_{n-1} + \frac{1-\delta_1}{\delta_1}$. We conclude the proof by noting that $a_n - a_{n-1} \leq \frac{1-\delta_1}{\delta_1}$ and that $\frac{1-\delta_1}{\delta_1} \in O(\Delta)$. \square

We have shown that the lowest subgame-perfect equilibrium payoffs of the two most patient players are arbitrarily close as Δ tends to zero. The intuition behind this result is that all players can collude against Player n by minmaxing him whenever the public correlating device generates a payoff of zero. Since Player $n-1$ is the most patient of the colluding players and since lowest subgame-perfect equilibrium payoffs are ordered according to discount factors, his lowest subgame-perfect equilibrium will determine by how much Player n 's equilibrium payoff can be pushed down.

We now show that the lowest subgame-perfect equilibrium payoffs of *any* two players are arbitrarily close to each other as Δ tends to zero (Lemma

1.5). We start by identifying bounds on Player $i > 1$'s lowest subgame-perfect equilibrium payoff. To do this, we find the largest time $\tau \geq 1$ such that the strategy profile $\sigma(\mu_i, \tau, i)$ is a subgame-perfect equilibrium and compute its equilibrium payoff for Player i . We then prove Lemma 1.5 by induction.

First, we introduce some useful notation. For every player $i \in \{1, \dots, n-1\}$, define

$$N_+^i := \{j > i : 1 - \delta_j + \delta_j a_j - \delta_j \mu_i > 0\}.$$

When proving that for a particular τ , $\sigma(\mu_i, \tau, i)$ is a subgame-perfect equilibrium, N_+^i should be thought of as the set of players for whom profitable deviations might exist depending on the value of τ . That is, N_+^i is the set of players for whom the right-hand side of (1.2) (when replacing μ with μ_i) is strictly positive. We will therefore choose τ to satisfy the no-deviation constraints of all players in N_+^i . When N_+^i is not empty, we identify the player from this set with the tightest constraint as j_i^* and we define \tilde{t}_i as follows:

$$j_i^* := \arg \min_{j \in N_+^i} \frac{\log \left((1 - \delta_j + \delta_j a_j - \delta_j \mu_i) / (1 - \mu_i) \right)}{\log \delta_j},$$

$$\tilde{t}_i := \frac{\log \left((1 - \delta_{j_i^*} + \delta_{j_i^*} a_{j_i^*} - \delta_{j_i^*} \mu_i) / (1 - \mu_i) \right)}{\log \delta_{j_i^*}}.$$

Let $t_i^* := \lfloor \tilde{t}_i \rfloor$ be the largest integer smaller or equal than \tilde{t}_i and define $r_i \in (0, 1)$ to be the fractional part of \tilde{t}_i :

$$r_i := \tilde{t}_i - t_i^*.$$

Note that t_i^* is the longest time τ such that j_i^* does not have a profitable one-shot deviation in $\sigma(\mu_i, \tau, i)$.

In Lemma 1.3 we show that for Δ sufficiently close to zero t_i^* is well defined and arbitrarily large and that the strategy profile $\sigma(\mu_i, t_i^*, i)$ is indeed subgame perfect.

Lemma 1.3. *Let $i \in \{2, \dots, n-1\}$, and assume that $N_+^i \neq \emptyset$. Given j_i^* , t_i^* and μ_i , $\exists \Delta_i^* > 0$ such that for $\Delta \in (0, \Delta_i^*)$, $\sigma(\mu_i, t_i^*, i)$ constitutes a subgame-perfect equilibrium.*

Proof. For notational convenience, we omit the i subscript on j_i^* , \tilde{t}_i , t_i^* , and r_i . First, recall that for Δ sufficiently close to zero, $\mu_i \leq 1$.¹¹ We now check that t^* is well defined. Note that $\exists \Delta_{ij} > 0$ and $\eta_{ij} < 1$ such that for $\Delta \leq \Delta_{ij}$, $\frac{1-\delta_j+\delta_j a_j-\delta_j \mu_i}{1-\mu_i} < \eta_{ij}$.¹² Because η_{ij} does not depend on Δ , this shows that $\lim_{\Delta \rightarrow 0} \tilde{t} = \infty$ and ensures that $\exists \Delta_i^* > 0$ such that t^* is well defined and strictly positive for $\Delta \in (0, \Delta_i^*)$.

Because i is being minmaxed if the public correlating device generates a payoff of zero, i does not have a profitable one-shot deviation. Also, no player will have a profitable one-shot deviation during the punishment phases of $\sigma(\mu_i, t_i^*, i)$, as those are subgame perfect.

We now check that no player $j \neq i$ has a profitable one-shot deviation, that is, we check that (1.1) (when replacing μ with μ_i and τ with t^*) holds for all players $j \neq i$:

$$(1 - \delta_j) + \delta_j a_j \leq \delta_j \left((1 - \delta_j^{t^*-1}) \mu_i + \delta_j^{t^*-1} \right). \quad (1.7)$$

We first check that (1.7) holds for players $j \leq i-1$ and then for players $j > i$:

- (i) No deviation from player $j \leq i-1$: Note that because $\mu_i \in [0, 1]$, we have that $\mu_i \leq (1 - \delta_j^{t^*-1}) \mu_i + \delta_j^{t^*-1}$. In order to show that (1.7) holds, we can therefore show that $(1 - \delta_j) + \delta_j a_j \leq \delta_j \mu_i$, which is equivalent to $\frac{1-\delta_j}{\delta_j} + a_j \leq a_{i-1} + \frac{1-\delta_1}{\delta_1}$. This inequality holds $\forall j \leq i-1$, as $\frac{1-\delta_j}{\delta_j} \leq \frac{1-\delta_1}{\delta_1}$ and $a_j \leq a_{i-1}$.

¹¹See footnote 10.

¹²Since $a_j \leq Q$, $\frac{1-\delta_j+\delta_j a_j-\delta_j \mu_i}{1-\mu_i} \leq \delta_j \frac{Q-\mu_i}{1-\mu_i} + \frac{1-\delta_j}{1-Q-(1-\delta_1)/\delta_1}$. For any x in $[0, 1)$, $\frac{Q-x}{1-x} \leq Q$, thus the right-hand-side of the previous inequality is bounded from above by $\delta_j Q + \frac{1-\delta_j}{1-Q-(1-\delta_1)/\delta_1}$, which tends to $Q < 1$ as Δ tends to zero.

(ii) No deviation from player $j > i$: We can rearrange (1.7) to get

$$\delta_j^{t^*} \geq \frac{1 - \delta_j + \delta_j a_j - \delta_j \mu_i}{1 - \mu_i}. \quad (1.8)$$

First, note that if $j \notin N_+^i$ then j has no incentive to deviate as $\delta_j^{t^*} > 0 \geq \frac{1 - \delta_j + \delta_j a_j - \delta_j \mu_i}{1 - \mu_i}$. Now let $j \in N_+^i$. Since t^* has been chosen such that (1.8) is satisfied for player j^* , (1.8) is also satisfied for all other players in N_+^i , and no player $j \in N_+^i$ will have an incentive to deviate.

We conclude that for Δ sufficiently close to zero, $\sigma(\mu_i, t_i^*, i)$ is a subgame-perfect equilibrium. \square

Remark 1.1 (Dispensability of public correlation). In Lemma 1.3, we show that $\sigma(\mu_i, t_i^*, i)$ is a subgame-perfect equilibrium and that t_i^* goes to infinity as Δ approaches zero. Instead of using the strategy $\sigma(\mu_i, t_i^*, i)$, which relies on public correlation, we can consider a deterministic strategy that alternates between $t_{i,1}^*$ zeros and $t_{i,2}^*$ ones, where $t_{i,1}^* + t_{i,2}^* = t_i^*$ and $t_{i,2}^*/t_i^*$ is arbitrarily close to μ_i , starting with a payoff of zero. This is possible because t_i^* goes to infinity. Intuitively, as Δ goes to zero, such a strategy will yield a payoff to any player arbitrarily close to the payoff from $\sigma(\mu_i, t_i^*, i)$, while having a period-zero incentive compatibility constraint less stringent than (1.7) since μ_i is promised on average over the first t_i^* periods and the first period payoff is a zero. This should ensure that Lemmas 1.3 and 1.4 still hold under such a deterministic strategy. \diamond

We now compute the payoff of player i from $\sigma(\mu_i, t_i^*, i)$ in order to bound the distance between a_i and a_{i-1} .

Lemma 1.4. $\forall i \in \{2, \dots, n-1\}$, we have that either:

- (i) $\forall j > i$, $|a_j - a_{i-1}| \in O(\Delta)$, or
- (ii) $|a_i - a_{i-1}| \in O(\Delta) + O(a_{j_i^*} - a_i)$, where $j_i^* > i$.

Proof. Again, for notational convenience, we omit the i subscript on j_i^* , t_i^* and r_i . If N_+^i is empty we directly have an indication of the distance between a_j and a_{i-1} by noting that no player $j > i$ has an incentive to deviate from $\sigma(\mu_i, \tau, i)$, irrespective of τ : if $N_+^i = \emptyset$, then $\forall j > i$, $0 \leq a_j - a_{i-1} \leq \frac{1-\delta_1}{\delta_1} - \frac{1-\delta_j}{\delta_j}$, which implies that $|a_j - a_{i-1}| \in O(\Delta)$.

Assume now that $N_+^i \neq \emptyset$, so that $\sigma(\mu_i, t^*, i)$ is a subgame-perfect equilibrium. We now compute Player i 's payoff from $\sigma(\mu_i, t^*, i)$ and compare it with his lowest subgame-perfect equilibrium payoff. The payoff to Player i from the strategy profile $\sigma(\mu_i, t^*, i)$ is:

$$\begin{aligned} (1 - \delta_i^{t^*})\mu_i + \delta_i^{t^*} &= \mu_i + \delta_i^{t^*} (1 - \mu_i) \\ &= \mu_i + \delta_i^{-r} \left(\frac{1 - \delta_{j^*} + \delta_{j^*} a_{j^*} - \delta_{j^*} \mu_i}{1 - \mu_i} \right)^{\frac{\rho_i}{\rho_{j^*}}} (1 - \mu_i) \\ &\geq a_i, \end{aligned}$$

where the last inequality holds because a_i is i 's lowest subgame-perfect equilibrium payoff. This inequality can be rewritten as

$$\frac{a_i - \mu_i}{1 - \mu_i} \leq \delta_i^{-r} \left(\frac{1 - \delta_{j^*} + \delta_{j^*} a_{j^*} - \delta_{j^*} \mu_i}{1 - \mu_i} \right)^{\frac{\rho_i}{\rho_{j^*}} - 1} \left(\frac{1 - \delta_{j^*} + \delta_{j^*} a_{j^*} - \delta_{j^*} \mu_i}{1 - \mu_i} \right),$$

where $\frac{\rho_i}{\rho_{j^*}} - 1 > 0$, as $i < j^*$. Recall from the proof of Lemma 1.3 that for $\Delta \leq \Delta_{ij^*}$, $(1 - \delta_{j^*} + \delta_{j^*} a_{j^*} - \delta_{j^*} \mu_i) / (1 - \mu_i) < \eta_{ij^*}$, where $\eta_{ij^*} < 1$ does not depend on Δ . For $\Delta \leq \Delta_{ij^*}$, we therefore have:

$$\frac{a_i - \mu_i}{1 - \mu_i} \leq \delta_i^{-r} \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} \left(\frac{1 - \delta_{j^*} + \delta_{j^*} a_{j^*} - \delta_{j^*} \mu_i}{1 - \mu_i} \right).$$

The previous inequality can be rewritten as:¹³

$$\begin{aligned} a_i - a_{i-1} &\leq \frac{1 - \delta_1}{\delta_1} + \delta_i^{-r} \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} \delta_{j^*} (a_i - a_{i-1}) + \\ &\quad \delta_i^{-r} \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} \left(1 - \delta_{j^*} + \delta_{j^*} (a_{j^*} - a_i) - \delta_{j^*} \frac{1 - \delta_1}{\delta_1} \right). \end{aligned} \quad (1.9)$$

¹³By canceling the $1 - \mu_i$ and adding and subtracting $\delta_{j^*} a_i$ inside the term in parentheses.

Because

$$\lim_{\Delta \rightarrow 0} \delta_i^{-r} \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} \delta_{j^*} = \lim_{\Delta \rightarrow 0} \delta_i^{-r} \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} = \eta_{ij^*}^{\frac{\rho_i}{\rho_{j^*}} - 1} < 1,$$

there exists a $\widetilde{\Delta}_i \geq 0$ and an $R < 1$ such that for $\Delta \leq \widetilde{\Delta}_i$ we have:

$$a_i - a_{i-1} \leq \frac{1 - \delta_1}{\delta_1} + R(a_i - a_{i-1}) + R \left(1 - \delta_{j^*} + \delta_{j^*} (a_{j^*} - a_i) - \delta_{j^*} \frac{1 - \delta_1}{\delta_1} \right).$$

To conclude, note that $\frac{1 - \delta_1}{(1 - R)\delta_1} + \frac{R}{1 - R} \left(1 - \delta_{j^*} - \delta_{j^*} \frac{1 - \delta_1}{\delta_1} \right)$ is of order Δ , and that $\frac{R}{1 - R} \delta_{j^*} (a_{j^*} - a_i) \in O(a_{j^*} - a_i)$, as $R < 1$ is a fixed constant. \square

Recall that the difference between the two most patient players' lowest subgame-perfect equilibrium payoffs, a_n and a_{n-1} , is of order Δ (Lemma 1.2). Moreover in Lemma 1.4 we established a bound for the distance between a_{i-1} and the lowest subgame-perfect equilibrium payoff of a more patient player. We can now establish by induction that the lowest subgame-perfect equilibrium payoffs of *any* two players are arbitrarily close to each other as Δ tends to zero.

Lemma 1.5. $|a_i - a_j| \in O(\Delta), \forall (i, j)$.

Proof. By Lemma 1.2, we know that this result is true for $i, j \in \{n - 1, n\}$. We now prove this result by induction. Assume that $\forall i, j \geq k, |a_i - a_j| \in O(\Delta)$. Our aim is to show that $\forall i \geq k, |a_i - a_{k-1}| \in O(\Delta)$.

If the first statement of Lemma 1.4 holds, then we have that $\forall j > k, |a_j - a_{k-1}| \in O(\Delta)$. Moreover, $|a_k - a_{k-1}| \leq |a_k - a_j| + |a_j - a_{k-1}|$ for any $j > k$. By induction, $|a_k - a_j| \in O(\Delta)$, thus we have $|a_k - a_{k-1}| \in O(\Delta)$.

If the second statement of Lemma 1.4 holds then $\exists k^* > k$ such that $|a_k - a_{k-1}| \in O(\Delta) + O(a_{k^*} - a_k)$. From our induction hypothesis, $|a_{k^*} - a_k| \in O(\Delta)$, which implies that $|a_k - a_{k-1}| \in O(\Delta)$. Using the triangle inequality, $\forall i \geq k, |a_i - a_{k-1}| \leq |a_i - a_k| + |a_k - a_{k-1}| \in O(\Delta)$.

This shows that $\forall i, j \geq k - 1, |a_i - a_j| \in O(\Delta)$. \square

Finally, we show that the lowest subgame-perfect equilibrium payoff of Player 1 is arbitrarily close to zero as Δ tends to zero. This is done by using a proof similar to the one of Lemma 1.4, and considering the strategy profile $\sigma(0, t_1^*, 1)$.

Lemma 1.6. $a_1 \in O(\Delta)$.

Proof. We follow the same line of reasoning as in the proof of Lemma 1.3 and Lemma 1.4, using the strategy $\sigma(0, t_1^*, 1)$. As in Lemma 1.3, $\sigma(0, t_1^*, 1)$ is well defined and constitutes a subgame-perfect equilibrium. Again, for notational convenience, we omit the subscript 1.

The strategy profile $\sigma(0, t^*, 1)$ yields a payoff of $\delta_1^{t^*} = \delta_1^{-r} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}}}$ to Player 1. Because a_1 is player 1's lowest subgame-perfect equilibrium payoff, we have

$$\begin{aligned} a_1 &\leq \delta_1^{-r} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}}} \\ &= \delta_1^{-r} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}} - 1} (1 - \delta_{j^*} + \delta_{j^*} (a_{j^*} - a_1)) \\ &\quad + \delta_1^{-r} \delta_{j^*} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}} - 1} a_1. \end{aligned}$$

Because

$$\lim_{\Delta \rightarrow 0} \delta_1^{-r} \delta_{j^*} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}} - 1} = \lim_{\Delta \rightarrow 0} \delta_1^{-r} (1 - \delta_{j^*} + \delta_{j^*} a_{j^*})^{\frac{\rho_1}{\rho_{j^*}} - 1} \leq \eta_{1j^*}^{\frac{\rho_1}{\rho_{j^*}} - 1},$$

and $\eta_{1j^*}^{\frac{\rho_1}{\rho_{j^*}} - 1} < 1$ there exists an $R < 1$ and $\Delta_1^* \geq 0$ such that for $\Delta \leq \Delta_1^*$ we have

$$a_1 \leq R \left(1 - \delta_{j^*} + \delta_{j^*} (a_{j^*} - a_1) \right) + R a_1,$$

or

$$a_1 \leq \frac{R}{1 - R} \left(1 - \delta_{j^*} + \delta_{j^*} (a_{j^*} - a_1) \right).$$

We know from Lemma 1.5 that $a_{j^*} - a_1 \in O(\Delta)$, which concludes the proof, as $R < 1$ does not depend on Δ . \square

We are now able to prove Theorem 1.1:

Proof of Theorem 1.1. From Lemma 1.5 and 1.6, we have that $\forall i \in \{1, \dots, n\}$, $|a_i - a_1| \in O(\Delta)$ and $a_1 \in O(\Delta)$. Using the triangle inequality, $|a_i| \leq |a_i - a_1| + |a_1| \in O(\Delta)$. \square

1.3 Conclusion

In this chapter, we considered the set of games where the classical folk theorem does not apply because of the low dimensionality of the set of stage-game payoffs. In such setups, it is not possible to create player-specific punishments which are necessary to sustain low values of equilibrium payoffs.

We extend the setting by allowing players to have different discount factors and prove that player-specific punishments as close as desired to the player's individual minmax can be constructed. Those punishments can be used to enforce any stage-game payoff as an equilibrium payoff. This generalizes the folk theorem to games which violate NEU but where players have different discount factors. They can also be used to yield equilibrium payoffs strictly outside the convex hull of the stage-game payoffs. However, the characterization of this multidimensional boundary for the complete equilibrium pay off set is left for future research.

In the next sections, we first show that our result does require all players to have different discount factors and does not hold if two “intermediate” players share the same discount factor. We then briefly discuss subsequent research that generalizes our result.

1.3.1 Two intermediate players have the same discount factor

In this section we confirm that our result does indeed require all players to have different discount factors by means of a counter-example similar to the one presented in Section 1.1.1, but with four players. We present a particular four-player game in which player 1 and player 2 have the same discount factor $\tilde{\delta} \in$

(δ_3, δ_4) , but such that in every stage game at least one of them is guaranteed a payoff of $1/2$.

In the game of Figure 1.3, player 1 chooses a row, player 2 chooses a column, player 3 chooses between the two left matrices or the two right ones and player 4 chooses between the two top matrices or the two bottom ones. Notice that in this game, the min-max payoff of each player is 0, and there is a mixed-strategy Nash equilibrium $(1/2, 1/2, 1/2, 1)$ which yields a payoff of $1/2$, so that assumptions A1 to A3 are satisfied.

0,0,0,0	1,1,1,1	1,1,1,1	0,0,0,0
0,0,0,0	1,1,1,1	1,1,1,1	0,0,0,0
0,0,0,0	1,1,1,1	0,0,0,0	0,0,0,0
0,0,0,0	1,1,1,1	1,1,1,1	1,1,1,1

Figure 1.3: A four-player stage game with one-dimensional payoffs

Let α_i denote the probability with which player i plays his first action (either top or left). The expected payoff to all players from strategy profile $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \in [0, 1]^4$ is

$$(1 - \alpha_2)\alpha_3 + \alpha_2(1 - \alpha_3)\alpha_4 + (1 - \alpha_1)(1 - \alpha_3)(1 - \alpha_4).$$

We now show that for any stage-game action profile, at least one of player 1 or player 2 has a deviation guaranteeing him a payoff of $1/2$.

Consider the payoff from a deviation for player 1. If player 1 plays the top row ($\alpha_1 = 1$), his payoff is

$$u_1^1 = (1 - \alpha_2)\alpha_3 + \alpha_2(1 - \alpha_3)\alpha_4,$$

while his payoff from playing the bottom row ($\alpha_1 = 0$) is

$$u_1^0 = (1 - \alpha_2)\alpha_3 + \alpha_2(1 - \alpha_3)\alpha_4 + (1 - \alpha_3)(1 - \alpha_4).$$

For player 1, playing the bottom row is the best deviation (bottom indeed

weakly dominates top).

For player 2, the payoff from playing the left column ($\alpha_2 = 1$) is

$$u_2^1 = (1 - \alpha_3)\alpha_4 + (1 - \alpha_1)(1 - \alpha_3)(1 - \alpha_4),$$

while his payoff from playing the right column ($\alpha_2 = 0$) is

$$u_2^0 = \alpha_3 + (1 - \alpha_1)(1 - \alpha_3)(1 - \alpha_4).$$

We now show that $\max\{u_1^0, u_2^0, u_2^1\} \geq 1/2$ for any quadruple $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$. First let $\beta_i = 1 - \alpha_i$. We can then rewrite u_1^0 , u_2^0 and u_2^1 as $\beta_2(1 - \beta_3) + (1 - \beta_2)\beta_3(1 - \beta_4) + \beta_3\beta_4 = (1 - 2\beta_3 + \beta_3\beta_4)\beta_2 + \beta_3$, $\beta_3(1 - \beta_4) + \beta_1\beta_3\beta_4$ and $1 - \beta_3 + \beta_1\beta_3\beta_4$, respectively. As β_2 only appears in u_1^0 , we can first minimize $\max\{u_1^0, u_2^0, u_2^1\}$ with respect to β_2 . Moreover, u_1^0 is linear in β_2 , so that it's minimum is $\beta_3 + \min(0, 1 - 2\beta_3 + \beta_3\beta_4)$.

We now notice that β_4 only appears in the expression $\beta_3\beta_4$ and that β_1 only appears in the expression $\beta_1\beta_3\beta_4$. Let $\gamma_4 = \beta_3\beta_4$ and $\gamma_1 = \beta_1\gamma_4$, our problem is equivalent to showing that the minimum for γ_1 , β_3 and γ_4 such that $0 \leq \gamma_1 \leq \gamma_4 \leq \beta_3 \leq 1$ of the maximum between $\beta_3 + \min(0, 1 - 2\beta_3 + \gamma_4)$, $\beta_3 - \gamma_4 + \gamma_1$ and $1 - \beta_3 + \gamma_1$ is greater than one half.

Consider first the case when $1 - 2\beta_3 + \gamma_4 \geq 0$. Our problem is to show that $\max\{\beta_3, \beta_3 - \gamma_4 + \gamma_1, 1 - \beta_3 + \gamma_1\} \geq 1/2$ whenever $1 - 2\beta_3 + \gamma_4 \geq 0$ and $0 \leq \gamma_1 \leq \gamma_4 \leq \beta_3 \leq 1$. Given that $\gamma_1 \leq \gamma_4$ then $\beta_3 \geq \beta_3 - \gamma_4 + \gamma_1$. First, if β_3 is the maximum of those three terms then $\beta_3 \geq 1 - \beta_3 + \gamma_1$, so that $2\beta_3 \geq 1 + \gamma_1 \geq 1$, or $\beta_3 \geq 1/2$. Second, if $1 - \beta_3 + \gamma_1$ is the maximum of those three terms then $1 - \beta_3 + \gamma_1 \geq \beta_3$, so that $\beta_3 \leq (1 + \gamma_1)/2$ and therefore $1 - \beta_3 + \gamma_1 \geq 1 - (1 + \gamma_1)/2 + \gamma_1 = (1 + \gamma_1)/2 \geq 1/2$.

Consider now the case when $1 - 2\beta_3 + \gamma_4 \leq 0$. Our problem is to show that $\max\{1 - \beta_3 + \gamma_4, \beta_3 - \gamma_4 + \gamma_1, 1 - \beta_3 + \gamma_1\} \geq 1/2$ whenever $1 - 2\beta_3 + \gamma_4 \leq 0$ and $0 \leq \gamma_1 \leq \gamma_4 \leq \beta_3 \leq 1$. Given that $\gamma_1 \leq \gamma_4$ then $1 - \beta_3 + \gamma_4 \geq 1 - \beta_3 + \gamma_1$. First if $1 - \beta_3 + \gamma_4 \geq \beta_3 - \gamma_4 + \gamma_1$ then $\beta_3 \leq 1/2 + \gamma_4 - \gamma_1/2$, so that

$1 - \beta_3 + \gamma_4 \geq (1 + \gamma_1)/2 \geq 1/2$. Second if $1 - \beta_3 + \gamma_4 \leq \beta_3 - \gamma_4 + \gamma_1$ then $\beta_3 \geq 1/2 + \gamma_4 - \gamma_1/2$, so that $\gamma_1 + \beta_3 - \gamma_4 \geq (1 + \gamma_1)/2 \geq 1/2$.

Hence there is always one player amongst player 1 and player 2 who can achieve a payoff of $1/2$ in the stage game. Therefore for any stage-game profile $(\alpha_1, \alpha_2, \alpha_3, \alpha_4) \in [0, 1]^4$, both player 1 and player 2 are guaranteed a repeated-game payoff of at least

$$(1 - \tilde{\delta})\frac{1}{2} + \tilde{\delta}u^*,$$

where u^* is the minimum payoff attainable in any subgame-perfect equilibrium for players 1 and 2.¹⁴ If $(\alpha_1, \alpha_2, \alpha_3, \alpha_4)$ is part of an equilibrium that gives players 1 and 2 their lowest subgame-perfect equilibrium payoff we then have:

$$u^* \geq (1 - \tilde{\delta})\frac{1}{2} + \tilde{\delta}u^*,$$

so that

$$u^* \geq \frac{1}{2}.$$

1.3.2 Generalization

In a more recent paper, Chen and Takahashi (2012) generalize our result. They aggregate the stage-game dimensionality assumption with the different discount factor assumption in a dynamic non-equivalent utility assumption (DNEU). DNEU simply states that when players have equivalent utilities they must have different discount factors.¹⁵ Chen and Takahashi (2012) dispense with the pure minmax assumption that we make and provide a more explicit construction of the dynamic player specific punishments, whereas we rely on the compactness of the equilibrium payoff set and use this to provide bounds on the difference between the lowest equilibrium payoffs of any two players. We note however that Chen and Takahashi (2012) rely on the compactness of

¹⁴Note that because they have the same stage game payoffs and the same discount factor, players 1 and 2 must have the same lowest subgame-perfect equilibrium payoff.

¹⁵In this chapter we considered a case in which all players have equivalent utilities, which is the most problematic case for the folk theorem. DNEU therefore reduces to having all players have different discount factors in that case.

the set of feasible repeated-game payoffs and use the lowest feasible payoff for each player without explicitly constructing them.

Chapter 2

Failure of Gradualism under Imperfect Monitoring

2.1 Introduction

In many economic settings players often have incentives to free ride and benefit from the contributions of others without having to incur a private cost. As a result, many Pareto optimal outcomes cannot be sustained as equilibria of strategic interactions. As discussed in Chapter 1 this problem can usually be overcome through repeated interactions, which allow players to reward and punish each other over time. The folk theorem then tells us that any individually rational payoff profile can be obtained in equilibrium, provided players are sufficiently patient. For example, in an infinitely repeated prisoner's dilemma, cooperation can be achieved with a simple grim-trigger strategy, which prescribes a return to the stage-game Nash equilibrium if one player deviates from a given cooperative path.

While the ability that players have to punish each other over time is central to reaching cooperation, in various economic situations this ability can be limited. In this chapter we consider two potential sources of limitations: irreversibility and imperfect monitoring.

Irreversibility occurs when after some degree of cooperation, threats to

return to a non-cooperative outcome can no longer be made. For example, in dynamic contribution games, irreversibility arises as past contributions are non refundable: players cannot threaten to reduce their overall level of contributions to the public project as past contributions cannot be claimed back. Irreversibility thus reduces the ability that players have to punish each other, which may in turn limit the scope for cooperation.

Imperfect monitoring occurs when players do not perfectly observe each other's actions. As a consequence deviations by players may be more difficult to detect, reducing the ability that players have to punish each other and again reducing the scope for cooperation.

The constraints that irreversibility and imperfect monitoring impose on cooperation have been studied independently and both strands of literature show a series of positive results, where cooperation is possible in equilibrium. We discuss those results in the following paragraphs. In this chapter we study how those two limitations interact: actions are irreversible and players do not perfectly observe each other's actions. In contrast with both strands of literature, we show that cooperation can be impossible to achieve when those limitations are considered together.

The main insight from the literature on irreversibility is that cooperation has to take the form of gradual increases in contribution levels over time. Marx and Matthews (2000) study a game of dynamic voluntary contribution to a public project where past contributions are not refundable and payoffs are linear in cumulative contributions, with a possible extra benefit when cumulative contributions are above a given threshold (the "completion point"). They construct an approximately efficient subgame-perfect equilibrium when there is little discounting. Lockwood and Thomas (2002) study a similar setting, with no extra benefit, and characterize the efficient equilibria for any discount factor by the means of a second-order difference equation. Gale (2001) introduces the notion of monotone games with positive spillovers, a more general setting, and looks at the case without discounting. All those

papers show that there can be cooperation when actions are irreversible, but that it has to take the form of gradual increases in contribution levels: because of irreversibility, the only threats that can be made are the reductions of future increases in the levels of contributions. This implies that cooperation has to be gradual. In a bargaining setting Compte and Jehiel (2004) show that, when players' outside options are history dependent and players have the option to terminate the game at any stage, equilibrium concessions will exhibit gradualism. The option that players have of terminating the game has the same role as the threat of discontinuing contributions to the public project.

Admati and Perry (1991) also study voluntary contributions to a public project and show that when past contributions are sunk players contribute gradually. The setting however is different as players move sequentially and do not enjoy intermediate flow payoffs. Instead, the benefit from the project is received only once the project is completed. Compte and Jehiel (2003) however show that when players value the project differently it can then be completed in only two stages.

In the papers reviewed, players always perfectly observe either individual contributions or total contributions. In particular players can condition their actions on the level of total contributions and detect any deviation from a given contribution path, possibly triggering a punishment phase. However, it is often the case that players cannot perfectly monitor each other's actions. Continuing with the public project example, players may only be able to observe the stage of development of the public project, which can be a noisy signal of total contributions.

In a repeated game with finite actions and signals, when signals are publicly observed and sufficiently informative, Fudenberg et al. (1994) show that cooperation can still be achieved in equilibrium, provided players are sufficiently patient, and establish a folk theorem for games with imperfect public monitoring. Cooperation is again possible with a continuum of actions and

signals. Green and Porter (1984) and Porter (1983) study collusion in Cournot games with imperfect public monitoring where the action set (the quantity to produce) and the set of signals (the market price) are a continuum. They show that collusion is possible and Porter (1983) characterize the optimal collusive trigger strategy. Abreu et al. (1986) study optimal strategies in the Green and Porter (1984) model, without restricting attention to trigger strategies. However they depart from the Green and Porter (1984) model by restricting attention to a finite set of actions.

With a particular form of private monitoring (“network monitoring”), Wolitzky (2013) studies the level of cooperation that can be achieved in a repeated public good game where players perfectly observe the actions of their neighbours in a network but cannot observe the other players’ level of cooperation. As with the other papers mentioned, this paper does not consider irreversibility constraints, so that the payoff structure remains stationary over time.

The main question we address in this chapter is whether cooperation can still be achieved when there is imperfect monitoring and the environment is non-stationary due to irreversibility constraints. We show that, under certain regularity assumptions about the payoff function and the monitoring technology, cooperation can no longer be achieved and players must play the unique stage-game Nash equilibrium for ever.¹ This result is striking as it shows a stark discontinuity with the perfect monitoring case: the introduction of a little noise in the monitoring technology can render cooperation impossible.²

We consider a model in which a two-player Prisoner’s Dilemma with continuous actions is played infinitely many times. In each period players choose a level of contribution (a number in \mathbb{R}_+). While it is strictly dominant not to contribute in the stage game, it is mutually beneficial to do so. Actions

¹We show the result for pure-strategy Nash equilibrium.

²In a recent paper Bonatti and Hörner (2011) study inefficiencies that arise in teams in a dynamic moral hazard setting with incomplete information about the quality of the project. Players do not observe the actions of others in their team, but only observe whether and when the project succeeds. Under this particular monitoring structure they show that agents will under invest in effort.

are irreversible, so contribution levels cannot decrease over time.³ Crucially, players do not perfectly observe each other's actions. Instead in each period they receive a noisy signal of the action profile played. The signal is publicly observed by all players and drawn from a compact subset of \mathbb{R}^k according to a known probability distribution. When the payoff function is continuously differentiable in actions and the monitoring technology is continuous in actions, we show that with irreversibility there can no longer be cooperation.⁴

Under perfect monitoring, with irreversibility, cooperation takes place in the form of gradual increases in contribution levels. At any point in time, the threat of maintaining contribution levels constant forever provides the necessary incentives to players to contribute today. That is, the losses from the withdrawal of future increases in contribution levels offset the instantaneous gain from a deviation.

However under imperfect monitoring this is no longer the case. If there are strictly positive contributions in equilibrium, it can first be shown that for a set of histories of positive measure, contributions will be arbitrarily close to an upper bound.⁵ Close to this upper bound, a player will have an incentive to deviate by slightly reducing his contribution today and then resuming to the prescribed strategy tomorrow. The gain from this deviation is instantaneous and of a similar order of magnitude as the deviation. The cost is two-folds: first, the deviation affects the distribution of signals. This effect is also of similar order as the deviation.⁶ Secondly, given that signals are affected, the deviating player receives a lower continuation value. This loss is however arbitrarily small as contributions are arbitrarily close to an upper bound. Under

³One interpretation of irreversibility is that the payoff-relevant variable is the stock of total contributions, which is irreversible as players can only contribute non-negative amounts in each period. See Lockwood and Thomas (2002, Section 4) for a discussion.

⁴We also show that the result holds for the "linear kinked" case, when the benefit from another player's contribution becomes null beyond a certain level of total contributions.

⁵This is because in order to provide incentives for players to increase their level of contributions today, contributions have to increase with strictly positive probability in the future. As contributions will be bounded from above in equilibrium, they must converge to a finite limit. Moreover this limit cannot be reached, as players would have an incentive to deviate just before reaching that limit.

⁶When the monitoring technology is smooth with respect to actions.

perfect monitoring this small loss is sufficient to deter deviations, as it occurs with probability 1. Under imperfect monitoring however this loss, coupled with the fact that a small deviation will have a small impact on the distribution of signals, will not be sufficient to provide incentives for cooperation.

We can think of a number of examples of strategic situations in which cooperation is mutually desirable but myopic incentives are to defect and actions are irreversible or very costly to reverse. For example in an industry with a declining demand, competing firms might have a mutual interest in reducing their capacity, which can be considered irreversible. However in a one-shot game it is strictly dominant for firms not to reduce their capacity. Similarly, parties over-exploiting a common resource might mutually benefit from a destruction of their capital in order to reduce over-exploitation, even though in a one-shot interaction it is dominant for each party not to destroy capital. Other examples include environmental cooperation, where the installation of costly abatement technology is irreversible, and disarmament between warring parties.⁷ In all these examples, it is possible that players might not perfectly observe each other's actions but only noisy signals of those actions.

The rest of this chapter is organized as follows: in Section 2.2 we describe the model under the imperfect public monitoring framework; the main results of the model with perfect monitoring are summarized in Section 2.3; in Section 2.4 we characterize the unique public Nash equilibrium, in which players (almost) never contribute; Section 2.5 presents a counterexample where cooperation is possible, but where the monitoring technology is not continuous with respect to the players' actions; and Section 2.6 concludes.

2.2 The model under imperfect public monitoring

In this section we present the main features of the model under the assumption of imperfect public monitoring. We follow the model of Lockwood and

⁷Those examples are taken from Marx and Matthews (2000) and Lockwood and Thomas (2002).

Thomas (2002) and add the assumption that actions are not perfectly monitored. We start with the properties of the stage game, which has the structure of a Prisoner's Dilemma with continuous actions: Players choose a level of contribution (for example to a public project) in \mathbb{R}_+ . While it is strictly dominant for players not to contribute in the stage game, players can both benefit from strictly positive levels of contributions. We then present the dynamic version of the model characterized by two main assumptions: actions are irreversible, so that the level of contributions has to be non-decreasing; and players do not perfectly observe each other's actions. Instead they observe a public noisy signal drawn from a known probability distribution on a compact subset of \mathbb{R}^k . We then describe the histories upon which players condition their actions and how they evaluate future streams of random payoffs.

2.2.1 The stage game

There are two players $i = 1, 2$.⁸ Each player i chooses an action $c_i \in \mathbb{R}_+$, interpreted as his level of contribution to a public project. Both players simultaneously choose an action and the payoff to player i from the action profile $(c_1, c_2) \in \mathbb{R}_+^2$ is $\pi(c_i, c_j)$. It is assumed that π is continuously differentiable, decreasing in its first argument and increasing in its second argument. There exist contribution levels $c_1 > 0$ and $c_2 > 0$ such that it is mutually desirable for both players to reach those levels, providing the game with a Prisoner's Dilemma structure. Furthermore, the function $\pi(c_1, c_2) + \pi(c_2, c_1)$ is assumed to have a unique global maximizer on \mathbb{R}_+^2 . Finally the marginal cost of contributing is restricted to be bounded away from zero:

Assumption 2.1 (Smoothness). *The function $(c_1, c_2) \mapsto \pi(c_1, c_2)$ is continuously differentiable.*⁹

⁸The main result of this paper can be generalized to the case of n players in a straightforward way.

⁹Note that this assumption excludes the linear kinked case considered by Lockwood and Thomas (2002). The result still holds in that case, and this is discussed in Remark 2.7.

Assumption 2.2 (Prisoner's dilemma structure). *The function π is decreasing in its first argument and increasing in its second argument: $\pi_1 \leq 0$ and $\pi_2 \geq 0$.¹⁰ Moreover there exist $c_1 > 0$ and $c_2 > 0$ such that $\pi(c_1, c_2) > \pi(0, 0)$ and $\pi(c_2, c_1) > \pi(0, 0)$.*

Assumption 2.3 (Global maximizer). *The function $(c_1, c_2) \mapsto \pi(c_1, c_2) + \pi(c_2, c_1)$ has a unique global maximizer (c_1^*, c_2^*) on \mathbb{R}_+^2 , such that $\pi(c_1, c_2) + \pi(c_2, c_1)$ is decreasing in $c_1 + c_2$ for $c_1 + c_2 \geq c_1^* + c_2^*$.*

Remark 2.1. Note that because $(c_1, c_2) \mapsto \pi(c_1, c_2) + \pi(c_2, c_1)$ is symmetric, if it has a unique maximiser then this maximiser is such that $c_1^* = c_2^* = c^*$. \diamond

Remark 2.2. The second part of Assumption 2.3 implies that the function $(c_1, c_2) \mapsto \pi(c_1, c_2) + \pi(c_2, c_1)$ is quasiconcave on the set $\{(c_1, c_2) \in \mathbb{R}_+^2 : c_1 + c_2 \geq c_1^* + c_2^*\}$. This assumption rules out the possibility that beyond the optimum, more contributions could be beneficial. In the example of environmental cooperation, it could be the case that even though players have over-invested, further contributions will produce a technological breakthrough for which the incremental benefit will outweigh the incremental cost.

This assumption can be dispensed with when focusing on symmetric equilibrium. \diamond

Assumption 2.4 (Strictly positive marginal cost). $\pi_1(x, y) < 0$ for $x > 0$.

Example 2.1. The benefit function $\pi(c_1, c_2) = -\frac{c_1^2}{2} + c_2$ satisfies assumptions 2.1 to 2.4. \diamond

2.2.2 The dynamic game and monitoring structure

The stage game is played infinitely many times. In each period $t = 0, 1, \dots$ both players simultaneously choose an action $c_i^t \in \mathbb{R}_+$ that cannot be lower than the action chosen in the previous period. One interpretation is that the payoff-relevant variable is the sum of all contributions and that increments in

¹⁰We use subscripts to denote the partial derivatives of π : $\pi_1(x, y) = \partial\pi(x, y)/\partial x$ and $\pi_2(x, y) = \partial\pi(x, y)/\partial y$.

contributions have to be non-negative, so that the total level of contribution of each player will be non-decreasing:

Assumption 2.5 (Irreversibility). $c_i^t \geq c_i^{t-1}$, $i = 1, 2$, $t \geq 1$.

There is a common discount factor $\delta \in (0, 1)$ and the players evaluate payoff streams using discounting. The sequence of action profiles $(c_i^t, c_j^t)_{t=0}^\infty$ generates a payoff for player i of:

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi(c_i^t, c_j^t).$$

We consider a game with imperfect public monitoring: at the end of each period, players only observe a public signal y drawn from a compact set $Y \subset \mathbb{R}^k$, $k \geq 1$. If $k = 1$ for example, the signal observed could be the sum of the players' contributions plus some random noise. If $k = 2$, the public signal could have a component for each of the players' actions. We denote by $\mathbb{P}_Y(\cdot \mid c_1, c_2)$ the probability measure on Y induced by the public signal, conditional on contribution levels c_1 and c_2 , so that for any measurable $E \subset Y$ we have:

$$\mathbb{P}(y \in E \mid c_1, c_2) = \int_E \mathbb{P}_Y(dy \mid c_1, c_2).$$

If $\mathbb{P}_Y(dy \mid c_1, c_2)$ can be written as $f(y \mid c_1, c_2)dy$ then the monitoring technology is absolutely continuous with respect to the Lebesgue measure with probability density function $f(y \mid c_1, c_2)$, but this is not necessarily the case.

We assume that the probability measure \mathbb{P}_Y is continuous with respect to contribution levels, so that the distribution of signals cannot change too much if changes in contribution levels are small:

Assumption 2.6 (Feller continuity). *There exists a constant K such that $|\mathbb{P}(E \mid c_1 + \Delta, c_2) - \mathbb{P}(E \mid c_1, c_2)| \leq K\Delta$ and $|\mathbb{P}(E \mid c_1, c_2) - \mathbb{P}(E \mid c_1, c_2 + \Delta)| \leq K\Delta$ for any measurable set $E \subset Y$.¹¹*

¹¹For example, if the monitoring technology is continuous with respect to the Lebesgue measure, we could impose the following restriction on its probability density function f : there exists a constant K such that $|f(y \mid c_1 + \Delta, c_2) - f(y \mid c_1, c_2)| \leq K\Delta$ and $|f(y \mid c_1, c_2) - f(y \mid c_1, c_2 + \Delta)| \leq K\Delta$ for any $y \in Y$.

Example 2.2. The public signal $y = \min\{c_1 + c_2 + \epsilon, M\}$,¹² where ϵ is a uniform random variable on the interval $[a, b]$ ($a < b \in \mathbb{R}$) and $M > 0$ would satisfy assumption 2.6. \diamond

The public signal y is the only information each player has about his opponent's play. Therefore if player i receives payoffs at the end of each period they cannot convey additional information about the other player's action: the realized payoff π^* is a function of one's own contribution and the public signal. The ex ante payoff function π is then the expectation of the ex post payoff:

$$\pi(c_i, c_j) = \int_Y \pi^*(c_i, y) \mathbb{P}_Y(dy | c_1, c_2), \quad \forall c_i, c_j. \quad (2.1)$$

Remark 2.3. Note that we do not provide assumptions on π^* and \mathbb{P}_Y such that π will satisfy assumptions 2.1 to 2.4, and it is not clear that for every π and every \mathbb{P}_Y we can find a π^* that satisfy (2.1). The advantage of working with the ex ante payoff rather than the ex post payoff is that it allows us to remain agnostic about the link between actions and signals - besides the continuity assumption 2.6. In particular we do not have to make likelihood ratio assumptions, such that higher actions are statistically associated with higher signals. As our result is a negative one it makes it more powerful.

Here is however a simple example that would satisfy our assumptions: $\pi(c_i, c_j) = -c_i^2/2 + c_j$, $y = (y_i, y_j)$ such that $\mathbb{E}_Y(y | c_i, c_j) = (c_i, c_j)$ and $\pi^*(c_i, y) = -c_i^2/2 + y_j$.¹³

A possible alternative that would not require to define ex post payoffs π^* would be to interpret the discount factor δ as the probability with which the interaction will terminate. Players then only receive their payoff once the interaction has ended. \diamond

Remark 2.4. Even though we consider actions in each period as being the

¹²The truncation of the public signal if $c_1 + c_2 + \epsilon \geq M$ is made in order to respect the compactness assumption of the support of y . This is not restrictive as M could be arbitrarily large, and as we show that in equilibrium players' contribution levels must be bounded from above (Lemma 2.3).

¹³ $\mathbb{E}_Y(y | c_i, c_j) = (c_i, c_j)$ would occur on a finite range to accommodate with the compact support assumption.

level of total contributions, another interpretation is that players choose a non-negative increment to their level of contribution. Marx and Matthews (2000) formulate their model in such a way, and this is also discussed in Lockwood and Thomas (2002, Section 4). In this setting, it may be natural to consider a public signal that does not depend on the total stock of contributions but on the increments in contributions. We could then interpret benefits as being a sum of incremental benefits, each increment depending on one's incremental contribution and the public signal. \diamond

2.2.3 Private and public histories

In each period, players only observe the public signal and the action they have played. A private history for player i is a sequence of actions and signals $h_i^t = (c_i^0, y^0; c_i^1, y^1; \dots; c_i^{t-1}, y^{t-1})$, and the set of all private histories for player i is $\mathcal{H}_i := \cup_{t \geq 0} (\mathbb{R}_+ \times Y)^t$, where $(\mathbb{R}_+ \times Y)^0 = \emptyset$. A public history h^t is a sequence of t public signals: $h^t = (y^0, y^1, \dots, y^{t-1}) \in Y^t$. The set of all public histories is $\mathcal{H} := \cup_{t \geq 0} Y^t$. We will also use the notations \mathcal{H}^t and \mathcal{H}_i^t to denote the set of public and private histories of length t respectively.

A pure strategy σ_i for player i is a measurable function that specifies a level of contribution $\sigma_i(h_i^t)$ after any $h_i^t \in \mathcal{H}_i$ and that satisfies irreversibility:

$$\sigma_i: \begin{cases} \mathcal{H}_i & \longrightarrow \mathbb{R}_+ \\ h_i^t & \longmapsto \sigma_i(h_i^t) \end{cases},$$

such that for any $h_i^t \in \mathcal{H}_i$ and any $y \in Y$ we have $\sigma_i(h_i^t; c_i^{t-1}, y) \geq c_i^{t-1}$, where $h_i^t; c_i^{t-1}, y$ denotes the concatenation of histories h_i^t and c_i^{t-1}, y . For any strategy σ_i , player i 's continuation strategy induced by $h_i^t \in \mathcal{H}_i^t$ is denoted by $\sigma_i|_{h_i^t}$ such that $\sigma_i|_{h_i^t}(h) := \sigma_i(h_i^t h)$, $\forall h \in \mathcal{H}_i$.

Remark 2.5 (Pure vs. mixed strategies). In this chapter we focus on pure strategies. An important characteristic of pure strategies is that on the equi-

librium path, they are realization equivalent to a public strategy.¹⁴ (See for example Mailath and Samuelson 2006, Lemma 7.1.2.) This plays an important role in the proof of Theorem 2.1, which relies on the fact that players know each other's contribution levels on the equilibrium path. This property will no longer hold when considering mixed strategies if players rely on their private information, and the proof Theorem 2.1 cannot be extended to this case.

As on the equilibrium path pure strategies depend only on the public history, all statements apply to public histories, as the aim is to characterize the set of pure-strategy Nash equilibria. \diamond

2.2.4 Stochastic process of public signals

Let $\Omega := Y^{\mathbb{N}}$ be the space of infinite sequences of public signals. Along with a monitoring technology, a pure strategy profile $\sigma = (\sigma_1, \sigma_2)$ determines recursively a stochastic process of public signals, which induces a probability distribution on Ω that we denote by \mathbb{P}_σ . Expectations with respect to that probability distribution will be denoted by \mathbb{E}_σ .

An element $\omega \in \Omega$ is an infinite sequence of public signals and we denote by $h^t(\omega)$ the first t elements of ω . Let $V(\sigma_i, \sigma_j)$ be the expected payoff of player i from the strategy profile $\sigma = (\sigma_1, \sigma_2)$ and $V(\sigma_i, \sigma_j | h^\tau)$ be the continuation payoff from σ after the public history h^τ .¹⁵

$$V(\sigma_i, \sigma_j) := (1 - \delta) \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \delta^t \pi(\sigma_i(h^t(\omega)), \sigma_j(h^t(\omega))) \right],$$

$$V(\sigma_i, \sigma_j | h^\tau) := (1 - \delta) \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \delta^t \pi(\sigma_i(h^{t+\tau}(\omega)), \sigma_j(h^{t+\tau}(\omega))) | h^\tau \right].$$

A profile of pure strategies (σ_1, σ_2) is a Nash equilibrium if for any $i \in \{1, 2\}$ and any strategy σ' we have that $V(\sigma_i, \sigma_j) \geq V(\sigma', \sigma_j)$, provided that

¹⁴Off path, the irreversibility constraint will imply that players may have to use their private information.

¹⁵Even though any finite history occurs with probability zero because there is a continuum of signals, the probability conditional on a history h^τ is well defined as it is the probability measure induced by the continuation strategy profile $\sigma |_{h^\tau}$.

$V(\sigma_i, \sigma_j)$ is well defined.¹⁶

2.2.5 Notation

For ease of readability we will use the following notations in the rest of the paper:

$$\pi(\sigma_i, \sigma_j | h^t) := \pi(\sigma_i(h^t), \sigma_j(h^t)),$$

and

$$\mathbb{P}_Y(dy | \sigma_i, \sigma_j, h^t) := \mathbb{P}_Y(dy | \sigma_i(h^t), \sigma_j(h^t)).$$

We also denote by $h^t|_\tau$, $1 \leq \tau \leq t-1$, the τ -truncation of a public history h^t : If $h^t = (y^0, y^1, \dots, y^{t-1})$ then $h^t|_\tau = (y^0, y^1, \dots, y^{\tau-1})$.

2.3 Main Results under Perfect Monitoring

In this section we briefly summarize the main result from the model with perfect monitoring, as in Lockwood and Thomas (2002). They first show that in any equilibrium with positive contributions contributions never reach their limit. They then show that in a symmetric equilibrium, any efficient path of contributions must solve the following difference equation:

$$\pi(c_t, c_{t+1}) = \frac{1}{\delta} [\pi(c_{t-1}, c_t) - \pi(c_t, c_t)] + \pi(c_t, c_t), \quad t > 0, \quad (2.2)$$

They then show that for any given discount factor, contribution levels in any equilibrium will be bounded from above by a bound strictly lower than the first-best level of contributions. However this inefficiency vanishes as the discount factor tends to one.

¹⁶As we have not made the assumption that π is bounded, it could be the case that V is not finite.

2.4 Failure of Cooperation in Equilibrium

In this section we show that the only pure-strategy Nash equilibrium of the dynamic game presented in Section 2.2 is to play $(0,0)$ (no contribution) after almost every history: cooperation cannot be achieved with positive probability. In Section 2.5, Assumption 2.6 is relaxed and it is shown that cooperation is again possible in equilibrium.

Theorem 2.1. *Let $\sigma = (\sigma_1, \sigma_2)$ be a pure-strategy Nash equilibrium of the dynamic contribution game with imperfect public monitoring. Then $\sigma_i(h^t) = 0$, $i = 1, 2$, for almost every public history $h^t \in \mathcal{H}$, under Assumptions 2.1 to 2.6.*

Remark 2.6 (Nash equilibrium vs. subgame perfect equilibrium). Statements are made almost surely with respect to the probability measure \mathbb{P}_σ . Hence the analysis that follows occurs on the equilibrium path. \diamond

In Section 2.4.1 we provide a brief discussion of the proof. Section 2.4.2 introduces some preliminary results, while the formal proof is presented in Section 2.4.3.

2.4.1 Outline of the Proof

The main idea behind the proof of Theorem 2.1 is to consider, in each period, the essential supremum of a player's level of contribution. Contributions can only take values above the essential supremum on set of measure zero and there always exists a set of histories of positive measure for which contributions are arbitrarily close to the essential supremum.

It is first shown that at any point in time, there is a set of histories of positive measure for which contribution levels are close to their essential supremum (for one of the players) and for which contributions are expected to increase in the future. Because of irreversibility, the only way to provide incentives for players to contribute today is through the promise of future increases in contributions. No longer increasing contributions after histories for which

contributions are close to the essential supremum would mean that along such histories there will be a last time where contributions increase, giving players an opportunity to profitably deviate. This result is analogous to Lockwood and Thomas (2002, Lemma 2.1 (ii)), who show that in an equilibrium with positive contributions, contributions cannot become constant after a certain time.

As the essential supremum of contributions converges (it is a sequence of increasing numbers bounded from above in equilibrium), we can consider a time after which it is arbitrarily close to its limit. The previous intuition now tells us that along histories for which contributions are close to the limit of the essential supremum, players are expected to increase contributions in the future. However as contributions are close to their upper bound, they cannot be expected to increase by a significant amount. To prove Theorem 2.1 we, then, show that along such histories a player can profitably deviate by increasing his contribution levels by less than what is prescribed in a putative equilibrium with positive contributions.

The gain from such a deviation consists of the instantaneous gain from reducing contribution levels. The cost is two-folds: first, a deviation will affect the distribution of signals; secondly, given that signals are affected, the deviating player will receive different (lower) continuation values.¹⁷ However, when the monitoring technology is continuous with respect to players' actions, a small deviation will have an impact on the distribution of public signals of a similar order to the gain.¹⁸ Furthermore, as contributions of the other player are close to an upper bound, losses from lower future continuation values will be arbitrarily small. Hence the cost of deviating consists of one element that is of similar order to the gain and another element that is arbitrarily small, making the deviation profitable.

¹⁷Note that the use of the word "lower" might suggest that some assumption has been made about the informativeness of the signals (such as a monotone likelihood ratio assumption). The proof does not rely on any such assumption. However if the player expects higher continuation values after deviating then the proof becomes trivial.

¹⁸With perfect monitoring, the impact of a small deviation is large, as this small deviation is detected with probability 1.

2.4.2 Preliminary Results

In this section we present some general properties of the possible equilibria of the game described in Section 2.2.

First we show that after almost every public history where contributions have increased, contributions are again expected to increase in the future. That is, the only force that provides incentives for players to increase their contributions today is the expectation of future increases in contributions from the other player.

Lemma 2.1. *In equilibrium, for almost every history $h^t \in \mathcal{H}^t$ such that $\sigma_i(h^t) > \sigma_i(h^t|_{t-1})$, $i \in \{1, 2\}$, both players are expected to increase their levels of contributions in the future:*

$$\mathbb{P}_\sigma \left(\{h \in \mathcal{H} : \sigma_j(h^t h) > \sigma_j(h^t)\} \mid h^t \right) > 0, \quad j \in \{1, 2\}.$$
¹⁹

Proof. Assume first that there is a set of histories of positive measure and length t such that $\sigma_i(h^t) > \sigma_i(h^t|_{t-1})$ but such that $\mathbb{P}_\sigma(\{h \in \mathcal{H} : \sigma_j(h^t h) > \sigma_j(h^t)\} \mid h^t) = 0$.²⁰ Even though player i has increased his contribution, he does not expect player j to do so in the future. As contributing is strictly dominated, player i could profitably deviate by not increasing his contribution.

If player i increases his contribution, he then expects player j to also do so in the future. But if player j increases his contributions it is because he similarly expects player i to increase his contributions in the future. Hence we also have $\mathbb{P}_\sigma(\{h \in \mathcal{H} : \sigma_i(h^t h) > \sigma_i(h^t)\} \mid h^t) > 0$. \square

The next lemma shows that in equilibrium the value function of players is bounded from below by their current flow payoff. If this was not the case a player could always choose to maintain his contribution level constant. As

¹⁹Recall from footnote 8 that we can condition on finite histories even though any finite history has measure zero, as a continuation strategy induces a probability distribution on the set of continuation histories.

²⁰Note that if a certain property holds for a set of histories of positive measure, then from the σ -additivity of \mathbb{P}_σ and because time is countable, there will be a certain t and a set of histories of positive measure and length t such that this property holds for those histories.

payoffs are non-decreasing in the other player's contribution, which cannot decrease due to irreversibility, this would guarantee a future payoff of at least the current flow payoff.

Lemma 2.2. *If (σ_1, σ_2) is an equilibrium, then for almost every $y \in Y$ and almost every $h^t \in \mathcal{H}$ we have that $V(\sigma_i, \sigma_j | h^t y) \geq \pi(\sigma_i, \sigma_j | h^t)$. Furthermore $V(\sigma_i, \sigma_j | h^t) \geq \pi(\sigma_i, \sigma_j | h^t)$.*

Proof. Assume that there is an equilibrium where $V(\sigma_i, \sigma_j | h^t y) < \pi(\sigma_i, \sigma_j | h^t)$ for a set of histories of positive measure and a given t . Consider the strategy σ' for player i that coincides with σ_i except after a history h^t where $V(\sigma_i, \sigma_j | h^t y) < \pi(\sigma_i, \sigma_j | h^t)$, in which case player i stops contributing forever: $\sigma'(h^t h) = \sigma_i(h^t)$. Then for any $h \in \mathcal{H}$, $\pi(\sigma', \sigma_j | h^t h) \geq \pi(\sigma_i, \sigma_j | h^t)$. Hence $V(\sigma_i, \sigma_j | h^t y) \geq \pi(\sigma_i, \sigma_j | h^t)$, a contradiction.

Moreover,

$$\begin{aligned} V(\sigma_i, \sigma_j | h^t) &= (1 - \delta)\pi(\sigma_i, \sigma_j | h^t) + \delta \mathbb{E}_\sigma[V(\sigma_i, \sigma_j | h^t y) | h^t] \\ &\geq (1 - \delta)\pi(\sigma_i, \sigma_j | h^t) + \delta \pi(\sigma_i, \sigma_j | h^t) \\ &= \pi(\sigma_i, \sigma_j | h^t), \end{aligned}$$

which completes the proof. \square

In the next lemma we show that in equilibrium players' contributions are bounded from above.

Lemma 2.3. *In equilibrium, for almost every history $h \in \mathcal{H}$ we have $\sigma_1(h) + \sigma_2(h) < 2c^*$.*

Proof. Assume that this is not the case and there is a set of histories of positive measure and length t such that $\sigma_1(h^t) + \sigma_2(h^t) \geq 2c^*$. The sum of the players' value functions, given h^t , is a discounted sum of terms $\pi(\sigma_1, \sigma_2 | h^t h) + \pi(\sigma_2, \sigma_1 | h^t h)$, $h \in \mathcal{H}$. We consider in turn two cases.

First assume that players increase their contributions with positive probability after such histories. As $\pi(c_1, c_2) + \pi(c_2, c_1)$ is above its global max-

imum (Assumption 2.3) it is decreasing in $c_1 + c_2$. Therefore $\pi(\sigma_1, \sigma_2 \mid h^t h) + \pi(\sigma_2, \sigma_1 \mid h^t h) \leq \pi(\sigma_1, \sigma_2 \mid h^t) + \pi(\sigma_2, \sigma_1 \mid h^t) \forall h \in \mathcal{H}$, where the inequality is strict on a set of histories of positive measure. This implies that $V(\sigma_i, \sigma_j \mid h^t) < \pi(\sigma_i, \sigma_j \mid h^t)$ for at least one $i \in \{1, 2\}$, a contradiction with Lemma 2.2.

Secondly if instead contributions remain constant then for each h^t such that $\sigma_1(h^t) + \sigma_2(h^t) \geq 2c^*$ there is a time s , $0 \leq s < t$, at which one of the players makes a last increase. This is a contradiction with Lemma 2.1, which states that if one player increases his contribution, contributions are expected to increase again in the future. \square

Lemma 2.3 implies that in equilibrium each player will not contribute above $2c^*$. We define this upper bound as $\bar{c} := 2c^*$.

2.4.3 Proof of Theorem 2.1

Assume that there is a pure-strategy equilibrium $\sigma = (\sigma_1, \sigma_2)$ where players make strictly positive contributions. Without loss of generality, in equilibrium, at least one of the players contribute in period 0.²¹ Suppose $\sigma_1(\emptyset) > 0$.

In what follows we consider the *essential supremum* of contribution levels, which is the supremum except on a set of measure zero and is defined as follows:²²

$$\text{ess sup}_{h^t \in \mathcal{H}^t} \sigma_i(h^t) := \inf \left\{ a \in \mathbb{R} : \mathbb{P}_\sigma(\{h^t \in \mathcal{H}^t : \sigma_i(h^t) > a\}) = 0 \right\}.$$

Let $(\hat{\sigma}_{it})_{t=0}^\infty$, $i = 1, 2$, be the deterministic sequence of a player's essential supremum contribution:

$$\hat{\sigma}_{it} = \text{ess sup}_{h^t \in \mathcal{H}^t} \sigma_i(h^t), \quad \forall t \geq 0, i = 1, 2. \quad (2.3)$$

By definition there always exists a set of histories of positive measure for which

²¹Consider a history h^t at which the first increase in contribution levels occurs. Then $(\sigma_1|_{h^t}, \sigma_2|_{h^t})$ is also an equilibrium of the original dynamic game.

²²For a brief discussion of the essential supremum, see for example Doob (1994, Part V.17).

contributions are close to the essential supremum. We denote by $\mathcal{H}_{i\epsilon}^t$ the set of histories of length t for which contributions of player i are within ϵ of the essential supremum:

$$\mathcal{H}_{i\epsilon}^t := \{h^t \in \mathcal{H}^t : |\hat{\sigma}_{it} - \sigma_i(h^t)| \leq \epsilon\} \quad \forall t, i = 1, 2. \quad (2.4)$$

Note that $\mathbb{P}_\sigma(\mathcal{H}_{i\epsilon}^t) > 0$ and that the set of histories for which contributions are strictly higher than the essential supremum is of measure zero.

Lemma 2.4. *In equilibrium, the deterministic sequence $(\hat{\sigma}_{it})_t$ converges to a limit $\hat{\sigma}_{i\infty} < \infty$, $i = 1, 2$.*

Proof. The sequence $(\hat{\sigma}_{it})_t$ is weakly increasing, because of irreversibility, and bounded from above by \bar{c} . It therefore converges. \square

In the next lemma we show that for any time t there is a set of histories of positive measure for which contributions are close to the essential supremum *and* will continue to increase in the future.

Lemma 2.5. *In an equilibrium where players contribute with positive probability, for any $i \in \{1, 2\}$, $\epsilon > 0$ and t , there is a set of histories $\tilde{\mathcal{H}}_{i\epsilon}^t$ of length t and of positive measure such that:*

- (i) *Contributions of player i are close to their essential supremum: $|\sigma_i(h^t) - \hat{\sigma}_{it}| \leq \epsilon$, $\forall h^t \in \tilde{\mathcal{H}}_{i\epsilon}^t$; and*
- (ii) *Player i increases his contribution with strictly positive probability in the future:*

$$\mathbb{P}_\sigma \left(\{h \in \mathcal{H} : \sigma_i(h^t h) > \sigma_i(h^t)\} \mid h^t \right) > 0, \quad \forall h^t \in \tilde{\mathcal{H}}_{i\epsilon}^t,$$

where $\hat{\sigma}_{it}$ is the essential supremum of player i 's contribution in period t , as defined in (2.3).

Proof. Let (σ_1, σ_2) be a pure-strategy Nash equilibrium with positive contributions. Assume that the result does not hold and consider the set $\mathcal{H}_{i\epsilon}^t$ of

histories of length t for which player i 's contribution is within ϵ of his essential supremum contribution, as defined in (2.4): $|\sigma_i(h^t) - \hat{\sigma}_{it}| \leq \epsilon, \forall h^t \in \mathcal{H}_{i\epsilon}^t$. Recall that $\mathbb{P}_\sigma(\mathcal{H}_{i\epsilon}^t) > 0$. If Lemma 2.5 does not hold then there exists a time t and an $\epsilon > 0$ such that $\mathbb{P}_\sigma(\{h \in \mathcal{H} : \sigma_i(h^t h) > \sigma_i(h^t)\} | h^t) = 0$, for almost every history in $\mathcal{H}_{i\epsilon}^t$. We now consider in turn two cases.

First assume that for almost every $h^t \in \mathcal{H}_{i\epsilon}^t$ we have $\sigma_i(h^t) = \sigma_i(\emptyset) > 0$ (recall that we assumed players contributed with positive probability in equilibrium). From irreversibility, for every $h^t \in \mathcal{H}^t$, $\sigma(h^t) \geq \sigma(\emptyset)$. Therefore $\forall h^t \in \mathcal{H}^t$ we have $|\sigma_i(h^t) - \hat{\sigma}_{it}| \leq \epsilon$, so that $\mathcal{H}^t \subseteq \mathcal{H}_{i\epsilon}^t$. As $\mathcal{H}_{i\epsilon}^t \subseteq \mathcal{H}^t$ we have $\mathcal{H}_{i\epsilon}^t = \mathcal{H}^t$: Player i keeps his contribution level constant at $\sigma_i(\emptyset)$ with probability one. As player i never increases his contribution, player j will best respond by also maintaining his contribution constant throughout the game. But then player i has an incentive to deviate and keep his contribution constant at zero in the first period, a contradiction with the fact that (σ_1, σ_2) is an equilibrium.

Assume now that there is a positive measure subset of $\mathcal{H}_{i\epsilon}^t$ such that for histories in that subset $\sigma_i(h^t) > \sigma_i(\emptyset)$. Then for each of such history there is a time $s < t$ such that $\sigma_i(h^t|_s) < \sigma_i(h^t|_{s+1}) = \dots = \sigma_i(h^t|_{t-1}) = \sigma_i(h^t)$. Player i makes his last increase along history h^t at time s , but would then have an incentive to deviate and not perform that last increase, again a contradiction with the fact that (σ_1, σ_2) is an equilibrium. \square

The following corollary ensues:

Corollary 2.1. *In any equilibrium where players contribute with positive probability, for any $i \in \{1, 2\}$ and any $\epsilon > 0$, there exists a finite t and a set of histories $\mathcal{H}_{i\epsilon}^{t*}$ of positive measure such that:*

- (i) $|\sigma_i(h^t) - \hat{\sigma}_{i\infty}| \leq \epsilon, \forall h^t \in \mathcal{H}_{i\epsilon}^{t*}$; and
- (ii) Player $j \neq i$ increases his contribution at h^t : $\sigma_j(h^t) > \sigma_j(h^t|_{t-1}), \forall h^t \in \mathcal{H}_{i\epsilon}^{t*}$.

Proof. Consider a time t' such that $|\hat{\sigma}_{it'} - \hat{\sigma}_{i\infty}| \leq \epsilon/2$ and the set $\tilde{\mathcal{H}}_{i\epsilon/2}^{t'}$ introduced in Lemma 2.5. We know that this set has positive measure, and that for any history in that set $|\sigma_i(h^{t'}) - \hat{\sigma}_{it'}| \leq \epsilon/2$. As t' is such that $|\hat{\sigma}_{it'} - \hat{\sigma}_{i\infty}| \leq \epsilon/2$, by the triangle inequality we have that $|\sigma_i(h^{t'}) - \hat{\sigma}_{i\infty}| \leq \epsilon$, $\forall h^{t'} \in \tilde{\mathcal{H}}_{i\epsilon/2}^{t'}$. Lemma 2.5 tells us that player i increases his contribution with positive probability after histories in $\tilde{\mathcal{H}}_{i\epsilon/2}^{t'}$. By Lemma 2.1 this also implies that player j will increase his contribution with positive probability at a time $t \geq t'$. \square

We now complete the proof of Theorem 2.1 by showing that it is profitable to deviate for (say) player 1 after a history in $\mathcal{H}_{2\epsilon}^{t*}$, where player 1's strategy specifies an increase in contribution levels while player 2's contribution level is ϵ -close to its upper bound:

$$\mathcal{H}_{2\epsilon}^{t*} := \{h^t \in \mathcal{H}^t : |\sigma_2(h^t) - \hat{\sigma}_{2\infty}| \leq \epsilon \text{ and } \sigma_1(h^t) > \sigma_1(h^t|_{t-1})\}.$$

To do so we consider a deviation σ' from σ_1 which prescribes lower increases in contribution levels after histories in $\mathcal{H}_{2\epsilon}^{t*}$ but agrees with σ_1 otherwise:²³

$$\sigma'(h^t) = \begin{cases} \sigma_1(h^t) - \nu(h^t) & \text{if } h^t \in \mathcal{H}_{2\epsilon}^{t*}, \\ \sigma_1(h^t) & \text{otherwise.} \end{cases}$$

To show that $\sigma = (\sigma_1, \sigma_2)$ cannot be an equilibrium, we will show that player 1 can profitably deviate to σ' after any history in $\mathcal{H}_{2\epsilon}^{t*}$:

$$V(\sigma', \sigma_2 | h^t) - V(\sigma_1, \sigma_2 | h^t) > 0, \quad \forall h^t \in \mathcal{H}_{2\epsilon}^{t*}, \quad (2.5)$$

where

$$V(\sigma, \sigma_2 | h^t) = (1 - \delta)\pi(\sigma, \sigma_2 | h^t) + \delta \int_Y V(\sigma_1, \sigma_2 | h^t y) \mathbb{P}_Y(dy | \sigma, \sigma_2, h^t),$$

²³For example we could take $\nu(h^t) = \frac{\sigma_1(h^t) - \sigma_1(h^t|_{t-1})}{2} > 0$.

$\sigma \in \{\sigma_1, \sigma'\}$.

Note that for any $y_0 \in Y$ we have that

$$\begin{aligned} & \int_Y V(\sigma_1, \sigma_2 | h^t y) [\mathbb{P}_Y(dy | \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy | \sigma', \sigma_2, h^t)] = \\ & \int_Y [V(\sigma_1, \sigma_2 | h^t y) - V(\sigma_1, \sigma_2 | h^t y_0)] [\mathbb{P}_Y(dy | \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy | \sigma', \sigma_2, h^t)], \end{aligned}$$

as

$$\int_Y V(\sigma_1, \sigma_2 | h^t y_0) [\mathbb{P}_Y(dy | \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy | \sigma', \sigma_2, h^t)] = 0.$$

A deviation to σ' is profitable for player 1 if (2.5) holds, which can be rewritten as:

$$\begin{aligned} & \int_Y [V(\sigma_1, \sigma_2 | h^t y) - V(\sigma_1, \sigma_2 | h^t y_0)] \times \\ & \quad [\mathbb{P}_Y(dy | \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy | \sigma', \sigma_2, h^t)] < \\ & \quad \frac{1 - \delta}{\delta} [\pi(\sigma', \sigma_2 | h^t) - \pi(\sigma_1, \sigma_2 | h^t)], \quad y_0 \in Y. \quad (2.6) \end{aligned}$$

We now look for an upper bound of the left-hand side of (2.6). From Lemma 2.2 we know that the next period's value function is bounded from below by current flow payoffs. Hence we have the following inequality:

$$\begin{aligned} & V(\sigma_1, \sigma_2 | h^t y) - V(\sigma_1, \sigma_2 | h^t y_0) \leq \\ & \quad \pi(\sigma_1(h^t), \hat{\sigma}_{2\infty}) - \pi(\sigma_1(h^t), \sigma_2(h^t)), \quad \forall y \in Y. \quad (2.7) \end{aligned}$$

Using the Mean Value Theorem (recall from Assumption 2.1 that π is continuous and differentiable), there exists a $\tilde{\sigma}_2 \in (\sigma_2(h^t), \hat{\sigma}_{2\infty}) \subseteq \mathbb{R}_+$ such that $\pi(\sigma_1(h^t), \hat{\sigma}_{2\infty}) - \pi(\sigma_1(h^t), \sigma_2(h^t)) = (\hat{\sigma}_{2\infty} - \sigma_2(h^t))\pi_2(\sigma_1(h^t), \tilde{\sigma}_2)$. As $h^t \in \mathcal{H}_{2\epsilon}^{t*}$, we have that $\hat{\sigma}_{2\infty} - \sigma_2(h^t) \leq \epsilon$. Moreover there is an upper bound γ on $\pi_2(\sigma_1(h^t), \tilde{\sigma}_2)$ that is independent of h^t , so that $\pi(\sigma_1(h^t), \hat{\sigma}_{2\infty}) - \pi(\sigma_1(h^t), \sigma_2(h^t)) \leq \gamma\epsilon$. This is because π_2 is a continuous function (Assumption 2.1) that, in equilibrium, takes values on the compact set $[0, \bar{c}] \times [0, \bar{c}]$

(Lemma 2.3). Combining this with (2.7), we have the following upper bound for left-hand side of (2.6):

$$\gamma\epsilon \int_Y \left| \mathbb{P}_Y(dy \mid \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy \mid \sigma', \sigma_2, h^t) \right|. \quad (2.8)$$

From Feller continuity (Assumption 2.6), the integral in (2.8) is bounded from above by $2K(\sigma_1(h^t) - \sigma'(h^t))$.²⁴ We can now bound the left-hand side of (2.6) from above (in absolute value) by $2K\gamma\epsilon(\sigma_1(h^t) - \sigma'(h^t))$. Hence, (2.6) holds if

$$2K\gamma\epsilon(\sigma_1(h^t) - \sigma'(h^t)) < \frac{1 - \delta}{\delta} \left[\pi(\sigma', \sigma_2 \mid h^t) - \pi(\sigma_1, \sigma_2 \mid h^t) \right].$$

Again by the Mean Value Theorem there exists a $\tilde{\sigma}$ in $(\sigma'(h^t), \sigma_1(h^t)) \subseteq \mathbb{R}_+$ such that

$$\pi(\sigma', \sigma_2 \mid h^t) - \pi(\sigma_1, \sigma_2 \mid h^t) = -\pi_1(\tilde{\sigma}, \sigma_2(h^t))(\sigma_1(h^t) - \sigma'(h^t)),$$

so that (2.6) holds if

$$2K\gamma\epsilon < -\frac{1 - \delta}{\delta} \pi_1(\tilde{\sigma}, \sigma_2(h^t)). \quad (2.9)$$

The right-hand side of (2.9) is positive as π is decreasing in its first argument. Moreover it is bounded away from zero, independently of h^t . This is because π_1 is continuous and $(\tilde{\sigma}, \sigma_2(h^t))$ take values on the compact set $[\sigma_1(\emptyset), \bar{c}] \times [0, \bar{c}]$, on which it is strictly negative (Assumption 2.4), as $\sigma_1(\emptyset)$ was assumed to be strictly positive (see the beginning of Section 2.4.3).

For ϵ small enough (2.9) therefore holds, which ensures that (2.6) holds and that $\sigma = (\sigma_1, \sigma_2)$ cannot be an equilibrium, as player 1 has a profitable deviation on the set of histories in $\mathcal{H}_{2\epsilon}^{t*}$. This concludes the proof of Theo-

²⁴Let Y^+ (resp. Y^-) be the set such that $\mathbb{P}_Y(dy \mid \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy \mid \sigma', \sigma_2, h^t)$ is positive (resp. negative). Then $\int_Y \left| \mathbb{P}_Y(dy \mid \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy \mid \sigma', \sigma_2, h^t) \right| = \int_{Y^+} [\mathbb{P}_Y(dy \mid \sigma_1, \sigma_2, h^t) - \mathbb{P}_Y(dy \mid \sigma', \sigma_2, h^t)] + \int_{Y^-} [\mathbb{P}_Y(dy \mid \sigma', \sigma_2, h^t) - \mathbb{P}_Y(dy \mid \sigma_1, \sigma_2, h^t)]$. From Feller continuity, each integral is bounded from above by $K(\sigma_1(h^t) - \sigma'(h^t))$.

rem 2.1, and the only pure-strategy Nash equilibrium is when players do not cooperate and contribution levels remain constant at zero.

Remark 2.7 (The linear kinked case). The smoothness assumption (Assumption 2.1) excludes the “linear kinked” case discussed in Lockwood and Thomas (2002), where payoffs are as follows:

$$\pi(c_1, c_2) = \begin{cases} \pi_1 c_1 + \pi_2 c_2 & \text{if } c_1 + c_2 \leq 2c^*, \\ \pi_1 c_1 + \pi_2(2c^* - c_1) & \text{if } c_1 + c_2 > 2c^*, \end{cases}$$

where $\pi_1 < 0$, $\pi_2 > 0$ and $\pi_1 + \pi_2 > 0$. Under the linear kinked case most of the proof still holds, with the following simplifications: the right-hand side of (2.6) becomes $\frac{1-\delta}{\delta} [\pi(\sigma', \sigma_2 | h^t) - \pi(\sigma_1, \sigma_2 | h^t)] = \frac{1-\delta}{\delta} \pi_1 [\sigma'(h^t) - \sigma_1(h^t)]$; and the right-hand side of (2.7) becomes $\pi(\sigma_1(h^t), \hat{\sigma}_{2\infty}) - \pi(\sigma_1(h^t), \sigma_2(h^t)) = \pi_2(\hat{\sigma}_{2\infty} - \sigma_2(h^t))$.²⁵ Therefore (2.6) holds if

$$2K\pi_2\epsilon < -\frac{1-\delta}{\delta}\pi_1,$$

which is satisfied for ϵ sufficiently small, as K , π_1 , π_2 and δ are constants. \diamond

2.5 A Counterexample: All-or-Nothing Monitoring

In this section we consider an example where the monitoring does no longer have the Feller continuity property (Assumption 2.6): in each period, with probability $1 - \epsilon$, both players observe each other’s actions; with probability ϵ , players do not observe each other’s actions. Note that this monitoring structure is public. We refer to this setting as the ϵ -almost perfect monitoring

²⁵Note that in equilibrium Lemma 2.3 tells us that $\sigma_1(h^t) + \sigma_2(h^t) < 2c^*$, so that $\pi(\sigma_1(h^t), \sigma_2(h^t)) = \pi_1\sigma_1(h^t) + \pi_2\sigma_2(h^t)$.

game and will show that cooperation can again be achieved.^{26, 27}

Let us consider a possible contribution path $\mathbf{c} := (c_0, c_1, \dots)$ and the following strategy profile $\sigma_{\mathbf{c}}$: first, both players play c_0 in period 0. Then in period t both players play c_t if c_{t-1} was observed in period $t-1$. Otherwise players do not increase their levels of contributions from the previous period.

The strategy profile $\sigma_{\mathbf{c}}$ prescribes that players keep their levels of contributions constant forever after observing a deviation from the prescribed contribution path \mathbf{c} , or when not being able to observe the previous period's contributions. Given this strategy profile, we can interpret ϵ as being the probability of a breakdown in cooperation. Hence in this particular setting, the effect of imperfect monitoring is to render players less patient than in the perfect monitoring case.

Let V_t denote the value from the strategy $\sigma_{\mathbf{c}}$ at time t when players have observed the sequence of action profiles $((c_0, c_0), \dots, (c_{t-1}, c_{t-1}))$ up to time t .²⁸ After observing $((c_0, c_0), \dots, (c_{t-1}, c_{t-1}))$, players should play the action profile (c_t, c_t) in period t , yielding a current flow payoff of $\pi(c_t, c_t)$. With probability $(1 - \epsilon)$, players then observe that the action profile (c_t, c_t) was played, leading to a continuation payoff of V_{t+1} . With the complementary probability, players do not observe the action profile (c_t, c_t) and then play (c_t, c_t) forever. Hence we can express V_t as a function of V_{t+1} as follows:

$$V_t = (1 - \delta)\pi(c_t, c_t) + \delta\{(1 - \epsilon)V_{t+1} + \epsilon\pi(c_t, c_t)\}.$$

The strategy profile $\sigma_{\mathbf{c}}$ is an equilibrium of the ϵ -almost perfect monitoring game if there are no profitable one-shot deviations. Here we only consider deviations after histories of the type $((c_0, c_0), \dots, (c_{t-1}, c_{t-1}))$, as after other

²⁶This monitoring technology is similar to the “network monitoring” considered in Wolitzky (2013), where players connected in a network will only observe the actions of their neighbors. In Wolitzky (2013) however the network monitoring is private, whereas it remains public in this example.

²⁷In order to have a signal with compact support, we could assume that observations are truncated for very large level of contributions. See footnote 12.

²⁸This history occurs with probability $(1 - \epsilon)^t$.

histories players maintain their contributions constant, which is an equilibrium. The best one-shot deviation possible after such a history is not to increase at all the contribution level between time t and $t + 1$, as after a deviation players continue to cooperate with probability zero. Such a one-shot deviation is not profitable after history $((c_0, c_0), \dots, (c_{t-1}, c_{t-1}))$ if:

$$(1 - \delta)\pi(c_{t-1}, c_t) + \delta\pi(c_{t-1}, c_t) \leq (1 - \delta)\pi(c_t, c_t) + \delta\{(1 - \epsilon)V_{t+1} + \epsilon\pi(c_t, c_t)\}. \quad (2.10)$$

The left-hand side of (2.10) is the payoff from deviating. The current payoff is $\pi(c_{t-1}, c_t)$, and since there is a zero probability of observing the profile (c_t, c_t) the contribution levels become constant forever, yielding a continuation value of $\pi(c_{t-1}, c_t)$. The right-hand side of (2.10) is the payoff from following the prescribed strategies, which is V_t .

Proposition 2.1. *There exists a $\delta > 0$, an $\epsilon \in (0, 1)$ and a sequence $\mathbf{c} := (c_0, c_1, \dots)$ such that $\sigma_{\mathbf{c}}$ is an equilibrium of the ϵ -almost perfect monitoring game such that there is a strictly positive probability of players contributing. The sequence $(c_t)_t$ satisfies the following difference equation:*

$$\pi(c_t, c_{t+1}) = \frac{1}{\delta(1 - \epsilon)}[\pi(c_{t-1}, c_t) - \pi(c_t, c_t)] + \pi(c_t, c_t), \quad t > 0, \quad (2.11)$$

with initial conditions $\bar{c}_{-1} = 0$ and $\bar{c}_0 = c_0$.

Proof. Recall that the strategy $\sigma_{\mathbf{c}}$ is an equilibrium of the ϵ -almost perfect monitoring if there are no one-shot deviations, that is if (2.10) holds for $t \geq 0$ (where $c_{-1} = 0$). First let us rewrite (2.10) both for t and $t + 1$, assuming that the inequalities hold with equality:

$$\pi(c_{t-1}, c_t) = (1 - \delta)\pi(c_t, c_t) + \delta\{(1 - \epsilon)V_{t+1} + \epsilon\pi(c_t, c_t)\}, \quad (2.12)$$

and

$$\pi(c_t, c_{t+1}) = V_{t+1}. \quad (2.13)$$

Multiplying (2.13) by $-\delta(1 - \epsilon)$ and adding it to (2.12), we obtain:

$$\pi(c_{t-1}, c_t) - \delta(1 - \epsilon)\pi(c_t, c_{t+1}) = (1 - \delta)\pi(c_t, c_t) + \delta\epsilon\pi(c_t, c_t).$$

Rearranging the terms then leads to (2.11). \square

Note that (2.11) is similar to the difference equation (2.4) of Lockwood and Thomas (2002) (see also (2.2) in this chapter) but with a modified discount factor $\tilde{\delta} = \delta(1 - \epsilon)$. Indeed, as was argued previously, the effect of ϵ in this setting is to render players less patient than in the perfect monitoring setting.

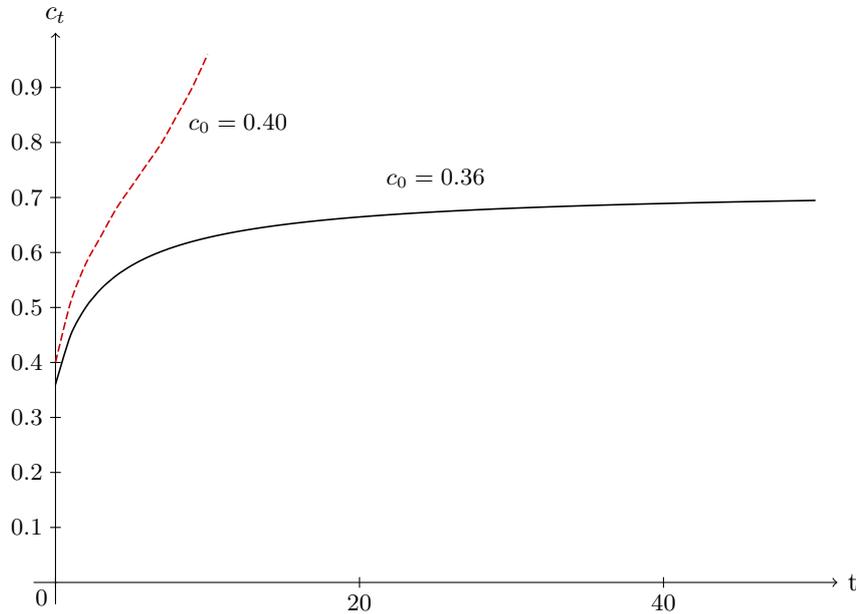
Figure 2.1 shows two different paths that solve the difference equation (2.11) when $\pi(x, y) = -x^2/2 + y$, $\delta = 0.8$ and $\epsilon = 0.10$, with initial values for c_0 being 0.36 and 0.4.²⁹ The function π and the value of δ have been chosen as in the example of Lockwood and Thomas (2002, Figure 1). Note that the solution to (2.11) does not converge when the initial value is 0.4 and that the limit of the sequence when the initial value is 0.36 (the highest initial value consistent with convergence) is 0.72, that is, $\delta(1 - \epsilon)$.

Remark 2.8 (Benefit jump and project completion). A key ingredient in the proof of Theorem 1 is that in a putative equilibrium with positive contributions, increments in contribution levels decrease to zero while remaining strictly positive on a set of histories of positive measure. When the monitoring technology is smooth, small deviations will be hard to detect when increments are close to zero, while consequences will be limited, giving players incentives to deviate and reduce their level of contributions. In this section we presented a counterexample where the monitoring distribution was not smooth and the result of Theorem 2.1 did not hold.

Another way to depart from the assumptions of Section 2.2 would be to consider a jump in the payoff function when the project is completed, as in Marx and Matthews (2000). In their paper the benefit of contributions to player i is $\lambda_i X$, where X is the sum of contributions, if X is below a certain

²⁹Note that the function $\pi(x, y) = -x^2/2 + y$ satisfies Assumptions 2.1 to 2.4.

Figure 2.1: Two solutions to the difference equation (2.11)



threshold X^* (the “completion point”) and $V_i \geq \lambda_i X^*$ when X is above the completion point. If $V_i - \lambda_i X^* > 0$ then it is possible that in a putative equilibrium with positive contributions, increments are bounded away from zero (before the project is completed). The type of deviation considered in the proof of Theorem 2.1 might then no longer be profitable. \diamond

2.6 Conclusion

In a dynamic game where players can contribute to a public project and contributions are irreversible, we have showed that under imperfect public monitoring, when the monitoring technology is sufficiently smooth, cooperation cannot be achieved and players do not contribute in equilibrium. This finding is in stark contrast with the perfect monitoring case, in which there exist equilibria with strictly positive contribution levels. However when the monitoring distribution is no longer smooth, small changes in contribution levels can be detected and cooperation is again possible.

The result relies on one player’s knowledge that contributions of the other

player are close to an upper bound. Such an argument cannot be used in the case of private monitoring, where contributions on the equilibrium path are no longer known. A possible approach in that case would be to consider conditions under which player's contributions converge uniformly. If contribution levels converge uniformly then there would then be a certain time after which player's cannot expect the other player's contribution to increase by more than an arbitrarily small amount, regardless of the past history. Our proof would then hold for private monitoring, or mixed strategies in the public monitoring case. While we do not know whether contributions converge uniformly, Egorov's theorem tells us that if a sequence of random variables converges almost surely then for any $\epsilon > 0$, there exists a subset of $B \subset \Omega$ of measure no more than ϵ and such that the convergence on $\Omega \setminus B$ is uniform.

Chapter 3

Repeated Games with Asymptotically Finite Horizon and Imperfect Monitoring

3.1 Introduction

Games that have a unique stage-game Nash equilibrium have the undesirable property that when repeated a finite number of times, the only subgame-perfect equilibrium is the repetition of the stage-game Nash equilibrium. That is, there are no intertemporal incentives and players often end up along an inefficient path, such as in the finitely repeated prisoner's dilemma.

Infinitely repeated games provide a satisfying answer to this problem. An infinite number of repetitions gives rise to new intertemporal incentives as there is no longer a last period from which incentives unravel. The folk theorem then tells us that every feasible and strictly individually rational payoff can then be supported as the outcome of a subgame-perfect equilibrium.

With constant discounting, infinitely repeated games are stationary. In

particular, the expected duration of the continuation game is the same at any time.¹ One might however argue that even though a game may not have deterministic ending, its expected length should decrease as time passes. For example, two competitors facing a declining demand might expect the probability of interacting to decline over time due to increasing risks of shutdowns. Accordingly, Bernheim and Dasgupta (1995) introduce repeated games with asymptotically finite horizons. In each period, there is a strictly positive probability that the game continues to the next period. However this probability converges to zero.

Bernheim and Dasgupta (1995) study a stage game in which action sets are compact and continuous and which has a unique interior and locally inefficient Nash equilibrium. The game is infinitely repeated and δ_t , the discount factor applied from period t to period $t + 1$, is such that $\delta_t > 0$ for all t and $\lim_{t \rightarrow \infty} \delta_t = 0$. They show that, provided discount factors do not converge to zero too fast, it is possible to have subgame-perfect equilibria in which in each period players receive a payoff strictly higher than the stage-game Nash equilibrium payoff. However, they also show that in any such equilibria, action profiles will converge to the stage-game Nash equilibrium. As time progresses, intertemporal incentives become weak because discount factors decline. To reduce the myopic incentives to deviate from a cooperative equilibrium, action profiles will then have to approach the stage-game Nash equilibrium.

In this chapter we show that when introducing an arbitrarily small amount of smooth noise, all the non-degenerate equilibria of Bernheim and Dasgupta (1995) break down, and the only equilibrium of the dynamic game is the infinite repetition of inefficient stage-game Nash equilibrium. The main difference with Chapter 2 is that we now consider games with an interior stage-game Nash equilibrium. This changes the order of magnitude of gains from deviations along any non-degenerate equilibrium path, as in any non-degenerate equilibrium actions will have to converge to the interior stage-game Nash equi-

¹If the discount factor is interpreted as the probability that the interaction continues in the next period.

librium.

3.2 Perfect monitoring

In this section we briefly present the model of Bernheim and Dasgupta (1995) and discuss their main results.

3.2.1 The stage game

There are two players, $i = 1, 2$.² Let A_i denote player i 's action set and $u_i : A_i \times A_j \rightarrow \mathbb{R}$ denote his utility function. We summarize below all assumptions on the stage game:

Assumption 3.1. $A_i := [a_i, \bar{a}_i]$ is a compact subset of \mathbb{R} . There is a unique stage-game Nash equilibrium a^N which is contained in the interior of $A := A_1 \times A_2$. Each payoff function $u_i : A \rightarrow \mathbb{R}$ is twice continuously differentiable in a and strictly quasi-concave in a_i . Each player's best reply function $\phi_i : A_j \rightarrow A_i$ is continuously differentiable in a neighborhood of a^N .³ The Jacobian matrix of partial derivatives Du has full rank at the stage-game Nash equilibrium.⁴

Remark 3.1. Note that assumption 3.1 exclude from the analysis games with a corner Nash equilibrium, such as the continuous-action prisoner's dilemma studied in Chapter 2. In such games, even as action profiles approach the stage-game Nash equilibrium, myopic incentives to deviate are too high. They are of order one, whereas when the stage-game Nash equilibrium is interior they are of second order from the envelope theorem.

Consider for example the game of Chapter 2 with the following payoff function: $\pi(c_i, c_j) = \pi_1 c_i + \pi_2 c_j$, where $\pi_1 < 0$ and $\pi_2 > 0$. The difference

²This is only to simplify exposition, but all results extend to the case of n players.

³Because u_i is strictly quasi-concave in a_i , each player i has a unique best response to any action a_j .

⁴In the case of two players, the Jacobian matrix evaluated at the Nash equilibrium will be $Du(a^N) = \begin{pmatrix} 0 & \frac{\partial u_1}{\partial a_2}(a^N) \\ \frac{\partial u_2}{\partial a_1}(a^N) & 0 \end{pmatrix}$. The full rank assumption then reduces to the assumption that at the Nash equilibrium $\frac{\partial u_1}{\partial a_2}$ and $\frac{\partial u_2}{\partial a_1}$ are non-zero.

equation (2.2) then reduces to $\Delta c_t = \left(\frac{-\pi_1}{\pi_2}\right)^t \frac{1}{\prod_{\tau=0}^{t-1} \delta_\tau} c_0$. By taking logs we can see that $\lim_{t \rightarrow \infty} \delta_t = 0$ implies that increments in contribution levels must become increasingly large, that is that $\lim_{t \rightarrow \infty} \Delta c_t = \infty$. \diamond

Assumption 3.2. *Utility functions u_i are three times continuously differentiable in a neighbourhood of a^N , $D^2 u_i(a^N)$ is negative definite, $i = 1, 2$ and $D\phi(a^N) - I$ is non-singular.*

Remark 3.2. While Bernheim and Dasgupta (1995) assume that utility functions are three time continuously differentiable, we only use the fact that they are twice continually differentiable in what follows. Moreover we will not require that $D^2 u_i(a^N)$ is negative definite but only that $\partial^2 u_i / \partial a_i^2 < 0$ at the Nash equilibrium. \diamond

Example 3.1. A Cournot duopoly with linear inverse demand and constant marginal cost satisfies assumptions 3.1 and 3.2. Let $A_1 = A_2 = [0, \bar{q}]$, where \bar{q} is sufficiently high, and let $u_i(q_i, q_j) = \max\{0, q_i(1 - q_i - q_j)\}$, $i = 1, 2$, $j \neq i$. \diamond

3.2.2 The dynamic structure

Time is discrete and the game is played infinitely many times: $t = 0, 1, \dots$. Players share a common sequence of discount factors $(\delta_t)_{t \geq 0}$, where $\delta_t \in (0, 1)$ is the discount rate from period t to $t + 1$. The game has an asymptotically finite horizon in the sense that $\lim_{t \rightarrow \infty} \delta_t = 0$.

Payoff streams are evaluated using unnormalized discounting. The payoff to player i from a stream of payoffs (u_i^0, u_i^1, \dots) , evaluated in period k , $k \geq 0$ is:⁵

$$\sum_{t=k}^{\infty} \left\{ \prod_{\tau=k}^{t-1} \delta_\tau \right\} u_t = u_k + \delta_k u_{k+1} + \delta_k \delta_{k+1} u_{k+2} + \delta_k \delta_{k+1} \delta_{k+2} u_{k+3} + \dots$$

Players discount the future in the sense that $\sum_{t=0}^{\infty} \left\{ \prod_{\tau=0}^{t-1} \delta_\tau \right\} < \infty$. For

⁵With the convention that $\prod_{\tau=k}^{k-1} \delta_\tau = 1$.

ease of notation we define $\beta_k^t := \prod_{\tau=k}^{t-1} \delta_\tau$ to be the rate at which period t payoffs are discounted in period k , $k \leq t$.

3.2.3 Cooperation in repeated games with asymptotically finite horizons and perfect monitoring

In this section we briefly present the main results of Bernheim and Dasgupta (1995). First, they find a sufficient condition on the rate of convergence of discount factors to zero to guarantee the existence of a non-degenerate subgame-perfect equilibrium, that is an equilibrium in which players obtain a payoff strictly greater than the stage-game Nash equilibrium payoff in each period. More specifically, the log of discount rates must grow, in absolute value, faster than 2^k :

Assumption 3.3. *There exist $c > 0$ and $\Lambda > 0$ such that $\prod_{k=0}^{\tau-1} \delta_k^{2^{\tau-1-k}} \geq c\Lambda^{2^\tau}$, $\tau \geq 1$. This is equivalent to having $\lim_{\tau \rightarrow \infty} \sum_{k=0}^{\tau-1} \frac{1}{2^{k+1}} \ln(\delta_k) > -\infty$.⁶*

Theorem 3.1 (Bernheim and Dasgupta 1995). *Under assumptions 3.1 and 3.3, there exists a subgame-perfect equilibrium in which players receive a payoff strictly higher than the stage-game NE payoff in each period.*

Under the additional regularity assumption 3.2 Bernheim and Dasgupta (1995) show that assumption 3.3 is not only sufficient but necessary for the existence of non-degenerate subgame-perfect equilibria.

They also establish a folk theorem. More specifically they prove the existence of a time T^* such that for any T , if discount factors are above a certain threshold for $T + T^*$ periods and only then start declining (at the rate implied by assumption 3.3) then for any feasible, interior and strictly individually rational payoff v there is a subgame-perfect equilibrium in which players get v for at least T periods. That is, the decline of the discount factor in the distant future will not affect current behaviour if current discount factors are close to one.

⁶We can see that this assumption is indeed not satisfied in Remark 3.1.

3.3 The model under public monitoring

We now introduce imperfect public monitoring to the model. This section is similar to Section 2.2 and therefore the discussion is kept to a minimum.

At the end of each period, and conditional on an action profile a , players observe a public signal y drawn from a compact set $Y \subset \mathbb{R}^m$ ($m \geq 1$) according to a probability measure $\mathbb{P}_Y(\cdot | a)$. For any measurable $E \subset Y$ we have:

$$\mathbb{P}(y \in E | a) = \int_E \mathbb{P}_Y(dy | a).$$

As in Chapter 2 we assume that the probability measure \mathbb{P}_Y is continuous with respect to action profiles:

Assumption 3.4 (Feller continuity). *There exists a constant K such that $|\mathbb{P}(E | a_1 + \Delta, a_2) - \mathbb{P}(E | a_1, a_2)| \leq K\Delta$ and $|\mathbb{P}(E | a_1, a_2) - \mathbb{P}(E | a_1, a_2 + \Delta)| \leq K\Delta$ for any measurable set $E \subset Y$.*

We denote player i 's realized payoff by u_i^* , which is a function of his current action and the public signal. The ex ante utility function u_i is then the expectation of the ex post payoff:

$$u_i(a_i, a_j) = \int_Y u_i^*(a_i, y) \mathbb{P}_Y(dy | a_1, a_2), \quad \forall a_i, a_j. \quad (3.1)$$

As in Chapter 2 we restrict attention to pure public strategies (see remark 2.5 for a discussion).⁷ A pure public strategy σ_i for player i is a measurable function that specifies an action $\sigma_i(h^t) \in A_i$ after any public history $h^t \in \mathcal{H}$:

$$\sigma_i: \begin{cases} \mathcal{H} & \longrightarrow A_i \\ h^t & \longmapsto \sigma_i(h_i^t) \end{cases}.$$

The monitoring technology, along with a strategy profile $\sigma = (\sigma_1, \sigma_2)$, induce a probability distribution on $\Omega := Y^{\mathbb{N}}$ that we denote by \mathbb{P}_σ . Expecta-

⁷Although we argue that our result would hold in the case of private monitoring, or public monitoring with private mixed strategies.

tions with respect to that probability distribution will be denoted by \mathbb{E}_σ .

Let $V_i(\sigma_i, \sigma_j)$ be the expected payoff of player i from the strategy profile $\sigma = (\sigma_1, \sigma_2)$ and $V_i(\sigma_i, \sigma_j | h^\tau)$ be the continuation payoff from σ after the public history h^τ :⁸

$$V_i(\sigma_i, \sigma_j) := \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \beta_0^t u_i(a^t) \right],$$

$$V_i(\sigma_i, \sigma_j | h^\tau) := \mathbb{E}_\sigma \left[\sum_{t=0}^{\infty} \beta_\tau^{\tau+t} u_i(a^{\tau+t}) | h^\tau \right].$$

A profile of pure strategies (σ_1, σ_2) is a Nash equilibrium if for any $i \in \{1, 2\}$ and any strategy σ' we have that $V(\sigma_i, \sigma_j) \geq V(\sigma', \sigma_j)$.⁹

3.4 Breakdown of cooperation with public monitoring

As in Chapter 2, cooperation breaks down under “smooth” imperfect public monitoring, even for an arbitrarily small amount of noise. There are however two different effects that now go in opposite directions. First, as discount factors converge to zero, intertemporal incentives eventually become weak, which should facilitate deviation from cooperative outcomes. Second, myopic incentives to deviate will be weaker than in Chapter 2. This is because in Chapter 2 contributing was strictly dominated. Here we will see that in any non-degenerate subgame-perfect equilibrium action profiles must converge almost surely to the stage-game Nash equilibrium, making potential gains from deviations only of second order from the envelope theorem. However, with imperfect public monitoring myopic incentives eventually take over:

Theorem 3.2. *Consider an infinitely repeated game with asymptotically finite horizon and imperfect public monitoring that satisfy assumptions 3.1, 3.2, 3.3*

⁸ Even though any finite history occurs with probability zero because there is a continuum of signals, the probability conditional on a history h^τ is well defined as it is the probability measure induced by the continuation strategy profile $\sigma |_{h^\tau}$.

⁹ Given that u_i is bounded, V_i is well defined.

and 3.4. Then with probability one the only Nash equilibrium is the infinite repetition of the unique stage-game Nash equilibrium:¹⁰

$$\mathbb{P}_{\sigma}\{a^t = a^N\} = 1, \forall t \geq 0.$$

3.4.1 Preliminary results

First, we show that in any non-degenerate Nash equilibrium of the dynamic game, for any t , there must be a set of histories of length at least t for which stage-game payoffs are strictly higher than the stage-game Nash equilibrium payoff. That is, incentives for cooperation must be maintained at all times. If that was not the case then there would be a time after which only the stage-game Nash equilibrium would be played. Unravelling would then occur through backward induction and our equilibrium would have to consist of an infinite repetition of the stage-game Nash equilibrium. This lemma is similar to Lemma 2.1 of the previous chapter, which stated that contribution levels had to always increase with positive probability.

Lemma 3.1. *Let (σ_1, σ_2) be a non-degenerate Nash equilibrium of the dynamic game. Then for any $t \geq 0$, there exists $t' \geq t$ such that $\mathbb{P}_{\sigma}(a^{t'} \neq a^N) > 0$.*

Proof. Assume this is not the case and let $T \geq 0$ be the smallest integer such that for any $t' \geq T$ we have $\mathbb{P}_{\sigma}(a^{t'} \neq a^N) = 0$. Given that the Nash equilibrium is non-degenerate it must be that $T > 0$. There is then a set of histories of positive measure such that an action profile different than a^N is played in period $T - 1$ while a^N is played with probability one in the future. Some players will therefore have an incentive to deviate in period $T - 1$, a contradiction with the fact that (σ_1, σ_2) is an equilibrium. \square

Note that there can still be histories after which the stage-game Nash equilibrium is played indefinitely. This set of histories however is of measure less than one.

¹⁰Note that because time is countable, this is equivalent to saying $\mathbb{P}_{\sigma}\{a^t = a^N, \forall t \geq 0\} = 1$.

Lemma 3.1 tells us that intertemporal incentives must remain throughout the game. However, as players are not mutually best-responding, there will be myopic incentives to deviate. As discount factors become low those myopic incentives become more significant relative to intertemporal incentives. To mitigate this action profiles have to approach the stage-game Nash equilibrium. The next lemma is the probabilistic version of Bernheim and Dasgupta (1995, Lemma 2.1.).

Lemma 3.2. *Let (σ_1, σ_2) be a non-degenerate Nash equilibrium of the dynamic game. Then the sequence of action profiles converges almost surely to the stage-game Nash equilibrium: $\mathbb{P}_\sigma(\lim_{t \rightarrow \infty} a^t = a^N) = 1$.*

Proof. In any equilibrium, the instantaneous gain from a deviation must be lower than the maximal punishment incurred after a deviation: that is, for any $t \geq 0$ and $i = 1, 2$ we have that $0 \leq \max_{a'_i} u_i(a'_i, a_j^t) - u_i(a^t) \leq \sum_{k=0}^{\infty} \beta_t^{t+1+k} (\bar{u}_i - \underline{u}_i)$ a.s., where $\bar{u}_i = \max_{a \in A} u_i(a)$ and $\underline{u}_i = \min_{a \in A} u_i(a)$. As δ_t converges to zero, for any $\epsilon > 0$ there exists T_ϵ such that for $t \geq T_\epsilon$ we have $\beta_t^{t+1+k} \leq \epsilon^{1+k}$, $k \geq 0$. Therefore $|\max_{a'_i} u_i(a'_i, a_j^t) - u_i(a^t)| \rightarrow_{t \rightarrow \infty} 0$ a.s., $i = 1, 2$. By the maximum theorem $a_j \mapsto \max_{a_i \in A_i} u_i(a_i, a_j)$ is continuous. As A is compact and there is a unique stage-game Nash equilibrium, this implies that $\mathbb{P}_\sigma\{\lim_{t \rightarrow \infty} a^t = a^N\} = 1$. \square

Corollary 3.1. *For any $\epsilon > 0$ and $\mu > 0$, there exists $T_{\epsilon, \mu} \geq 0$ such that for $t \geq T_{\epsilon, \mu}$ (i) $\delta_t < \epsilon$ and (ii) $\|a^t - a^N\| \leq \mu$ (a.s.).¹¹*

We denote by $\mathcal{H}_{\epsilon, \mu}$ the set of histories of length $T_{\epsilon, \mu}$ that satisfy the two properties of corollary 3.1. The first condition tells us that that discount factors are arbitrarily close to zero while the second condition tells us that action profiles are arbitrarily close to the stage-game Nash equilibrium.

Finally, we will make use of the following lemma which gives us a lower bound on the instantaneous gains from a deviation when the action profile is close to the stage-game Nash equilibrium:

¹¹Note that all norms in \mathbb{R}^2 are equivalent.

Lemma 3.3. *Under assumption 3.2, there exists $b, \bar{\mu} > 0$ such that for all a with $\|a - a^N\| < \bar{\mu}$ we have that*

$$\max_{x \in A_1} u_1(x, a_2) - u_1(a) + \max_{y \in A_2} u_2(a_1, y) - u_2(a) \geq b \|a - a^N\|^2$$

Lemma 3.3 tells us that as action profiles approach the stage-game Nash equilibrium, the sum of gain from each player best responding is at least of order two. The proof is given in Section 3.6

3.4.2 Proof of theorem 3.2

The idea behind the proof of theorem 3.2 is similar to the proof of theorem 2.1, although orders of magnitude are different. In Chapter 2 the gain from a deviation was of order one, whereas here it is only of order two as action profiles approach the interior stage-game Nash equilibrium. The cost from a small deviation will nonetheless be smaller than the gain: first, as the monitoring technology is continuous, a small deviation induces a small change in future payoffs; second as action profiles are close to the stage-game Nash equilibrium, future losses are small; last, discount factors are arbitrarily close to zero.

Let us assume there is a non-generate equilibrium (σ_1, σ_2) of the asymptotically finite repeated game with imperfect monitoring. We will show that there exist a profitable deviation for histories in $\mathcal{H}_{\epsilon, \mu}$ when ϵ and μ are sufficiently small.

Note that at least one of the players must have myopic incentives to deviate for histories in $\mathcal{H}_{\epsilon, \mu}$, given that the action profile is different from the stage-game Nash equilibrium. Consider the following deviations for player i , $i = 1, 2$:

$$\sigma'_i(h^t) = \begin{cases} \phi_i(\sigma_j(h^t)) & \text{if } h^t \in \mathcal{H}_{\epsilon, \mu}, \\ \sigma_i(h^t) & \text{otherwise.} \end{cases}$$

Strategy σ'_i prescribes best responding to player j for any history in $\mathcal{H}_{\epsilon, \mu}$, while agreeing with σ_i otherwise. We now show that at least one of σ'_1 or σ'_2

must be a profitable deviation on $\mathcal{H}_{\epsilon,\mu}$, that is that:

$$V_i(\sigma'_i, \sigma_j | h^t) - V_i(\sigma_i, \sigma_j | h^t) > 0, \quad \forall h^t \in \mathcal{H}_{\epsilon,\mu}, \quad (3.2)$$

for at least one of the players, where for $\sigma \in \{\sigma_i, \sigma'_i\}$,

$$V_1(\sigma, \sigma_j | h^t) = u_i(\sigma, \sigma_j | h^t) + \delta_t \int_Y V_i(\sigma_i, \sigma_j | h^t y) \mathbb{P}_Y(dy | \sigma, \sigma_j, h^t).$$

To do so we show that

$$\sum_i \left\{ V_i(\sigma'_i, \sigma_j | h^t) - V_i(\sigma_i, \sigma_j | h^t) \right\} > 0, \quad \forall h^t \in \mathcal{H}_{\epsilon,\mu}, \quad (3.3)$$

which implies that (3.2) holds for at least one player.

Equation (3.3) can be rewritten as follows, for $y_0 \in Y$:

$$\begin{aligned} \sum_i \delta_t \int_Y [V_i(\sigma_i, \sigma_j | h^t y) - V_i(\sigma_i, \sigma_j | h^t y_0)] \times \\ [\mathbb{P}_Y(dy | \sigma_i, \sigma_j, h^t) - \mathbb{P}_Y(dy | \sigma'_i, \sigma_j, h^t)] \\ < \sum_i \left\{ u_i(\sigma'_i, \sigma_j | h^t) - u_i(\sigma_i, \sigma_j | h^t) \right\}. \end{aligned} \quad (3.4)$$

From Lemma 3.3 we know that for μ sufficiently small the right-hand side of (3.4) is bounded from below by $b\mu^2$, while the integrals on the left-hand side of (3.4) are bounded from above by $a\mu^2$, $a > 0$, from Feller continuity (assumption 3.4) and because all action profiles are within μ of a^N (Corollary 3.1). As δ_t is arbitrarily small (3.4) holds, which implies that at least one of the players has a profitable deviation.

3.5 Conclusion

In this chapter we consider the class of repeated games with asymptotically finite horizons introduced by Bernheim and Dasgupta (1995) and show that non-degenerate equilibria are not robust to the introduction of an arbitrarily

small amount of smooth noise in the monitoring. Below we briefly discuss in turns two questions related to this chapter and the previous one. First, we question what other games exhibit the property that non-degenerate equilibria break down with the introduction of an arbitrarily small amount of smooth noise. Second, we discuss the challenges with extending our results to the private monitoring case.

3.5.1 Tapering off

Takahashi (2005) introduces the notion of “tapering-off” for K -coordination games:

Definition 3.1 (Takahashi (2005)). A K -coordination game *tapers off* if the greatest payoff variation conditional on the first t periods of an efficient history converges to 0 at a rate faster than K^{-t} .¹²

What drives our results in Chapters 2 and 3 is a condition similar to the above tapering-off condition. In Chapter 2, because of irreversibility, the greatest payoff variation conditional on the first t periods converges to zero. This implies that the greatest variation in continuation values also converge to zero. In this chapter, because discount factors converge to zero, it is also the case that the greatest variation in continuation values converges to zero. When the greatest variation in continuation values converges to zero, the size of punishments available to players also converge to zero over time. While this might not be a problem to induce cooperation under perfect monitoring, with an arbitrarily small amount of smooth noise in the monitoring technology there will be a time for which punishment become too small to offset deviations. Small deviations will have small impacts on the monitoring technology, while future consequences will also be minimal.

¹²A K -coordination game is a game in which each player can decrease other player’s payoff by at most K times his own cost of punishment.

3.5.2 Private monitoring

In Chapter 2 the proof of our main result relies on the essential supremum of contribution levels and the knowledge that one player has about the other player being close to his essential supremum. Therefore the proof could not be extended to private monitoring, although we argued that if we had uniform convergence it could. In this chapter the closeness of action profiles to the unique stage-game Nash equilibrium depends on the rate of convergence of discount factors to zero, which is common knowledge. Therefore action profiles in any non-degenerate equilibrium converge uniformly to the stage-game Nash equilibrium. This suggests that the result of this chapter can be extended to private monitoring.

3.6 Proof of Lemma 3.3

Let $u_i^d(a_i, a_j) : A \rightarrow \mathbb{R}$ be the function which returns the highest payoff player i can get by deviating from action profile (a_i, a_j) , that is $u_i^d(a_i, a_j) = u_i(\phi_i(a_j), a_j)$. Consider the following second order Taylor expansions around a^N :

$$\begin{aligned} u_1(a_1, a_2) &= u_1(a^N) + (a_2 - a_2^N) \frac{\partial u_1}{\partial a_2}(a^N) + (a_1 - a_1^N)^2 \frac{1}{2} \frac{\partial^2 u_1}{\partial a_1^2}(a^N) \\ &\quad + (a_2 - a_2^N)^2 \frac{1}{2} \frac{\partial^2 u_1}{\partial a_2^2}(a^N) + (a_1 - a_1^N)(a_2 - a_2^N) \frac{\partial^2 u_1}{\partial a_1 \partial a_2}(a^N) \\ &\quad + o(\|a - a^N\|^2), \end{aligned}$$

where the first order term in a_1 is zero from the first order condition $\frac{\partial u_1}{\partial a_1}(a^N) = 0$ and

$$\begin{aligned}
u_1^d(a_1, a_2) &= u_1(\phi_1(a_2), a_2) \\
&= u_1(\phi_1(a_2^N), a_2^N) + (a_2 - a_2^N) \left[\frac{\partial u_1}{\partial a_1}(\phi_1(a_2^N), a_2^N) \phi_1'(a_2^N) + \frac{\partial u_1}{\partial a_2}(\phi_1(a_2^N), a_2^N) \right] \\
&\quad + (a_2 - a_2^N)^2 \frac{1}{2} \left[\frac{\partial^2 u_1}{\partial a_1 \partial a_2}(\phi_1(a_2^N), a_2^N) \phi_1'(a_2^N) + \frac{\partial^2 u_1}{\partial a_2^2}(\phi_1(a_2^N), a_2^N) \right] \\
&\quad + o(\|a - a^N\|^2) \\
&= u_1(a^N) + (a_2 - a_2^N) \frac{\partial u_1}{\partial a_2}(a^N) + \\
&\quad (a_2 - a_2^N)^2 \frac{1}{2} \left[\frac{\partial^2 u_1}{\partial a_1 \partial a_2}(a^N) \phi_1'(a_2^N) + \frac{\partial^2 u_1}{\partial a_2^2}(a^N) \right] \\
&\quad + o(\|a - a^N\|^2),
\end{aligned}$$

where the last equality is obtained because from the envelope theorem we have that $\frac{\partial u_1}{\partial a_1}(\phi_1(a_2^N), a_2^N) = 0$.

Therefore

$$\begin{aligned}
u_1^d(a_1, a_2) - u_1(a_1, a_2) &= (a_2 - a_2^N)^2 \frac{1}{2} \frac{\partial^2 u_1}{\partial a_1 \partial a_2}(a^N) \phi_1'(a_2^N) \\
&\quad - (a_1 - a_1^N)^2 \frac{1}{2} \frac{\partial^2 u_1}{\partial a_1^2}(a^N) - (a_1 - a_1^N)(a_2 - a_2^N) \frac{\partial^2 u_1}{\partial a_1 \partial a_2}(a^N) + o(\|a - a^N\|^2),
\end{aligned}$$

which can be further simplified into

$$\begin{aligned}
u_1^d(a_1, a_2) - u_1(a_1, a_2) &= -\frac{1}{2} \frac{\partial^2 u_1}{\partial a_1^2}(a^N) \left[(a_2 - a_2^N)^2 (\phi_1'(a_2^N))^2 \right. \\
&\quad \left. + (a_1 - a_1^N)^2 - 2(a_1 - a_1^N)(a_2 - a_2^N) \phi_1'(a_2^N) \right] + o(\|a - a^N\|^2),
\end{aligned}$$

by noting that $\frac{\partial^2 u_i}{\partial a_i \partial a_j} = -\phi_i'(a_2^N) \frac{\partial^2 u_i}{\partial a_2^2}$ from differentiating the first order condition.

We obtain a similar expression for $u_2^d(a_1, a_2) - u_2(a_1, a_2)$ and sum both

gains to obtain:

$$u_1^d(a) - u_1(a) + u_2^d(a) - u_2(a) = \frac{1}{2}(a - a^N)^t [D\phi(a^N) - I]^t U [D\phi(a^N) - I](a - a^N) + o(\|a - a^N\|^2),$$

where

$$D\phi(a^N) - I = \begin{pmatrix} -1 & \phi_1'(a^N) \\ \phi_2'(a^N) & -1 \end{pmatrix},$$

and

$$U = - \begin{pmatrix} \frac{\partial^2 u_1}{\partial a_1^2}(a^N) & 0 \\ 0 & \frac{\partial^2 u_2}{\partial a_2^2}(a^N) \end{pmatrix}.$$

The matrix $[D\phi(a^N) - I]^t U [D\phi(a^N) - I]$ is positive definite as we know that when $a \neq a^N$ then at least one player has a profitable deviation and therefore $u_1^d(a) - u_1(a) + u_2^d(a) - u_2(a) > 0$. Moreover its determinant is non-zero as $D\phi(a^N) - I$ and U are non-singular. U is non-singular as utility functions are strictly concave at the Nash equilibrium, so that $\frac{\partial^2 u_i}{\partial a_i^2}(a^N) \neq 0$, $i = 1, 2$. Let λ_{min} denote the smallest eigenvalue of $[D\phi(a^N) - I]^t U [D\phi(a^N) - I]$, which is strictly positive. We then have the following inequality:

$$u_1^d(a) - u_1(a) + u_2^d(a) - u_2(a) \geq \frac{1}{2} \lambda_{min} \|a - a^N\|^2 + o(\|a - a^N\|^2).$$

Chapter 4

Revision Games with One-Sided Incomplete Information

4.1 Introduction

In this chapter we study a revision game with asymmetric information. Revision games (see Kamada and Kandori, 2011 and Calcagno et al., 2013) model a situation in which players can prepare their actions during a pre-play phase. At the deadline, the action profile last prepared by players is implemented and players receive the corresponding stage-game payoff. Revision opportunities are stochastic and arrive according to independent Poisson processes. There is therefore a positive probability that a player might no longer be able to revise his prepared action before the deadline. While Kamada and Kandori (2011) show that with a continuum of actions the addition of a pre-play phase can increase the set of equilibrium payoffs, Calcagno et al. (2013) show that it can also narrow it down in finite games. In particular in coordination games with two Pareto ranked Nash equilibria the pre-play preparation phase select the Pareto dominant equilibrium, even when it is risk-dominated.

In this chapter we introduce incomplete information and consider a re-

vision game with one-sided incomplete information in which players seek to coordinate on an action which depends on a state of the world known only to Player 1.

We show that equilibria in which information is transmitted to the uninformed player exist if and only if the preparation phase is sufficiently long and characterize such equilibria: first, Player 1 will not signal his private information close to the deadline (Proposition 4.1); second, far away from the deadline, Player 2 can prefer to be miscoordinated with Player 1, provided Player 1 has not yet signalled the state of the world (Proposition 4.2).

The intuition behind Proposition 4.1 is that close to the deadline the informed player strictly prefers to be coordinated on the wrong action rather than being miscoordinated: as time to the deadline approaches, the risk of miscoordination becomes too important relative to the benefit from being coordinated, since the probability of at least one player receiving a subsequent revision opportunity is small.

Proposition 4.2 tells us that there is an option value from being miscoordinated early on in the game for the uninformed player. This occurs even though Player 2's belief about Player 1's action being the correct action is strictly greater than one half.¹ The intuition for this result is that when players are miscoordinated, it requires only one revision opportunity to coordinate on the correct action. When players are already coordinated, it requires either zero or two revision opportunities for players to coordinate on the correct state of the world. When the beliefs of the uninformed player are still close to the $(1/2, 1/2)$ prior he will prefer to be miscoordinated.

We note however that as time to the deadline becomes arbitrarily large, the length of time during which Player 2 prefers to be miscoordinated will remain finite, leaving enough subsequent time for players to coordinate on the correct action. In particular, in an informative equilibrium payoffs will converge to

¹We start from uniform priors and symmetric payoffs (that is, payoffs do not depend on which state of the world is correct but only on whether players coordinate or not on the correct action). In an informative equilibrium, Player 2's belief about Player 1's initial action will therefore weakly increase until Player 1 eventually changes his prepared action.

the Pareto-optimal payoffs as time to the deadline becomes arbitrarily large.

4.2 Setting

	A	B		A	B
A	d,d	$0,1$	A	$1,1$	$1,0$
B	$1,0$	$1,1$	B	$0,1$	d,d
	$\omega = a$			$\omega = b$	

Figure 4.1: A Bayesian coordination game ($d > 1$)

In this section we describe the characteristics of a revision game with one-sided incomplete information in which players seek to coordinate on an action that depends on the state of the world. We first describe the Bayesian game and preparation stage. We then illustrate how the uninformed player revises his beliefs. Finally we describe histories, strategies, and define the equilibrium concept used in the Bayesian revision game.

4.2.1 The Bayesian game

We consider the following Bayesian game, presented in Figure 4.1. There are two players, $N = \{1, 2\}$, two states of the world $\Omega = \{a, b\}$ and two actions for each player: $S_1 = S_2 = \{A, B\}$. Players weakly prefer to be coordinated than not coordinated, and would like to coordinate on the correct state of the world. We assume that Player 1 knows the state of the world while Player 2 is uninformed and has a uniform prior belief: $p(a) = p(b) = 1/2$. Note that both states play an interchangeable role as payoffs in both matrices are similar and prior beliefs are uninformative.

4.2.2 Timing and revision opportunities

Time goes from $-T < 0$ to 0. The game is played at $t = 0$. A positive time $t > 0$ denotes the time remaining until the deadline, while a negative time $-t \in [-T, 0]$ denotes the time in the game. At $-T$, an action profile $(s_1, s_2) \in S_1 \times S_2$ is exogenously given. From $-T$ to 0, players receive

opportunities to revise their actions according to two independent Poisson processes with arrival rates λ_1 and λ_2 . In particular, the probability that revision opportunities are simultaneous is zero.

4.2.3 Belief updating of Player 2

We now assume that the first revision from Player 1 is interpreted as a signal about the correct state of the world to illustrate how Player 2's beliefs evolve. Let $X \in \{A, B\}$ denote the initial action of Player 1 and let $p^x(t)$ denote Player 2's belief that the state is $x \in \{a, b\}$ when $t > 0$ is remaining until the deadline. If Player 1 has revised his action, given that this is interpreted as a signal, Player 2's belief falls to $p^x(t) = 0$.

Consider now the case in which Player 1 has not revised his action since the beginning of the game (that is, for a time interval of length $T - t$). If Player 1 has not yet revised his prepared action, it could be because (i) the state is x , which occurs with probability $1/2$, or (ii) the state is y and Player 1 did not get a revision opportunity, which occurs with probability $\frac{1}{2}e^{-\lambda_1(T-t)}$.² Therefore from Bayes rule we have that:

$$p^x(t) = \frac{1}{1 + e^{-\lambda_1(T-t)}}. \quad (4.1)$$

The beliefs of Player 2 when Player 1 has not revised his prepared action are illustrated in Figure 4.2 for a preparation phase of length 1 and when Player 1 has on average either 0.5, 1 or 2 revision opportunities per unit of time. Note that when Player 1 has not revised his prepared action we have that:

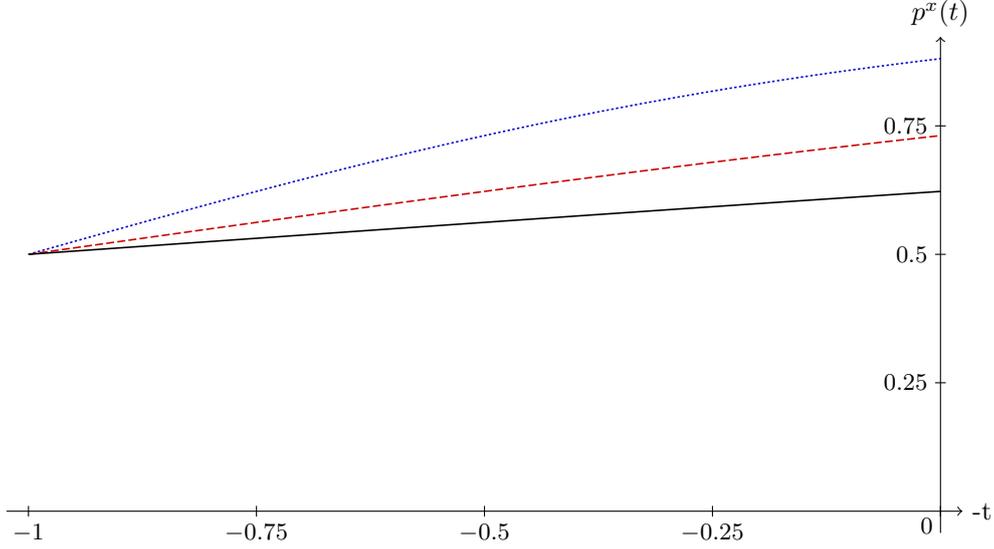
$$\frac{1 - p^x(t)}{p^x(t)} = e^{-\lambda_1(T-t)}, \quad (4.2)$$

and

$$\frac{\partial p^x(t)}{\partial t} = -\lambda_1 p^x(t)(1 - p^x(t)). \quad (4.3)$$

²If $x \in \{a, b\}$ we use y to denote $\{a, b\} \setminus \{x\}$, and similarly for X and Y in $\{A, B\}$.

Figure 4.2: Player 2's beliefs when Player 1 has not revised his prepared action ($T = 1$, $\lambda_1 = 0.5, 1, 2$)



4.2.4 Histories

Let $-t_i^k$ be the time at which Player i receives his k^{th} revision opportunity, and $-t_i^0 := -T$ and let $k_i(t)$ be the number of revision opportunities received by Player i in $[-T, -t)$. Let $X_i^k \in \{A, B\}$ be the action prepared by Player i at time $-t_i^k$ and $X_i(t)$ be the action prepared by Player i at time $-t$. Finally let $r_i(t) \in \{0, 1\}$ indicate whether Player i received a revision opportunity at time $-t$.

A history of the revision game when at time $-t$ takes the following form:

$$h(t) = \left\{ (t_1^k, X_1^k)_{k=0}^{k_1(t)}, (t_2^k, X_2^k)_{k=0}^{k_2(t)}, r_1(t), r_2(t) \right\}.$$

That is, a history indicates all the revision opportunities received up to and including time $-t$ and specifies the prepared actions up to, but not including, time $-t$.

A player only knows about revision opportunities of the other player if he observes a change in the prepared action. Therefore a private history for

player i is take the form of:

$$h_i(t) = \left\{ (t_i^k, X_i^k)_{k=0}^{k_i(t)}, (t_j^{f(k)}, X_j^{f(k)})_{k=0}^{k_j(t)}, r_i(t) \right\}.$$

where $f(0) = 0$ and $t_j^{f(k)}$ is the first time after $t_j^{f(k-1)}$ such that $X_j^{f(k)} \neq X_j^{f(k-1)}$. The set of all private histories for Player i is denoted by H_i .

4.2.5 Strategies

A strategy for player i is a mapping $\sigma_i : H_i \rightarrow \{\emptyset\} \times \Delta(S_i)$ such that $\sigma_i(h_i(t)) = \emptyset$ if $r_i(t) = 0$. (That is, a player can choose an action only when having a revision opportunity.) A pair of strategies (σ_1, σ_2) , along with the Poisson processes, generate a measure $\mathbb{P}_{\sigma_1, \sigma_2}$ on the set of prepared actions at the deadline, $(X_1(0), X_2(0))$.

4.2.6 Equilibrium

A strategy profile (σ_1^*, σ_2^*) is a perfect Bayesian equilibrium of the revision game if for any history $h_i(t)$ such that $r_i(t) = 1$ and any strategy σ_i , we have that

$$\mathbb{E}_{\sigma_1^*, \sigma_2^*} \left[u_i(X_i(0), X_j(0)) \mid h_i(t) \right] \geq \mathbb{E}_{\sigma_1, \sigma_2^*} \left[u_i(X_i(0), X_j(0)) \mid h_i(t) \right],$$

$i = 1, 2$, and Player 2's beliefs are derived from Bayes' rule whenever possible.

4.3 Informative equilibria

We focus on “informative equilibria”, in which Player 1 signals the state of the world through his prepared action. We show that such equilibria exist if and only if the preparation stage is sufficiently long.

Theorem 4.1. *There exists $\tau_1 > 0$ such that an informative equilibrium exists*

if and only if $T \geq \tau_1$. Moreover τ_1 is given by:

$$\tau_1 = \frac{1}{\lambda_1 + \lambda_2} \ln \left[\frac{\lambda_1 + \lambda_2 d}{(d-1)\lambda_2} \right]. \quad (4.4)$$

We then characterize informative equilibria when the preparation stage is sufficiently long. First, we show that close to the deadline Player 1 does no longer signal the state of the world (Proposition 4.1). Therefore there is always a positive probability that player 1 chooses to disregard his private information. This shows why there cannot be an informative equilibrium if the preparation stage is too short. Second, we characterise Player 2's behavior in an informative equilibrium. We show that if the time to the deadline is sufficiently long, the uninformed player will prefer to be miscoordinated with the informed player, provided that the informed player has not revised his prepared action (Proposition 4.2). This shows that an informative equilibrium exists.

4.3.1 No signalling close to the deadline

We first show that close to the deadline, Player 1 will prefer to be coordinated on the wrong action rather signal the correct state of the world. This is because close to the deadline the probability that a future revision opportunity arises is too small relative to the benefit of being coordinated on the correct action.

Note that this result implies the only-if part of Theorem 4.1, that is that there cannot be an informative equilibrium when $T \leq \tau_1$.

Proposition 4.1. *In any equilibrium, there is a time left to the deadline τ such that for $t \leq \tau$ Player 1 prefers to be coordinated with Player 2 on the wrong action than being miscoordinated. Moreover $\tau \geq \tau_1$, where τ_1 is given by (4.4).*

Proof. Let us assume that at any revision opportunity Player 2 seeks to coordinate with Player 1. This is the most favourable case for Player 1 and will therefore gives us the lower bound τ_1 .

Let us assume that Player 2 is choosing the wrong action and let $V(c, 0, t)$ denote Player 1's value when players are coordinated on the wrong action and let $V(nc, 1, t)$ denote Player 1's value when players are miscoordinated and Player 1 is preparing the correct action. We find $V(c, 0, t)$ and $V(nc, 1, t)$ using dynamic programming. In a short time interval dt :

- **Player 1 receives a revision opportunity** with probability $1 - e^{-\lambda_1 dt} \sim -\lambda_1 dt$. He can then choose between being coordinated on the wrong action or signal the correct state of the world.
- **Player 2 receives a revision opportunity** with probability $1 - e^{-\lambda_2 dt} \sim -\lambda_2 dt$ and will coordinate with Player 1 if players are miscoordinated.

Therefore the value of being coordinated on the wrong action for Player 1 satisfies the following equation:

$$V(c, 0, t) \sim \lambda_1 dt \max\{V(c, 0, t - dt), V(nc, 1, t - dt)\} + (1 - \lambda_1 dt)V(c, 0, t - dt).$$

By subtracting $V(c, 0, t - dt)$ from both sides, dividing by dt , and letting dt go to zero, we obtain the following Bellman equation:

$$V_t(c, 0, t) = \lambda_1 \max\{V(nc, 1, t) - V(c, 0, t), 0\}, \quad (4.5)$$

where $V_t(c, 0, t)$ is the derivative of $V(c, 0, t)$ with respect to the time left until the deadline.

Note that $V_t(c, 0, t) \geq 0$, so that the value weakly decreases as the deadline approaches.

Similarly, we have the following Bellman equation for $V(nc, 1, t)$:

$$V_t(nc, 1, t) = \lambda_1 \max\{V(c, 0, t) - V(nc, 1, t), 0\} + \lambda_2 [d - V(nc, 1, t)]. \quad (4.6)$$

The second term corresponds to Player 2 having a revision opportunity, in which case he will coordinate with Player 1 on the correct action, yielding a

payoff of d for both players. Note that we also have $V_t(nc, 1, t) \geq 0$.

When $V(c, 0, t) - V(nc, 1, t) \geq 0$, it is optimal for Player 1 to be coordinated with Player 2 on the wrong action. In that case, (4.5) and (4.6) become

$$V_t(c, 0, t) = 0, \quad (4.7)$$

and

$$V_t(nc, 1, t) = \lambda_1[V(c, 0, t) - V(nc, 1, t)] + \lambda_2[d - V(nc, 1, t)]. \quad (4.8)$$

Because $V(c, 0, t)$ is constant and $V(nc, 1, t)$ is decreasing as the deadline approaches, if $V(c, 0, t) - V(nc, 1, t) \geq 0$ then $V(c, 0, t') - V(nc, 1, t') \geq 0$ for $t' \leq t$: when it is optimal for Player 1 to remain coordinated on the wrong action, it continues to be so as the deadline approaches. This implies that

$$V(c, 0, t) = V(c, 0, 0) = 1, \quad (4.9)$$

and we can therefore rewrite (4.8) as

$$V_t(nc, 1, t) + (\lambda_1 + \lambda_2)V(nc, 1, t) = \lambda_1 + \lambda_2 d. \quad (4.10)$$

Along with the terminal condition $V(nc, 1, 0) = 0$,³ this gives us

$$V(nc, 1, t) = \frac{\lambda_1}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t}) + \frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})d. \quad (4.11)$$

With probability $\frac{\lambda_1}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})$, Player 1 gets the first revision opportunity before the deadline and chooses to coordinate on the wrong action. With probability $\frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})$ Player 2 gets the first revision opportunity before the deadline and coordinates with Player 1 on the correct action, yielding a payoff of d . Finally with the complementary probability no player gets a revision opportunity before the deadline and Player 1 gets a payoff of

³See either the top-right entry of the left matrix or the bottom-left entry of the right matrix in Figure 4.1.

0.

The time τ_1 after which Player 1 prefers to be coordinated on the wrong action is then defined by $V(nc, 1, \tau_1) = V(c, 0, \tau_1) = 1$, that is:

$$\frac{\lambda_1}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)\tau_1}) + \frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)\tau_1})d = 1, \quad (4.12)$$

which gives us

$$\tau_1 = \frac{1}{\lambda_1 + \lambda_2} \ln \left[\frac{\lambda_1 + \lambda_2 d}{(d-1)\lambda_2} \right]. \quad (4.13)$$

□

Note that Player 1 is willing to remain miscoordinated longer as revision opportunities become more frequent, that is $\frac{\partial \tau_1}{\partial \lambda_1} < 0$ and $\frac{\partial \tau_1}{\partial \lambda_2} < 0$.⁴ If λ_1 increases then Player 1 will have more opportunities to coordinate with Player 2 in the future and is therefore willing to remain uncoordinated longer. Similarly, if λ_2 increases then there are more chances that Player 2 will be able to coordinate on the correct action with Player 1 in the future and therefore Player 1 is willing to signal the correct action longer.

Proposition 4.1 tells us that close to the deadline Player 1 will prefer to be coordinated with Player 2 on the wrong action rather than signal the correct action through miscoordination. This is because close to the deadline the risk of miscoordination becomes too important, as it is unlikely that Player 2 will have a revision opportunity and be able to coordinate with Player 1 on the correct action. Therefore if given a revision opportunity Player 1 will prefer to coordinate with Player 2 on the wrong action.

Note that Proposition 4.1 also gives us Player 2's best reply in an informative equilibrium when Player 1 deviates from his equilibrium behaviour (by changing his prepared action a second time before $-\tau_1$):

⁴We have that $\frac{\partial \tau_1}{\partial \lambda_2} = -\frac{1}{(\lambda_1 + \lambda_2)^2} \ln\left(\frac{\lambda_2 d + \lambda_1}{\lambda_2 d - \lambda_1}\right) - \frac{\lambda_1}{\lambda_2(\lambda_1 + \lambda_2)(\lambda_2 d + \lambda_1)} < 0$. Furthermore $\frac{\partial \tau_1}{\partial \lambda_1} = \frac{1}{(\lambda_1 + \lambda_2)^2} \left[\frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 d} + \ln \frac{(d-1)\lambda_2}{\lambda_1 + \lambda_2 d} \right]$. The term in brackets is a strictly increasing function of d for any $\lambda_1 > 0$, $\lambda_2 > 0$ and $\lim_{d \rightarrow \infty} \frac{\lambda_1 + \lambda_2}{\lambda_1 + \lambda_2 d} + \ln \frac{(d-1)\lambda_2}{\lambda_1 + \lambda_2 d} = 0$. Therefore $\frac{\partial \tau_1}{\partial \lambda_1} < 0$ for any $d > 1$, $\lambda_1 > 0$ and $\lambda_2 > 0$.

Corollary 4.1. *In an informative equilibrium, when Player 1 has signalled the state of the world, Player 2 will choose to prepare the corresponding action until τ_2 is left to the deadline, irrespective of Player 1's prepared action, where τ_2 is given by*

$$\tau_2 = \frac{1}{\lambda_1 + \lambda_2} \ln \left[\frac{\lambda_2 + \lambda_1 d}{(d-1)\lambda_1} \right]. \quad (4.14)$$

4.3.2 The uninformed player can prefer miscoordination

We now characterise Player 2's behaviour in an informative equilibrium. In particular, we show that Player 2 might choose to be miscoordinated with Player 1 early on in the game. This is because when players are miscoordinated, coordination on the right action requires only one revision opportunity. When players are coordinated, coordination requires either zero revision opportunities or two revision opportunities. Early on in the game both cases occur with probability close to one half and the uninformed player will prefer to be miscoordinated.

Proposition 4.2. *Consider an informative equilibrium. There exists a length of time T^* such that for any $T \geq T^*$, there is a unique time $t_2^*(T)$ such that for $t \geq t_2^*(T)$ Player 2 prefers to be miscoordinated with Player 1, if Player 1 has not yet revised his initial action. The values of $t_2^*(T)$ and T^* are given by (4.24) and (4.25) and limit beliefs at the threshold $t_2^*(T)$ satisfy*

$$\lim_{T \rightarrow \infty} p^x(t_2^*(T)) = \frac{1}{2} \left[1 + \frac{U(X, Y, 1, \tau_1)}{d} \right] > \frac{1}{2}, \quad (4.15)$$

where X is Player 1's initial prepared-action, $U(X, Y, 1, \tau_1)$ is the payoff to Player 2 when τ_1 is remaining to the deadline, the prepared action profile is (X, Y) and Player 2 believes the state is x .⁵

Equation (4.15) gives us the limit threshold belief after which Player 2 will

⁵That is, the third variable in $U(X, Y, 1, \tau_1)$ is either 1 if Player 1's prepared action is correct and 0 otherwise. Note however that Player 2 cannot know for sure that the state is X given that if the state is X Player 1 will not revised his initial action.

want to coordinate with Player 1 if Player 1 has not yet revised his initial prepared action. This threshold does not depend on the length of the preparation phase, which tells us that $\lim_{T \rightarrow \infty} T - t_2^*(T)$ is a constant and that $\lim_{T \rightarrow \infty} t_2^*(T) - \tau_1 = \infty$.

Proof. Consider the following strategy for Player 1: prepare the correct action until τ_1 remains to the deadline, after which try to coordinate with Player 2. That is,

$$\sigma_1(h^t) = \begin{cases} \text{match the state of the world} & \text{if } t \geq \tau_1 \text{ and } r_1(t) = 1, \\ x_2(t) & \text{if } t < \tau_1 \text{ and } r_1(t) = 1. \end{cases}$$

Note that this strategy is optimal for Player 1 as long as Player 2 is willing to coordinate with Player 1 for $t \in [\tau_1, \tau_1 + \epsilon)$, $\epsilon > 0$.⁶ We will show that this is always the case, and therefore that an informative equilibrium exists, proving the if part of Theorem 4.1.

Let $U(X_1(t), X_2(t), p(t), t)$ denote Player 2's value function when t is left until the deadline, the prepared action profile is $(X_1(t), X_2(t))$ and Player 2's belief about $X_1(t)$ is $p(t)$. Let X denote Player 1's initial action. We first look at Player 2's best responses when the deadline is close, and then when it is far.

Case 1: the deadline is close

We first consider times close to the deadline, that is $t \leq \tau_1$. We know that Player 1 will try to coordinate with Player 2 irrespective of the state of the world. When Player 1 has revised his prepared action before τ_1 and therefore signalled the state of the world to Player 2, we have

$$U(Y, Y, 1, t) = d, \tag{4.16}$$

⁶This is because from Proposition 4.1 τ_1 is the last time Player 1 is willing to prepare the correct action given that Player 2 will coordinate with Player 1.

and

$$U(Y, X, 1, t) = \frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})d + 1 - \frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t}). \quad (4.17)$$

With probability $\frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})$ Player 2 is the first to obtain a revision opportunity and will coordinate to get a payoff of d . With the complementary probability either Player 1 revises his action first or no player receives a revision opportunity, in which case Player 2's payoff at the deadline is 1.

If Player 1 has not revised his action before τ_1 then Player 2's belief is given by $p^x(t) = p^x(\tau_1) = \frac{1}{1 + e^{-\lambda_1(t - \tau_1)}} \in (1/2, 1)$ and Player 2's value is

$$U(X, X, p^x(\tau_1), t) = p^x(\tau_1)d + p^y(\tau_1). \quad (4.18)$$

We now check that when Player 1 has not revised his action prior to τ_1 then Player 2 prefers to be coordinated with Player 1 for $t \leq \tau_1$. Indeed, we have

$$\begin{aligned} U(X, Y, p^x(\tau_1), t) &= \frac{\lambda_2}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})[p^x(\tau_1)d + 1 - p^x(\tau_1)] + \\ &\quad \frac{\lambda_1}{\lambda_1 + \lambda_2}(1 - e^{-(\lambda_1 + \lambda_2)t})[p^x(\tau_1) + p^y(\tau_1)d] + \\ &\quad (1 - e^{-(\lambda_1 + \lambda_2)t})p^x(\tau_1). \end{aligned} \quad (4.19)$$

If Player 2 is the first to have a revision opportunity then he will coordinate with Player 1 on X and have the expected payoff $p^x(\tau_1)d + 1 - p^x(\tau_1)$. If Player 1 is the first to have a revision opportunity then he will coordinate with Player 2 and the expected payoff will be $p^x(\tau_1) + p^y(\tau_1)d < p^x(\tau_1)d + 1 - p^x(\tau_1)$ since $p^x(\tau_1) > 1/2$. Finally if no player can revise then Player 2 gets a payoff of 1 only if the state is x . Hence $U(X, Y, p^x(\tau_1), t) < U(X, X, p^x(\tau_1), t)$ for any $t \leq \tau_1$.

Note that what we have done until now is enough to guarantee the existence of an informative equilibrium, therefore proving Theorem 4.1, even though we have not yet established the best response of Player 2 when the deadline is

far.

Case 2: the deadline is far

Let us now consider times far from the deadline, that is $t \geq \tau_1$. As above we still have $U(Y, Y, 1, t) = d$. When Player 1 has signalled the state by changing his prepared action, we know that he will now wait until τ_1 is left before trying to coordinate with Player 2 on the wrong action again. Therefore if Player 2 obtains a revision opportunity before $-\tau_1$ his payoff will be d . If not his payoff will be $U(Y, X, 1, \tau_1)$, where $U(Y, X, 1, \tau_1)$ is given by (4.17), and we have that

$$U(Y, X, 1, t) = (1 - e^{-\lambda_2(t-\tau_1)})d + e^{-\lambda_2(t-\tau_1)}U(Y, X, 1, \tau_1). \quad (4.20)$$

To find Player 2's value functions when Player 1 has not revised his prepared action yet, $U(X, X, p^x(t), t)$ and $U(X, Y, p^x(t), t)$, we first assume that if Player 2 has a revision opportunity he will choose to remain coordinated with Player 1: $U(X, X, p^x(t), t) \geq U(X, Y, p^x(t), t)$. We know that this is true for $t = \tau_1$. We then get the following partial differential equation for $U(X, X, p^x(t), t)$:

$$U_t(X, X, p^x(t), t) + \lambda_1 p^y(t) U(X, X, p^x(t), t) = \lambda_1 p^y(t) U(Y, X, 1, t),$$

where $U_t(X, X, p^x(t), t) = \frac{\partial U(X, X, p^x(t), t)}{\partial t}$. This is because in a small time interval dt , Player 2 expects Player 1 to change action if the state is y and if Player 1 receives a revision opportunity, which occurs with probability $\sim p^y(t)\lambda_1 dt$. If that is the case then Player 2 gets the value $U(Y, X, 1, t)$, which

is given in (4.20). Given the the boundary condition (4.18) we obtain:⁷

$$\begin{aligned}
U(X, X, p^x(t), t) = & p^x(t)d + \\
& p^y(t) \left\{ \left[1 - \frac{\lambda_1}{\lambda_1 - \lambda_2} e^{-\lambda_2(t-\tau_1)} + \frac{\lambda_2}{\lambda_1 - \lambda_2} e^{-\lambda_1(t-\tau_1)} \right] d + \right. \\
& \left. \frac{\lambda_1}{\lambda_1 - \lambda_2} (e^{-\lambda_2(t-\tau_1)} - e^{-\lambda_1(t-\tau_1)}) U(Y, X, 1, \tau_1) + \right. \\
& \left. e^{-\lambda_1(t-\tau_1)} \right\}. \quad (4.21)
\end{aligned}$$

If the state is x then players will receive a payoff of d . We now explain the term within braces, which corresponds to the payoff when the state is y . The first term in brackets is the probability that during a time interval of length $t - \tau_1$ Player 1 obtains a revision opportunity and Player 2 obtains a subsequent revision opportunity, in which case players receive a payoff of d . With probability $\frac{\lambda_1}{\lambda_1 - \lambda_2} (e^{-\lambda_2(t-\tau_1)} - e^{-\lambda_1(t-\tau_1)})$, Player 1 will have a revision opportunity while Player 2 will not get a subsequent revision opportunity, in which case Player 2's payoff is his value from being uncoordinated and choosing the wrong action at time τ_1 , $U(Y, X, 1, \tau_1)$. Finally, with probability $e^{-\lambda_1(t-\tau_1)}$ Player 1 does not get a revision opportunity during that time interval and players remain coordinated on the wrong action, which yields a payoff of 1.

We now find $U(X, Y, p^x(t), t)$, while still assuming that $U(X, X, p^x(t), t) \geq U(X, Y, p^x(t), t)$, so that if Player 2 has a revision opportunity he will change his prepared action and coordinate with Player 1. Again with probability $\sim p^y(t)\lambda_1 dt$ Player 1 will change his prepared action to y , this time giving a payoff of d . We therefore obtain the following partial differential equation for

⁷Note that this expression is well defined when $\lambda_1 = \lambda_2$ as

$$\lim_{\lambda_1 \rightarrow \lambda} \frac{1}{\lambda_1 - \lambda} (\lambda_1 e^{-\lambda z} - \lambda e^{-\lambda_1 z}) = e^{-\lambda z} (1 + \lambda z)$$

and

$$\lim_{\lambda_1 \rightarrow \lambda} \frac{\lambda_1}{\lambda_1 - \lambda} (e^{-\lambda z} - e^{-\lambda_1 z}) = \lambda z e^{-\lambda z}.$$

$U(X, Y, p^x(t), t)$:

$$\begin{aligned} U_t(X, Y, p^x(t), t) + (\lambda_1 p^y(t) + \lambda_2)U(X, Y, p^x(t), t) \\ = \lambda_1 p^y(t)d + \lambda_2 U(X, X, p^x(t), t). \end{aligned}$$

Along with the boundary condition for $U(X, Y, p^x(\tau_1), \tau_1)$ given by (4.19) we find that

$$\begin{aligned} U(X, Y, p^x(t), t) = U(X, X, p^x(t), t) + e^{-\lambda_2(t-\tau_1)} \times \\ \left\{ p^x(t) \left[U(X, Y, 1, \tau_1) - d + e^{-\lambda_1(T-\tau_1)} [U(X, Y, 0, \tau_1) - d - 1 - U(Y, X, 1, \tau_1)] \right] + \right. \\ \left. p^y(t) \left[d + U(Y, X, 1, \tau_1) \right] \right\}. \quad (4.22) \end{aligned}$$

We now consider the difference $U(X, Y, p^x(t), t) - U(X, X, p^x(t), t)$:

$$\begin{aligned} U(X, Y, p^x(t), t) - U(X, X, p^x(t), t) = e^{-\lambda_2(t-\tau_1)} \times \\ \left\{ p^x(t) \left[U(X, Y, 1, \tau_1) - d + e^{-\lambda_1(T-\tau_1)} [U(X, Y, 0, \tau_1) - d - 1 - U(Y, X, 1, \tau_1)] \right] + \right. \\ \left. p^y(t) \left[d + U(Y, X, 1, \tau_1) \right] \right\}. \quad (4.23) \end{aligned}$$

We can see that if $U(X, Y, p^x(t), t) - U(X, X, p^x(t), t) < 0$ then $U(X, Y, p^x(t'), t') - U(X, X, p^x(t'), t') < 0$ for any $t' < t$. That is, if the uninformed player prefers to be coordinated with the informed player, he will continue to do so in the future. This is because the term in braces is the average of a negative term, weighted by $p^x(t)$,⁸ and a positive term, weighted by $p^y(t)$. As t decreases $p^x(t)$ increases, and therefore $U(X, Y, p^x(t'), t') - U(X, X, p^x(t'), t')$ remains negative.

Given that $U(X, Y, p^x(\tau_1), \tau_1) - U(X, X, p^x(\tau_1), \tau_1) < 0$, this shows that if there exists a time for which $U(X, Y, p^x(t), t) - U(X, X, p^x(t), t) > 0$ then this

⁸It is negative since d is the highest payoff possible.

will still be the case further away from the deadline. Therefore if $t_2^*(T)$ of Proposition 4.2 exists it is unique and given by the equation

$$p^x(t_2^*(T)) \left[U(X, Y, 1, \tau_1) - d + e^{-\lambda_1(T-\tau_1)} [U(X, Y, 0, \tau_1) - d - 1 - U(Y, X, 1, \tau_1)] \right] + p^y(t_2^*(T)) \left[d + U(Y, X, 1, \tau_1) \right] = 0. \quad (4.24)$$

Note that since $\lim_{T \rightarrow \infty} e^{-\lambda_1(T-\tau_1)} = 0$ and $U(X, Y, 1, \tau_1) = U(Y, X, 1, \tau_1)$, if $t_2^*(T)$ exists its limit then satisfies (4.15):

$$\lim_{T \rightarrow \infty} p^x(t_2^*(T)) = \frac{1}{2} \left[1 + \frac{U(X, Y, 1, \tau_1)}{d} \right].$$

We now show that when T is sufficiently large then Player 2 will prefer to be miscoordinated with Player 1 close to T . To do so consider again the difference $U(X, Y, p^x(t), t) - U(X, X, p^x(t), t)$ when $t = T$ (at the start of the game, when $p^x = 1/2$)

$$U(X, Y, p^x(T), T) - U(X, X, p^x(T), T) = \frac{1}{2} e^{-\lambda_2(T-\tau_1)} \times \left\{ U(X, Y, 1, \tau_1) + U(Y, X, 1, \tau_1) + e^{-\lambda_1(T-\tau_1)} [U(X, Y, 0, \tau_1) - d - 1 - U(Y, X, 1, \tau_1)] \right\}.$$

The sign of $U(X, Y, p^x(T), T) - U(X, X, p^x(T), T)$ is similar to the sign of the term in braces, which is an increasing function of T as the bracket is negative, and converges to $U(X, Y, 1, \tau_1) + U(Y, X, 1, \tau_1) > 0$ as $T \rightarrow \infty$. Therefore there is a unique T^* such that when $T \geq T^*$ then Player 2 prefers to remain miscoordinated with Player 1 close to T . Moreover, T^* is given by

$$U(X, Y, 1, \tau_1) + U(Y, X, 1, \tau_1) + e^{-\lambda_1(T^*-\tau_1)} [U(X, Y, 0, \tau_1) - d - 1 - U(Y, X, 1, \tau_1)] = 0 \quad (4.25)$$

□

Propositions 4.1 and 4.2 therefore prove the existence of an informative equilibrium when $T \geq \tau_1$. Player 1's strategy is such that:

- At any revision opportunity such that $t \geq \tau_1$, prepare the action that corresponds to the correct state of the world.
- At any revision opportunity such that $t \leq \tau_1$, coordinate with Player 2.

Player 2's strategy is such that:

- If his belief about a state of the world is 1, play the corresponding action until τ_2 remains to the deadline. After that coordinate with Player 1.
- If Player 2's beliefs are interior:
 - Miscoordinate with Player 1 if $t \geq t_2^*$.
 - Coordinate with Player 1 if $t \leq t_2^*$.

Player 2's beliefs jump to zero or one if Player 1 changes his prepared action before τ_1 is left until the deadline. When Player 1 does not revise his prepared action then Player 2's belief is given by (4.1) until τ_1 and then remains constant.

Moreover the payoffs from such an informative equilibrium converge to the efficient payoff, d , as T becomes arbitrarily large. This is because as T becomes arbitrarily large then so does $t_2^* - \tau_1$.

4.4 Conclusion

In this chapter we study a revision game with one-sided incomplete information in which players seek to coordinate on an action which depends on the state of the world. We show that close to the deadline, the informed player will not signal his private information, while far away from the deadline, the uninformed player prefers to be miscoordinated with the informed player. While it would be interesting to introduce two-sided incomplete information and have players stochastically receive information about the true state of the world,

the model and value functions quickly become intractable. It would also be interesting to look at other classes of games, such as opposing interest games or zero-sum games. In particular, in such games, no information transmission should occur unless the deadline is relatively close.

Bibliography

- Abreu, D., P. K. Dutta, and L. Smith (1994). The folk theorem for repeated games: A Neu condition. *Econometrica* 62(4), 939–948.
- Abreu, D., D. Pearce, and E. Stacchetti (1986). Optimal cartel equilibria with imperfect monitoring. *Journal of Economic Theory* 39(1), 251–269.
- Admati, A. and M. Perry (1991). Joint projects without commitment. *The Review of Economic Studies* 58(2), 259.
- Bernheim, B. and A. Dasgupta (1995). Repeated games with asymptotically finite horizons. *Journal of Economic Theory* 67(1), 129 – 152.
- Bonatti, A. and J. Hörner (2011). Collaborating. *The American Economic Review* 101(2), 632–663.
- Calcagno, R., Y. Kamada, S. Lovo, and T. Sugaya (2013). Asynchronicity and coordination in common and opposing interest games. *Theoretical Economics*.
- Chen, B. (2008). On effective minimax payoffs and unequal discounting. *Econ. Lett.* 100(1), 105–107.
- Chen, B. and S. Takahashi (2012). A folk theorem for repeated games with unequal discounting. *Games and Economic Behavior* 76(2), 571 – 581.
- Compte, O. and P. Jehiel (2003). Voluntary contributions to a joint project with asymmetric agents. *Journal of Economic Theory* 112(2), 334 – 342.

- Compte, O. and P. Jehiel (2004). Gradualism in bargaining and contribution games. *Review of Economic Studies* 71(4), 975–1000.
- Doob, J. (1994). *Measure theory*, Volume 143. Springer.
- Fudenberg, D. and D. Levine (1983). Subgame-perfect equilibria of finite and infinite horizon games. *J. Econ. Theory* 31(2), 251–268.
- Fudenberg, D., D. Levine, and E. Maskin (1994). The folk theorem with imperfect public information. *Econometrica* 62(5), 997–1039.
- Fudenberg, D. and E. Maskin (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54(3), 533–554.
- Gale, D. (2001). Monotone Games with Positive Spillovers. *Games and Economic Behavior* 37(2), 295–320.
- Green, E. and R. Porter (1984). Noncooperative collusion under imperfect price information. *Econometrica*, 87–100.
- Kamada, Y. and M. Kandori (2011). Revision games. *Working paper*.
- Lehrer, E. and A. Pauzner (1999). Repeated games with differential time preferences. *Econometrica* 67(2), 393–412.
- Lockwood, B. and J. P. Thomas (2002). Gradualism and irreversibility. *Review of Economic Studies* 69(2), 339–356.
- Mailath, G. J. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.
- Marx, L. and S. Matthews (2000). Dynamic voluntary contribution to a public project. *Review of Economic Studies* 67(2), 327–358.
- Porter, R. (1983). Optimal cartel trigger price strategies. *Journal of Economic Theory* 29(2), 313–338.

-
- Sugaya, T. (2010). Characterizing the limit set of PPE payoffs with unequal discounting. Mimeo.
- Takahashi, S. (2005). Infinite horizon common interest games with perfect information. *Games and Economic Behavior* 53(2), 231–247.
- Wen, Q. (1994). The “folk theorem” for repeated games with complete information. *Econometrica* 62(4), 949–954.
- Wolitzky, A. (2013). Cooperation with network monitoring. *The Review of Economic Studies* 80(1), 395–427.