

Fear of Miscoordination and the Robustness of Cooperation in Dynamic Global Games with Exit

Sylvain Chassang ^{*,†}

Princeton University

July 16, 2009

Abstract

This paper develops a framework to assess how fear of miscoordination affects the sustainability of cooperation. Building on theoretical insights from Carlsson and van Damme (1993), it explores the effect of small amounts of private information on a class of dynamic cooperation games with exit. Lack of common knowledge leads players to second guess each other's behavior and makes coordination difficult. This restricts the range of equilibria and highlights the role of miscoordination payoffs in determining whether cooperation is sustainable or not. The paper characterizes the range of perfect Bayesian equilibria as the players' information becomes arbitrarily precise. Unlike in one-shot two-by-two games, the global games information structure does not yield equilibrium uniqueness.

KEYWORDS: cooperation, fear of miscoordination, global games, dynamic games, exit games, local dominance solvability.

JEL CLASSIFICATION CODES: C72, C73

*I am deeply indebted to Abhijit Banerjee, Glenn Ellison, Bob Gibbons, and Muhamet Yildiz for time, advice, and encouragement. I am grateful to the co-editor and two anonymous referees for their considerable help in revising the paper. I thank Daron Acemoglu, George-Marios Angeletos, Christophe Chamley, Christian Hellwig, Sergei Izmalkov, Gerard Padro i Miquel, Stephen Morris, Ichiro Obara, Alessandro Pavan, Casey Rothschild, Satoru Takahashi, Olivier Tercieux, Nicolas Vieille and seminar participants at IESE, Insead, IZA, Kellogg, London Business School, MIT, Paris School of Economics, Princeton, Penn State, Stanford, Tinbergen Institute, UCLA and the 2007 Stony Brook Workshop on Global Games for many useful conversations.

[†]E-mail : chassang@princeton.edu / Land-Mail : Department of Economics, Princeton University, Princeton, NJ 08544.

1 Introduction

This paper analyzes the impact of small amounts of incomplete information on a class of dynamic cooperation games with exit. These exit games are infinite horizon two-player games with a fixed discount factor. Every period, players unilaterally choose whether to stay or exit from a joint partnership. Staying is the cooperative action in the sense that staying increases the payoffs of one's partner. In each period t , players' payoffs are affected by an i.i.d. state of the world w_t about which the players obtain noisy signals. This corresponds to a global games information structure. Because the players have different assessments of their environment, there will be miscoordination in equilibrium.¹ This fuels a process by which the players attempt to second guess each other's moves, potentially at the cost of reducing the scope for cooperation.

Within the class of dynamic global games with exit, the paper's main result is a characterization of perfect Bayesian equilibria (PBEs) and sequentially rationalizable strategies as the players' signals become arbitrarily precise. Specifically, the set of sequentially rationalizable strategies is bounded by extreme Markovian equilibria satisfying the following fixed point problem: players stay if and only if it is the risk-dominant action in a one-shot game augmented with the continuation value of playing in the same way in the future.²

In contrast with the case of static games studied by Carlsson and van Damme (1993) or Frankel, Morris and Pauzner (2003), the global games information structure does not lead to equilibrium uniqueness in exit games. Indeed, because the time horizon is infinite, the players can hold multiple self-fulfilling expectations about the value of future interaction. Despite multiplicity, the dominance solvability of static global games carries over in the weaker form of local dominance solvability.³ Moreover, equilibria are locally unique under the global games information structure, whereas there is a continuum of equilibria under complete information.

From the perspective of applications, the fact that the global games perturbation does

¹Here "miscoordination" means that ex post, at least one player would like to change her play unilaterally.

²See Harsanyi and Selten (1988) for a definition and an intuitive discussion of risk-dominance.

³See Moulin (1984) or Guesnerie (2002).

not yield uniqueness does not imply that it is irrelevant. By introducing a realistic risk of miscoordination in equilibrium, it places additional and intuitive restrictions on sustainable levels of cooperation. Losses upon miscoordination, which play no role under complete information, become a central determinant of the players' ability to cooperate. In contrast with a trembling hand or a quantal response approach, this happens even as players become arbitrarily well-informed, and the likelihood of actual miscoordination becomes vanishingly small. With applications in mind, the paper provides a simple criterion for cooperation to be robust in games with approximately constant payoffs.

Because termination payoffs upon exit take a fairly general form, trigger strategies of a repeated game naturally map into strategies of an appropriate exit game. Values upon exit are simply equilibrium values following misbehavior. In that sense, the results of this paper are also relevant for the study of repeated games. While the global games perturbation does not fully resolve the problem of multiplicity, it adds realistic constraints on the sustainability of cooperation: for cooperation to be robust to the global games perturbation, the value of continued cooperation needs to be greater than the deviation temptation plus an additional penalty depending on losses upon miscoordination. Because the framework is very tractable, this provides an operational alternative to focusing on the full information Pareto frontier.

From a methodological perspective, the paper has two main contributions. The first is to show that because of the exit assumption, the lattice theory techniques developed in Milgrom and Roberts (1990), Vives (1990) or Echenique (2004) can be gainfully applied to dynamic cooperation games, even under private information. The second contribution is to show how the Abreu, Pearce, and Stacchetti (1990) approach to dynamic games can be used to study the impact of a global games information structure in a broader set of circumstances than one-shot coordination games. The analysis proceeds in two steps: the first step is to recognize that one-shot action profiles in a perfect Bayesian equilibrium must be Nash equilibria of an augmented one-shot game incorporating continuation values; the second step is to apply global games selection results that hold uniformly over the family of possible augmented games, and derive a fixed point equation for equilibrium continuation

values.

This paper contributes to the literature on the effect of private information in infinite horizon cooperation games. Since Green and Porter (1984), Abreu, Pearce, and Stacchetti (1986, 1990), Radner, Myerson, and Maskin (1986), or Fudenberg, Levine and Maskin (1994), much of this literature has focused on settings in which there is imperfect but public monitoring, so that the relevant histories are always common knowledge, and coordination is never an issue. Under private monitoring, the relevant histories are no longer common knowledge and Mailath and Morris (2002, 2006) have highlighted the importance of miscoordination problems in such circumstances. In particular, they show that even very small departures from public monitoring generate higher order uncertainty that puts significant restrictions on the set of equilibria. The present paper considers an alternative model of miscoordination in which current payoffs rather than past actions are the source of private information. This framework delivers tractable results that can be readily used in applied work investigating the impact of miscoordination fear on cooperation.

This paper also fits in the growing literature on dynamic global games. Much of this literature, however, avoids intertemporal incentives. Levin (2001) studies a global game with overlapping generations. Chamley (1999), Morris and Shin (1999), and Angeletos, Hellwig and Pavan (2006) consider various models of dynamic regime change but shut down dynamic incentives and focus on the endogenous information dynamics that result from agents observing others' actions and new signals of the state of the world. In this sense, these models are models of dynamic herds rather than models of repeated interaction. In two papers that do not rely on private noisy signals as the source of miscoordination, but carry a very similar intuition, Burdzy, Frankel, and Pauzner (2001), and Frankel and Pauzner (2000) obtain uniqueness of equilibrium for a model in which players' actions have inertia and fundamentals follow a random walk. This uniqueness result hinges strongly on the random walk assumption and does not rule out multiplicity in settings where fundamentals follow different processes. Closer to the topic of this paper are Giannitsarou and Toxvaerd (2007) and Ordoñez (2008), both of which extend results from Frankel, Morris, and Pauzner (2003)

and prove an equilibrium uniqueness result for a family of dynamic, finite horizon, recursively supermodular, global games. From the perspective of the present paper, which is concerned with infinite horizon games, their uniqueness result is akin to equilibrium uniqueness in a finitely-repeated dominance-solvable game. Finally, Chassang and Takahashi (2009) explore the more abstract question of robustness to incomplete information in the context of repeated games. Rather than characterizing equilibria of a game with a specific incomplete information structure of interest, they use the approach of Kajii and Morris (1997) and explore the robustness of equilibria to *all* small enough incomplete information perturbations.

The paper is organized as follows. Section 2 presents the setup. Section 3 delineates the mechanics of the paper in the context of a simple example. Section 4 extends the analysis to more general exit games and establishes the main selection results. Section 5 discusses potential applications and alternative models of miscoordination. Proofs are contained in Appendix A, unless mentioned otherwise. Appendix B, available online, contains additional results.

2 Framework

2.1 Exit games

Consider an infinite-horizon game with discrete time $t \in \mathbb{N}$ and two players $i \in \{1, 2\}$ who share the same discount factor $\beta \in (0, 1)$. In every period, the two players simultaneously choose an action from $\mathcal{A} = \{Stay, Exit\}$. Payoffs are indexed by a state of the world $w_t \in \mathbb{R}$. Given a state of the world w_t , player i faces flow payoffs,

$$\begin{array}{c|cc}
 & S & E \\
 \hline
 S & g^i(w_t) & W_{12}^i(w_t) \\
 E & W_{21}^i(w_t) & W_{22}^i(w_t)
 \end{array}
 , \text{ where } i \text{ is the row player.}$$

The sequence of states of the world $\{w_t\}_{t \in \mathbb{N}}$ is an i.i.d. sequence of real numbers drawn from a distribution with density f , c.d.f. F , and convex support. All payoffs functions,

$g^i, W_{12}^i, W_{21}^i, W_{22}^i$ are continuous in w_t . Whenever a player chooses to exit, the game ends and players get continuation values equal to zero. This is without loss of generality since termination payoffs can be included in the flow-payoffs upon exit W_{12}^i, W_{21}^i and W_{22}^i .

At time t , the state of the world w_t is unknown, but each player gets a signal $x_{i,t}$ of the form

$$x_{i,t} = w_t + \sigma \varepsilon_{i,t}$$

where $\sigma \geq 0$ and $\{\varepsilon_{i,t}\}_{i \in \{1,2\}, t \in \mathbb{N}}$ is an i.i.d. sequence of independent random variables taking values in the interval $[-1, 1]$. For simplicity w_t is ex-post observable.

For all $\sigma \geq 0$, let Γ_σ denote this dynamic game with imperfect information. The paper is concerned with equilibria of Γ_σ when the noise level σ is strictly positive but arbitrarily small. According to this notation, Γ_0 denotes the complete information exit game in which the state w_t is publicly observable. Additional assumptions will be introduced in Section 4.

2.2 Solution concepts

Because of the exit structure, at any decision point it must be that players have always chosen to stay in the past. Hence, a history $h_{i,t}$ is simply characterized by a sequence of past and current signals, and past outcomes: $h_{i,t} \equiv \{x_{i,1}, \dots, x_{i,t}; w_1, \dots, w_{t-1}\}$. Let \mathcal{H} denote the set of all such sequences. A pure strategy is a mapping $s : \mathcal{H} \mapsto \{S, E\}$. Denote by Ω the set of pure strategies. For any set of strategies $S \subset \Omega$, let $\Delta(S)$ denote the set of probability distributions over S that have a countable support. The two main solution concepts we will be using are *perfect Bayesian equilibrium* and *sequential rationalizability*. To define these concepts formally, it is convenient to denote by $h_{i,t}^0 \equiv \{x_{i,1}, \dots, x_{i,t-1}; w_1, \dots, w_{t-1}\}$ the histories before players receive period t 's signal but after actions of period $t - 1$ have been taken. A strategy s_{-i} of player $-i$, conditional on the history $h_{-i,t}^0$ having been observed, will be denoted by $s_{-i|h_{-i,t}^0}$. A conditional strategy $s_{-i|h_{-i,t}^0}$ of player $-i$, along with player i 's conditional belief $\mu|h_{i,t}^0$ over $h_{-i,t}^0$, induce a mixed strategy of player $-i$, denoted by $(s_{-i|h_{-i,t}^0}, \mu|h_{i,t}^0)$. Player i 's sequential best-response correspondence, denoted by $BR_{i,\sigma}$, is defined as follows.

Definition 1 (sequential best-response) $\forall s_{-i} \in \Delta(\Omega)$, $s_i \in BR_{i,\sigma}(s_{-i})$ if and only if there exists a set of beliefs μ for player i such that:

- (i) At any history $h_{i,t}^0$, the conditional strategy $s_{i|h_{i,t}^0}$ is a best-reply of player i to the mixed strategy $(s_{-i|h_{-i,t}^0}, \mu|h_{i,t}^0)$;
- (ii) Whenever a history $h_{i,t}^0$ is attainable given s_{-i} , $\mu|h_{i,t-1}$ and player i 's action at $h_{i,t-1}$, then the belief $\mu|h_{i,t}^0$ over $h_{-i,t}^0$ is obtained from $\mu|h_{i,t-1}$ by Bayesian updating;
- (iii) Beliefs $\mu|h_{i,t}$ are obtained from $\mu|h_{i,t}^0$ by Bayesian updating.

Given this definition of sequential best-response, a strategy s_i of player i is associated with a perfect Bayesian equilibrium of Γ_σ if and only if $s_i \in BR_{i,\sigma} \circ BR_{-i,\sigma}(s_i)$. Sequential rationalizability is defined as follows.

Definition 2 (sequential rationalizability) A strategy s_i belongs to the set of sequentially rationalizable strategies of player i if and only if

$$s_i \in \bigcap_{n \in \mathbb{N}} (BR_{i,\sigma}^\Delta \circ BR_{-i,\sigma}^\Delta)^n(\Omega), \quad \text{where } BR_{i,\sigma}^\Delta \equiv BR_{i,\sigma} \circ \Delta.$$

Given strategies s_i , s_{-i} and beliefs upon unattainable histories, let $V_i(h_{i,t})$ denote the value player i expects from playing the game at history $h_{i,t}$. Pairs of strategies and pairs of value functions will respectively be denoted by $\mathbf{s} \equiv (s_i, s_{-i})$ and $\mathbf{V} \equiv (V_i, V_{-i})$.

3 An example

This section focuses on a simple game where two partners repeatedly choose to keep putting effort in their joint project or quit. While this example is fairly restrictive (besides the assumption that state w_t is i.i.d., as maintained throughout, payoffs are symmetric and satisfy strong complementarity properties), it highlights in detail the main steps of the analysis and

the technical difficulties that must be resolved to extend the global games framework to an infinite horizon.

3.1 Payoffs

Consider the exit game with symmetric flow payoffs given by,

	S	E
S	w_t	$w_t - c + \beta V_E$
E	$b + V_E$	V_E

where payoffs are given for the row player only, β is the discount factor, and $c > b \geq 0$.

This game can be thought of as a stylized partnership game in which players repeatedly choose to keep putting effort in their partnership or quit. Value V_E is the discounted present value of the players' constant outside option.⁴ The state w_t represents the expected returns from putting effort in the partnership at time t . Parameter c represents the losses from staying in the partnership when the other player walks out; parameter b (which can be set to 0) represents a potential benefit from cheating on a cooperating partner. When player i exits she obtains her outside option immediately. When player i stays but her partner exits, she obtains her outside option only in the next period.

States of the world w_t are drawn from a distribution with density f and support \mathbb{R} . It is assumed that $\mathbb{E}|w_t| < \infty$ and $V_E > 0$. As in Section 2.1, the complete information version of this game is denoted by Γ_0 , while Γ_σ denotes the game with i.i.d. global games perturbations. Define $M \equiv \frac{1}{1-\beta} \mathbb{E} \max(w_t, V_E + b)$. Any feasible value for playing game Γ_σ is strictly lower than M .

3.2 The complete information case

As a benchmark, this section studies the complete information case, where $\sigma = 0$. Note that the option to exit allows player i to guarantee herself a minimum value $V_i > V_E$.

⁴For instance, $V_E = \frac{1}{1-\beta} w_E$ where w_E is the flow payoff generated by the players' outside option.

Furthermore, independently of what she does, player i is always better off when player $-i$ stays. For this reason, staying will be interpreted as the cooperative action in what follows. Finally, whenever one player exits while the other stays, one player would always prefer to change her decision ex post. Circumstances in which one player stays while the other exits are referred to as miscoordination.

Under complete information, the set of subgame perfect equilibria admits a least cooperative equilibrium and a most cooperative equilibrium, both of which take a simple threshold-form. In the least cooperative equilibrium, players exit if and only if $w_t \leq (1 - \beta)V_E + c$. Note that when $w_t > (1 - \beta)V_E + c$, it is dominant for players to stay. The most cooperative equilibrium is characterized by a threshold \underline{w} such that players stay if and only if $w_t \geq \underline{w}$. This cooperation threshold \underline{w} is the lowest state for which staying can be an equilibrium action. It is associated with the greatest equilibrium continuation value \bar{V} and characterized by the following equations:

$$(1) \quad \underline{w} + \beta\bar{V} = b + V_E$$

$$(2) \quad \bar{V} = \mathbb{E}[(w_t + \beta\bar{V})\mathbf{1}_{w_t > \underline{w}}] + F(\underline{w})V_E,$$

where equation (2) is equivalent to $\bar{V} = \frac{1}{1 - \beta(1 - F(\underline{w}))} [\mathbb{E}(w_t \mathbf{1}_{w_t > \underline{w}}) + F(\underline{w})V_E]$.

Note that parameter c does not enter equations (1) or (2). This means that under complete information the Pareto efficient equilibrium is entirely unaffected by losses upon miscoordination. In contrast, Section 3.3 shows that once private information is introduced, losses upon miscoordination become critical determinants of cooperation.

Under complete information, the partnership game generically admits a continuum of equilibria. Whenever x is such that

$$x < (1 - \beta)V_E + c \quad \text{and} \quad b + V_E < x + \beta V(x),$$

where $V(x) = \frac{1}{1 - \beta(1 - F(x))} [\mathbb{E}(w_t \mathbf{1}_{w_t > x}) + F(x)V_E]$, then the pair of threshold-form strategies such that players stay whenever $w_t \geq x$ and exit whenever $w_t < x$ is an equilibrium. Thresh-

old \underline{w} is the lowest such value of x . When \underline{w} is not a local maximum of $x + \beta V(x)$, then there exists $\eta > 0$ such that all $x \in [\underline{w}, \underline{w} + \eta]$ are equilibrium thresholds.⁵

3.3 The incomplete information case

When players do not observe the state of the world w_t but instead observe a noisy private signal $x_{i,t} = w_t + \sigma \varepsilon_{i,t}$, miscoordination is possible in equilibrium. Players attempt to second guess each other's behavior and assess the miscoordination risk associated with each action. In equilibrium, this risk is particularly high around states of the world at which the players change their behavior. This pushes players towards cautiousness, and reduces the scope for cooperation. The analysis of the dynamic global game Γ_σ proceeds in two steps:

1. The first step shows that for a natural order over strategies the set of rationalizable strategies is bounded by extreme Markovian equilibria. This result relies on the exit game structure and exploits a partial form of monotone best-reply that is sufficient to apply the methods of Milgrom and Roberts (1990) and Vives (1990).
2. The second step characterizes such Markovian equilibria as noise σ goes to 0. Using the dynamic programming approach of Abreu, Pearce and Stacchetti (1990), equilibria of Γ_σ can be analyzed by studying families of one-shot global games augmented with appropriate continuation values. Along with selection results that hold uniformly over families of static global games, this yields a simple asymptotic characterization of Markovian equilibria.

3.3.1 Monotone best-response and extreme equilibria

The first step of the analysis exploits the exit game structure along with payoff complementarities to show that game Γ_σ satisfies a partial form of monotone best-response.

⁵For instance, in the game where $f \sim \mathcal{N}(3, 1)$, $V_E = 5$, $c = 3$, $b = 1$ and $\beta = 0.7$, then $\underline{w} = -1$ and any $x \in [-1, 4.5]$ is an equilibrium threshold.

Definition 3 The partial order \preceq on pure strategies is defined by

$$s' \preceq s \iff \{\forall h \in \mathcal{H}, s'(h) = Stay \Rightarrow s(h) = Stay\}.$$

In words, a strategy s is greater than s' with respect to \preceq if and only if players stay more under strategy s .

Consider a strategy s_{-i} of player $-i$ and player i 's best-reply, $s_i \in BR_{i,\sigma}(s_{-i})$. Pick a history $h_{i,t}$ and denote by V_i the continuation value player i expects at this history. Note that $V_i > V_E$, since for states w_t large enough, players will strictly prefer staying to taking the outside option V_E . Player i 's expected payoffs, $\Pi_S^i(V_i)$ and Π_E^i , from staying and exiting are as follows:

$$\begin{aligned} \Pi_S^i(V_i) &= \mathbb{E}\left[(w_t + \beta V_i)\mathbf{1}_{s_{-i}(h_{-i,t})=S} + (w_t - c + \beta V_E)\mathbf{1}_{s_{-i}(h_{-i,t})=E} \mid h_{i,t}, s_{-i}\right] \\ \Pi_E^i &= \mathbb{E}\left[(b + V_E)\mathbf{1}_{s_{-i}(h_{-i,t})=S} + V_E\mathbf{1}_{s_{-i}(h_{-i,t})=E} \mid h_{i,t}, s_{-i}\right]. \end{aligned}$$

Note that player i 's beliefs about history $h_{-i,t}$ depend both on history $h_{i,t}$ and player $-i$'s strategy, s_{-i} . Player i chooses to stay at history $h_{i,t}$ if and only if $\Pi_S^i(V_i) \geq \Pi_E^i$. We are interested in how i 's best-reply changes when s_{-i} increases. An increase in s_{-i} affects player i 's choice between staying and exiting through three distinct channels:

- increasing s_{-i} changes player i 's continuation value V_i ;
- increasing s_{-i} changes player i 's beliefs about the history $h_{-i,t}$ observed by player $-i$;
- keeping beliefs about $h_{-i,t}$ and continuation values constant, increasing s_{-i} affects player i 's static incentives to stay.

It will be shown that the effects of increasing s_{-i} on continuation values and static incentives both contribute towards making player i stay more as well. The effect on beliefs, however, is ambiguous. For this reason, the analysis *initially* focuses on Markovian strategies, for which this ambiguous effect on beliefs cancels out.

Definition 4 (Markovian strategies) For all $i \in \{1, 2\}$, a strategy s_i is said to be Markovian if $s_i(h_{i,t})$ depends only on player i 's current signal, $x_{i,t}$.

When player $-i$'s strategy is Markovian, then in period t the action taken by player $-i$ depends only on her current signal $x_{-i,t}$, and not on her past history $h_{-i,t}^0$. Furthermore, while player i 's beliefs about $h_{-i,t}^0$ depend on player $-i$'s strategy s_{-i} , player i 's beliefs about $x_{-i,t}$ depend only on her own signal $x_{i,t}$. Hence, when s_{-i} is Markovian, player i 's expected payoffs given her actions need only be conditioned on her own history, and not on player $-i$'s strategy. Given a Markovian strategy s_{-i} , the following facts hold.

Fact 1 (static complementarity) Keeping V_i constant, $\Pi_S^i(V_i) - \Pi_E^i$ is increasing in s_{-i} .

The proof is straightforward given that $V_i > V_E$, $c > b$ and

$$(3) \quad \Pi_S^i(V_i) - \Pi_E^i = \mathbb{E} [(\beta(V_i - V_E) + c - b)\mathbf{1}_{s_{-i}=S} + w_t - c - (1 - \beta)V_E | h_{i,t}].$$

This corresponds to the fact that for any $V_i > V_E$, the one-shot game

	S	E
S	$w_t + \beta V_i$	$w_t - c + \beta V_E$
E	$b + V_E$	V_E

is supermodular and, in particular, exhibits increasing differences in actions.

Fact 2 (dynamic complementarity)

(i) Keeping V_i fixed, both $\Pi_S^i(V_i)$ and Π_E^i are increasing in s_{-i} .

(ii) Keeping s_{-i} fixed, $\Pi_S^i(V_i) - \Pi_E^i$ is increasing in V_i .

Indeed, we have that

$$\begin{aligned} \Pi_S^i(V_i) &= \mathbb{E} [(c + \beta(V_i - V_E))\mathbf{1}_{s_{-i}=S} + w_t - c + \beta V_E | h_{i,t}] \\ \Pi_E^i &= \mathbb{E} [b\mathbf{1}_{s_{-i}=S} + V_E | h_{i,t}], \end{aligned}$$

which implies Fact 2 (i) given that $V_i > V_E$. Fact 2 (ii) follows directly from equation (3).

Fact 2 (i) relies on the fact that when player $-i$ stays more, player i 's continuation value increases. This corresponds to staying being the cooperative action. Fact 2 (ii) is specific to the exit-game structure. In a standard repeated game, where players might deviate to an inferior equilibrium, but cannot simply end the game, increasing all future continuation values does not increase players' incentives to cooperate in the current period.

Together, Facts 1 and 2 are sufficient to establish that game Γ_σ exhibits monotone best-response with respect to Markovian strategies: if s_{-i} is Markovian and increases, then the best-reply $BR_{i,\sigma}(s_{-i})$ shifts up as well. From Fact 1, it follows that, keeping continuation values constant, when strategy s_{-i} increases, player i 's best-reply will increase as well. From Fact 2, it follows that as s_{-i} increases, player i 's continuation value V_i increases, which reinforces player i 's incentives to stay. Monotone best-response for Markovian strategies can be strengthened to show that whenever player $-i$ moves from a strategy s_{-i} to \hat{s}_{-i} , monotone best-reply will hold as long as one of these strategies is Markovian (see Proposition 1 for a formal statement). Monotone best-reply does not generally hold when both strategies are non-Markovian.

This partial form of monotone best-reply is sufficient to replicate the construction of Milgrom and Roberts (1990) or Vives (1990), and show that the set of *all* sequentially rationalizable strategies is bounded by a highest and a lowest Markovian equilibrium. Indeed, “staying always” and “exiting always” are Markovian strategies that clearly bound the set of all possible strategies. By iteratively applying the best-reply mapping, one can bracket the set of sequentially rationalizable strategies between increasing and decreasing sequences of Markovian strategies that converge to extreme Markovian equilibria. Let $\mathbf{s}_\sigma^H = (s_{i,\sigma}^H, s_{-i,\sigma}^H)$ and $\mathbf{s}_\sigma^L = (s_{i,\sigma}^L, s_{-i,\sigma}^L)$ denote the extreme Markovian equilibria of game Γ_σ . Note that since the game is symmetric, these extreme Markovian equilibria must be symmetric. Let us denote by V_σ^H (resp. V_σ^L) the value associated with equilibrium \mathbf{s}_σ^H (resp. \mathbf{s}_σ^L). Values V_σ^H and V_σ^L are respectively the highest and the lowest possible equilibrium values of the exit game Γ_σ . Since extreme equilibria \mathbf{s}_σ^L and \mathbf{s}_σ^H are symmetric, we focus on symmetric

Markovian equilibria for the rest of this section. Appendix B.1 shows that in fact, when payoffs are symmetric, all Markovian equilibria must be symmetric for σ small enough.

3.3.2 Dynamic selection

Since the set of sequentially rationalizable strategies is bounded by extreme, symmetric Markovian equilibria, it is sufficient to focus on symmetric Markovian equilibria to characterize the range of PBEs of Γ_σ . The analysis follows the dynamic programming approach of Abreu, Pearce and Stacchetti (1990). Given a symmetric Markovian equilibrium \mathbf{s}_σ , let us denote by V_σ the value of playing that equilibrium. In any period t , \mathbf{s}_σ induces a one-shot action profile that is a Nash equilibrium of the static coordination game

	S	E
S	$w_t + \beta V_\sigma$	$w_t - c + \beta V_E$
E	$b + V_E$	V_E

where players observe a noisy signal $x_{i,t} = w_t + \sigma \varepsilon_{i,t}$. Let us denote by $\Psi_\sigma(V_\sigma)$ this one-shot incomplete information game. Game $\Psi_\sigma(V_\sigma)$ is essentially a global game that fits into the framework of Carlsson and van Damme (1993). The only difference is that here, both the information structure and the payoffs upon continuation depend on noise parameter σ . Carlsson and van Damme (1993)'s selection results hold when the signal $x_{i,t}$ becomes arbitrarily precise, but keeping fixed the payoff structure. The following uniform selection result resolves this technical difficulty (Fact 3 is a corollary of Lemma A.1 given in Appendix A).⁶

Fact 3 (uniform selection) *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, and all $V \in [V_E, M]$, the one-shot incomplete information game $\Psi_\sigma(V)$ has a unique Nash equilibrium. This Nash equilibrium is characterized by a threshold $x_\sigma^*(V)$ such that for all $i \in \{1, 2\}$,*

⁶Steiner (2008) uses a similar result in the context of a static coordination game in which many workers are assigned to many sectors and must all decide whether to stay or be assigned to an other sector. Steiner (2008) shows that the analysis of the overall matching game can be reduced to the analysis of many 2×2 coordination games with an endogenous value for exit. The overall game admits a unique symmetric equilibrium, which is also the only equilibrium when the number of sectors grows large.

player i stays if and only if $x_{i,t} \geq x_\sigma^*(V)$. Furthermore as σ goes to 0, $x_\sigma^*(V)$ converges uniformly over $[V_E, M]$ to

$$x^{RD}(V) = (1 - \beta)V_E + \frac{b + c}{2} + \beta \frac{V_E - V}{2},$$

the risk-dominance threshold of the one-shot augmented game $\Psi_0(V)$.

Consider $\sigma > 0$ small, so that Fact 3 holds. Given a value $V \in [V_E, M]$, let $\phi_\sigma(V)$ be the value of playing the one-shot game $\Psi_\sigma(V)$ according to its unique equilibrium. We have

$$\begin{aligned} \phi_\sigma(V) = \mathbb{E} & \left[(w_t + \beta V) \mathbf{1}_{\substack{x_{i,t} > x_\sigma^*(V) \\ x_{-i,t} > x_\sigma^*(V)}} + V_E \mathbf{1}_{\substack{x_{i,t} < x_\sigma^*(V) \\ x_{-i,t} < x_\sigma^*(V)}} \right. \\ & \left. + (b + V_E) \mathbf{1}_{\substack{x_{i,t} < x_\sigma^*(V) \\ x_{-i,t} > x_\sigma^*(V)}} + (w_t - c + \beta V_E) \mathbf{1}_{\substack{x_{i,t} > x_\sigma^*(V) \\ x_{-i,t} < x_\sigma^*(V)}} \right]. \end{aligned}$$

By stationarity, the value V_σ of playing Markovian equilibrium \mathbf{s}_σ is also the value of playing the one-shot game $\Psi_\sigma(V_\sigma)$ according to its unique equilibrium. Hence value V_σ must satisfy the fixed point equation, $V_\sigma = \phi_\sigma(V_\sigma)$. Conversely, any fixed point V_σ of ϕ_σ is associated with a Markovian equilibrium of Γ_σ such that players stay and exit according to threshold $x_\sigma^*(V_\sigma)$. Values V_σ^H and V_σ^L are the greatest and smallest fixed points of ϕ_σ .

As σ goes to 0, the threshold $x_\sigma^*(V)$ converges uniformly to $x^{RD}(V)$. Furthermore, as σ goes to 0, the likelihood that $x_{i,t} > x_\sigma^*(V)$ while $x_{-i,t} < x_\sigma^*(V)$ goes to 0 uniformly over $V \in [V_E, M]$. This implies that the value mapping ϕ_σ converges uniformly to the mapping Φ defined by

$$(4) \quad \Phi(V) \equiv \mathbb{E} \left[(w_t + \beta V) \mathbf{1}_{w_t > x^{RD}(V)} + V_E \mathbf{1}_{w_t < x^{RD}(V)} \right].$$

The limiting map Φ and its fixed points are easy to compute, and provide an accurate characterization of the equilibria of the dynamic exit game Γ_σ for the case of σ small. Indeed, the fixed points of ϕ_σ generically converge to the fixed points of Φ : the fact that ϕ_σ converges uniformly to Φ implies that any converging sequence of fixed points $(V_\sigma)_{\sigma > 0}$ of ϕ_σ must converge to a fixed point V of Φ ; conversely, any fixed point V of Φ such that

$\Phi'(V) \neq 1$ is the limit of some sequence $(V_\sigma)_{\sigma>0}$ of fixed points of ϕ_σ .⁷ Let V^L and V^H denote the extreme fixed points of Φ . The associated thresholds $x^{RD}(V^H)$ and $x^{RD}(V^L)$ characterize the highest and lowest levels of cooperation that can be sustained in game Γ_σ as σ goes to 0. Since the set of rationalizable strategies is bounded by Markovian equilibria, whenever Φ has a unique fixed point, game Γ_σ has an asymptotically unique equilibrium.




Figure 1: Equilibria of the partnership game depending on η . $V_E = 5$, $\beta = 0.7$, $\mu = 3$, $c = 3$ and $b = 1$. V_{min} and V_{max} are the extreme equilibrium values under complete information.

Note that the global games perturbation does not necessarily lead to uniqueness in infinite horizon games. Figure 1 plots Φ when the state w_t is drawn from a Gaussian distribution $\mathcal{N}(\mu, \eta^2)$ for different values of η . While the range of equilibria under complete information is roughly the same for the values of η considered, the impact of a global games information structure on the set of equilibria depends significantly on the shape of distribution f . In particular, in the example of Figure 1, Φ admits multiple fixed points when the distribution f has low variance, and a unique fixed point when f has high variance (Appendix B.3 provides a sufficient condition for unique selection along similar lines). Although the unique selection result of Carlsson and van Damme (1993) applies to each static game augmented with continuation values, this augmented game, and which equilibrium is risk-dominant,

⁷More generally, this property holds as long as Φ crosses strictly through the 45° line. See Appendix A.3 for details.

depends on the expectations of agents over future play. It follows that multiple levels of cooperation may be sustainable.

While the exit game Γ_σ can admit multiple asymptotic Markovian equilibria, there are generically finitely many of them. This contrasts with the complete information case in which there is a continuum of equilibria. Section 4.4 also highlights that under the global games perturbation, the set of equilibria is very structured. In particular, the stability and basin of attraction of equilibria with respect to iterated best-reply is essentially characterized by the stability and basin of attraction of fixed points of Φ . This provides additional insights on which equilibria may or may not be selected.

In the context of this paper, the global games perturbation is perhaps best understood as a way to model fear of miscoordination. Because players have different assessments of their environment, one partner may choose to exit while the other stays, and miscoordination can occur in equilibrium.⁸ More importantly, because miscoordination is driven by noise in the information structure, the likelihood of miscoordination depends both on the current state and on what strategies players are using. Miscoordination is most likely in states close to the critical threshold at which players change their behavior. For this reason, in equilibrium, losses upon miscoordination are an important determinant of the sustainability of cooperation, even though players are well informed and the ex ante likelihood of actual miscoordination is small. Fear of miscoordination, rather than miscoordination itself, affects the players' ability to cooperate.

Taking into account fear of miscoordination can affect comparative statics significantly. Equation (4) determines how V^H and $x^{RD}(V^H)$ vary with parameters of interest such as b and c . We have that

$$(5) \quad \frac{\partial \Phi(V)}{\partial b} = -\frac{\partial x^{RD}}{\partial b} f(x^{RD}(V)) (x^{RD}(V) + \beta V - V_E)$$

$$(6) \quad \frac{\partial \Phi(V)}{\partial c} = -\frac{\partial x^{RD}}{\partial c} f(x^{RD}(V)) (x^{RD}(V) + \beta V - V_E).$$

⁸Note that because player i is always better off when player $-i$ chooses to stay, it is difficult for players to exchange credible messages. For this reason the paper does not consider the possibility of cheap talk. See Baliga and Morris (2002) for a detailed discussion of this question.

Since $\frac{\partial x^{RD}}{\partial b} > 0$, $\frac{\partial x^{RD}}{\partial c} > 0$, and $x^{RD}(V) + \beta V > V_E + b$ by definition of $x^{RD}(V)$, it follows that Φ is decreasing in both b and c over the range $[V_E, M]$. Since Φ has finite extreme fixed points, downward shifts of Φ also shift its extreme fixed points downwards. This implies that V^H is strictly decreasing in c and b . It follows that $x^{RD}(V^H)$ is strictly increasing in c and b . Furthermore, since $\frac{\partial x^{RD}}{\partial c} = \frac{\partial x^{RD}}{\partial b}$, it follows from (5) and (6) that changes in deviation temptation b or losses upon miscoordination c have the same impact on the feasible amount of cooperation. This contrasts with the full-information environment where only b affects the sustainability of cooperation.

4 General analysis

This section completes and extends the analysis of Section 3. The framework includes games with asymmetric payoffs and satisfying a weak form of strategic complementarity. Section 4.1 describes the assumptions under which the analysis of Section 3 extends. Section 4.2 shows that under appropriate assumptions, exit games satisfy a partial form of monotone best-response and are bounded by extreme Markovian equilibria. Sections 4.3 uses the dynamic programming approach of Abreu, Pearce and Stacchetti (1990) along with global games selection results to derive a simple fixed point equation characterizing Markovian equilibria. Section 4.4 explores the question of dominance solvability and characterizes the stability and basins of attraction of Markovian equilibria with respect to iterated best-reply.

4.1 Assumptions

The assumptions that follow serve different purposes. Assumption 1 ensures that the values players can obtain are bounded. Assumptions 2, 3 and 4 ensure that the conditions of Carlsson and van Damme hold for one-shot games augmented with the players' possible continuation values. Assumptions 4 and 5 generate strategic complementarities both within

and across time periods. For convenience, recall the general form of flow payoffs,

	S	E	
S	$g^i(w_t)$	$W_{12}^i(w_t)$, where i is the row player.
E	$W_{21}^i(w_t)$	$W_{22}^i(w_t)$	

Assumption 1 (boundedness) *There exists a function $D : \mathbb{R} \rightarrow \mathbb{R}^+$, such that for all $w \in \mathbb{R}$, $D(w) \geq \max_{i,k,l \in \{1,2\}} \{|g^i(w)|, |W_{kl}^i(w)|\}$, and $\int_{-\infty}^{+\infty} D(w)f(w)dw < +\infty$.*

This assumption is fairly unrestrictive but still necessary given that in many natural examples w_t has unbounded support. Let m_i and M_i respectively denote the min-max and maximum values of player i in the complete information game Γ_0 . The maximum value M_i will be used in Assumption 2 while the min-max value m_i appears in Assumptions 4 and 5.

Assumption 2 (dominance) *There exist real numbers $\underline{w} < \bar{w}$, in the support of f , such that for all $i \in \{1, 2\}$,*

$$g^i(\underline{w}) + \beta M_i - W_{21}^i(\underline{w}) < 0 \quad \text{and} \quad W_{12}^i(\underline{w}) - W_{22}^i(\underline{w}) < 0 \quad (\text{Exit dominant})$$

$$\text{and} \quad W_{12}^i(\bar{w}) - W_{22}^i(\bar{w}) > 0 \quad \text{and} \quad g^i(\bar{w}) + \beta m_i - W_{21}^i(\bar{w}) > 0 \quad (\text{Staying dominant}).$$

Assumption 3 (increasing differences in the state of the world) *For all $i \in \{1, 2\}$, $g^i(w_t) - W_{21}^i(w_t)$ and $W_{12}^i(w_t) - W_{22}^i(w_t)$ are strictly increasing over $w_t \in [\underline{w}, \bar{w}]$, with a slope greater than some real number $r > 0$.*

Definition 5 *For any functions $V_i, V_{-i} : \mathbb{R} \rightarrow \mathbb{R}$, let $G(V_i, V_{-i}, w_t)$ denote the complete information one-shot game⁹*

	S	E	
S	$g^i(w_t) + \beta V_i(w_t)$	$W_{12}^i(w_t)$	
E	$W_{21}^i(w_t)$	$W_{22}^i(w_t)$	

⁹In what follows value functions will frequently take a unique value. In those cases we will identify the function and the value it takes.

where i is the row player. Let $\Psi_\sigma(V_i, V_{-i})$ denote the corresponding one-shot global game in which players observe signals $x_{i,t} = w_t + \sigma\varepsilon_{i,t}$.

Assumption 4 (coordination) For all states of the world w_t , the one-shot game $G(m_i, m_{-i}, w_t)$ has a pure strategy Nash equilibrium and all pure equilibria belong to $\{(S, S), (E, E)\}$.

Recall that m_i is player i 's min-max value in the full information game Γ_0 . When Assumptions 2 and 3 hold, Assumption 4 is equivalent to the fact that for all $i \in \{1, 2\}$, if the state w_i is such that $W_{12}^i(w_i) - W_{22}^i(w_i) = 0$, then we have that $g^i(w_i) + \beta m_i - W_{21}^i(w_i) > 0$ and $g^{-i}(w_i) + \beta m_{-i} - W_{21}^{-i}(w_i) > 0$. In words, whenever the state is high enough for player i to stay although player $-i$ exits, then player $-i$'s best-reply under complete information is to stay as well. This can be seen as a single-crossing property of the kind identified by Milgrom and Shannon (1994). It ensures that augmented one-shot games $\Psi_\sigma(\mathbf{V})$ exhibit strategic complementarities.¹⁰ It is strictly weaker than assuming that such one-shot games are supermodular. It is easy to check that Assumption 4 holds for the partnership game since $m_i > V_E$ and $c > b$.

Together Assumptions 2, 3 and 4 correspond to Carlsson and van Damme's assumption that states of the world are connected to dominance regions by a path that is entirely contained in the risk-dominance region of one of the equilibria. Assumption 4 ensures that at any state of the world w and for any pair of individually rational continuation values \mathbf{V} , either (S, S) or (E, E) is the risk-dominant equilibrium of $G(\mathbf{V}, w)$. Assumption 3 implies that there exists a risk-dominant threshold $x^{RD}(\mathbf{V})$ such that (S, S) is risk-dominant in $G(\mathbf{V}, w)$ if and only if $w \geq x^{RD}(\mathbf{V})$.

Assumption 5 (staying benefits one's partner) For all players $i \in \{1, 2\}$ and all states of the world $w \in [\underline{w}, \bar{w}]$, $g^i(w) + \beta m_i - W_{12}^i(w) \geq 0$ and $W_{21}^i(w) - W_{22}^i(w) \geq 0$.

¹⁰Note that if Assumption 4 is satisfied, then for any function $\mathbf{V}(w_t) = (V_i, V_{-i})$ taking values in $[m_i, +\infty) \times [m_{-i}, +\infty)$, the game $G(\mathbf{V}, w_t)$ also has a pure strategy equilibrium, and its pure equilibria also belong to $\{(S, S), (E, E)\}$. Indeed, whether (E, E) is an equilibrium or not does not depend on the value of (V_i, V_{-i}) . Furthermore, if (S, S) is an equilibrium when $\mathbf{V} = (m_i, m_{-i})$, then it is also an equilibrium when the continuation values of player i and $-i$ are greater than m_i and m_{-i} .

Recall that $[\underline{w}, \bar{w}]$ corresponds to states of the world where there need not be a dominant action. Assumption 5 means that under full information, over the range $[\underline{w}, \bar{w}]$, and independently of her own action, player i is weakly better off whenever player $-i$ stays.

This assumption is necessary to obtain dynamic complementarities. If it did not hold, staying more in the future would reduce current continuation values and lead players to stay less in the current period. In the partnership game example, this assumption corresponds to Fact 2 (i) of Section 3.3. This assumption rules out exit games in which “exiting is good”, such as wars of attrition or bargaining games, and restricts attention to games where staying is indeed a cooperative action.

4.2 Monotone best-response and extreme equilibria

This section exploits the exit structure along with Assumptions 4 and 5 to show that for noise σ small, game Γ_σ exhibits a partial form of monotone best-response. In turn this suffices to show the existence of extreme Markovian equilibria that bound the set of sequentially rationalizable strategies. The definitions of partial order \preceq and Markovian strategies given in Section 3 still apply here.

Assumption 4 implies that given continuation values \mathbf{V} , the one-shot augmented game $\Psi_\sigma(\mathbf{V})$ exhibits monotone best-response for σ small enough.¹¹ Assumption 5, that staying benefits one’s partner, implies that strategic complementarities hold across periods as well. This suffices to show a partial form of monotone best-response and the existence of extreme Markovian strategies.

Consider a strategy s_{-i} of player $-i$ and a history $h_{i,t}$ observed by player i . From the perspective of player i , at history $h_{i,t}$, the one-period action profile $s_{-i}(x_{-i,t}, h_{-i,t}^0)$ of player $-i$ can be represented as a mapping from player $-i$ ’s current signal $x_{-i,t}$ to lotteries over $\{stay, exit\}$. Let us denote by $a_{-i|h_{i,t}} : \mathbb{R} \rightarrow \Delta\{stay, exit\}$ this one-shot action profile. The

¹¹The restriction on noise σ comes from the fact Assumption 4 only implies a single-crossing property à la Milgrom and Shannon (1994). Single-crossing is not a sufficient condition for monotone best-response under incomplete information.

order \preceq on dynamic strategies extends to one-shot action profiles as follows:

$$a' \preceq a \iff \{a.s. \forall x \in \mathbb{R}, \text{Prob}[a'(x) = \text{Stay}] \leq \text{Prob}[a(x) = \text{Stay}]\}.$$

Note that if s_{-i} is Markovian, then $s_{-i}(x_{-i,t}, h_{-i,t}^0)$ does not depend on $h_{-i,t}^0$, and $a_{-i|h_{i,t}}$ is effectively a mapping from \mathbb{R} to $\{\text{stay}, \text{exit}\}$. For any mapping V_i that maps player i 's current signal, $x_{i,t} \in \mathbb{R}$, to a continuation value $V_i(x_{i,t})$, and any mapping $a_{-i} : \mathbb{R} \rightarrow \Delta\{\text{stay}, \text{exit}\}$, one can define $BR_{i,\sigma}(a_{-i}, V_i)$, as the one-shot best-response correspondence of player i when she expects a continuation value V_i and player $-i$ uses action profile a_{-i} . The next lemma establishes that the best-reply mappings for one-shot action profiles and for dynamic strategies admit highest and lowest elements: a basic property necessary to apply the tools of lattice theory.

Lemma 1 *For any $\sigma > 0$, we have that*

- (i) *For any one-shot action profile a_{-i} and any value function V_i , $BR_{i,\sigma}(a_{-i}, V_i)$ admits a lowest and a highest element with respect to \preceq . These are respectively denoted $BR_{i,\sigma}^L(a_{-i}, V_i)$ and $BR_{i,\sigma}^H(a_{-i}, V_i)$;*
- (ii) *Whenever a strategy s_{-i} of Γ_σ is Markovian, $BR_{i,\sigma}(s_{-i})$ admits a lowest and a highest element with respect to \preceq . These strategies are Markovian and are respectively denoted $BR_{i,\sigma}^L(s_{-i})$ and $BR_{i,\sigma}^H(s_{-i})$.*

The next lemma establishes that the one-shot best-reply mapping satisfies some monotonicity properties with respect to \preceq .

Lemma 2 *There exist $\bar{\sigma} > 0$ and $\nu > 0$ such that*

- (i) *For all constant value functions $V_i \in [m_i - \nu, M_i + \nu]$, and all $\sigma \in (0, \bar{\sigma})$, $BR_{i,\sigma}^H(a_{-i}, V_i)$ and $BR_{i,\sigma}^L(a_{-i}, V_i)$ are increasing in a_{-i} with respect to \preceq ;*
- (ii) *If V and V' are continuation values functions such that for all $h_{i,t} \in \mathcal{H}$, $V(h_{i,t}) \leq V'(h_{i,t})$, then for any a_{-i} , $BR_{i,\sigma}^H(a_{-i}, V) \preceq BR_{i,\sigma}^H(a_{-i}, V')$ and $BR_{i,\sigma}^L(a_{-i}, V) \preceq BR_{i,\sigma}^L(a_{-i}, V')$.*

Point (i) of Lemma 2 is a consequence of Assumption 4. Point (ii) relies on the exit structure: a general increase in future values increases the incentives to stay. Along with Assumption 5, which ensures that strategic complementarities hold across time periods, this allows us to show that Γ_σ exhibits monotone best-response as long as there is a Markovian strategy on one side of the inequality.

Proposition 1 (partial monotone best response) *There exists $\bar{\sigma}$ such that for all $\sigma \in (0, \bar{\sigma})$, whenever s_{-i} is a Markovian strategy, then, for **all** strategies s'_{-i} ,*

$$s'_{-i} \preceq s_{-i} \Rightarrow \{\forall s'' \in BR_{i,\sigma}(s'_{-i}), s'' \preceq BR_{i,\sigma}^H(s_{-i})\}$$

$$\text{and } s_{-i} \preceq s'_{-i} \Rightarrow \{\forall s'' \in BR_{i,\sigma}(s'_{-i}), BR_{i,\sigma}^L(s_{-i}) \preceq s''\}.$$

Proof: Let us show the first implication. Consider a Markovian strategy s_{-i} and any strategy s'_{-i} such that $s'_{-i} \preceq s_{-i}$. Define V_i and V'_i the continuation value functions respectively associated with player i 's best-response to s_{-i} and s'_{-i} . Since s_{-i} is Markovian, V_i is a constant function. Assumption 5, that staying benefits one's partner, implies that at all histories $h_{i,t}$, $V'_i(h_{i,t}) \leq V_i(h_{i,t})$. By point (ii) of Lemma 2, we have that

$$(7) \quad BR_{i,\sigma}^H(a'_{-i}, V'_i(h_{i,t})) \preceq BR_{i,\sigma}^H(a'_{-i}, V_i(h_{i,t})).$$

Since $V_i(h_{i,t})$ is constant we can apply point (i) of Lemma 2. For this, we need to show that $a'_{-i|h_{i,t}} \preceq a_{-i|h_{i,t}}$. This follows from s_{-i} being Markovian and the fact that $s'_{-i} \preceq s_{-i}$. Indeed, whenever $Prob\{a'_{-i|h_{i,t}} = \text{stay}\} > 0$, we must have $Prob\{a_{-i|h_{i,t}} = \text{stay}\} = 1$. Lemma 2 yields that

$$(8) \quad BR_{i,\sigma}^H(a'_{-i}, V_i(h_{i,t})) \preceq BR_{i,\sigma}^H(a_{-i}, V_i(h_{i,t})).$$

Combining equations (7) and (8) we obtain that indeed, for all $s'' \in BR_{i,\sigma}(s'_{-i})$, $s'' \preceq BR_{i,\sigma}^H(s_{-i})$. An identical proof holds for the other inequality. ■

As will be highlighted below, Proposition 1 is the key step to prove the existence of extreme

Markovian equilibria. Furthermore, these extreme equilibria have a simple structure. A strategy s_i is said to take a threshold-form if there exists a value x such that for all histories $h_{i,t}$, $s_i(h_{i,t}) = S$ if and only if $x_{i,t} \geq x$. A strategy of threshold x will be denoted s_x . The following lemma shows that the best-reply to a threshold-form strategy is unique, and is a threshold-form strategy.¹²

Lemma 3 *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$ and any $x \in \mathbb{R}$, there exists $x' \in \mathbb{R}$ such that $BR_{i,\sigma}(s_x) = \{s_{x'}\}$. Moreover, x' is continuous in x .*

Together, Proposition 1 and Lemma 3 imply the following theorem.

Theorem 1 (extreme strategies) *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, sequentially rationalizable strategies of Γ_σ are bounded by a highest and lowest Markovian Nash equilibria, respectively denoted by $\mathbf{s}_\sigma^H = (s_{i,\sigma}^H, s_{-i,\sigma}^H)$ and $\mathbf{s}_\sigma^L = (s_{i,\sigma}^L, s_{-i,\sigma}^L)$.*

Those equilibria take threshold forms : for all $i \in \{1, 2\}$ and $j \in \{H, L\}$, there exists $x_{i,\sigma}^j$ such that $\mathbf{s}_{i,\sigma}^j$ prescribes player i to stay if and only if $x_{i,t} \geq x_{i,\sigma}^j$.

Although Γ_σ is not supermodular, Proposition 1 and Lemma 3 are sufficient for the construction of Milgrom and Roberts (1990) or Vives (1990) to hold. The strategies corresponding to staying always, and exiting always are threshold-form Markovian strategies that bound the set of possible strategies. The idea is then to iteratively apply the best response mappings to these extreme strategies. Proposition 1 and Lemma 3 guarantee that these iterated strategies form converging sequences of Markovian threshold-form strategies.

Let us denote by \mathbf{x}_σ^H and \mathbf{x}_σ^L the pairs of thresholds associated with the highest and lowest equilibria. Note that $\mathbf{s}_\sigma^L \preceq \mathbf{s}_\sigma^H$, but $\mathbf{x}_\sigma^L \geq \mathbf{x}_\sigma^H$, as staying more corresponds to using a lower threshold. Let \mathbf{V}_σ^H and \mathbf{V}_σ^L be the value pairs respectively associated with \mathbf{s}_σ^H and \mathbf{s}_σ^L . Assumption 5 implies the following lemma.

Lemma 4 *\mathbf{s}_σ^H and \mathbf{s}_σ^L are respectively associated with the highest and lowest possible pairs of rationalizable value functions, \mathbf{V}_σ^H and \mathbf{V}_σ^L . More precisely, if s_{-i} is a rationalizable strategy,*

¹²Such results typically require a monotone likelihood ratio assumption – see for instance Athey (2002). Here, this becomes unnecessary as σ goes to 0.

the value function $V_{i,\sigma}$ associated with player i 's best-reply to s_{-i} is such that at all histories $h_{i,t}$, $V_{i,\sigma}^L \leq V_{i,\sigma}(h_{i,t}) \leq V_{i,\sigma}^H$.

The next section characterizes these extreme Markovian equilibria as σ goes to 0.

4.3 Dynamic selection

We can now state the main selection result of the paper. It shows that continuation values associated with Markovian equilibria of Γ_σ must be fixed points of a mapping $\phi_\sigma(\cdot)$ that converges uniformly to an easily computable mapping Φ from \mathbb{R}^2 to \mathbb{R}^2 . This provides explicit bounds for the set of rationalizable values and shows that the set of Markovian equilibria – which is a continuum under full information – typically shrinks to a finite number of elements under a global games information structure. The structure of the proof, given in the Appendix, follows the analysis of Section 3.3.2.

Theorem 2 *Under Assumptions 1, 2, 3, 4 and 5 there exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, there exists a continuous mapping $\phi_\sigma(\cdot) : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, mapping value pairs to value pairs such that,*

- (i) \mathbf{V}_σ^L and \mathbf{V}_σ^H are the lowest and highest fixed points of $\phi_\sigma(\cdot)$;
- (ii) A vector of values $(V_i, V_{-i}) \in \mathbb{R}^2$ is supported by a Markovian equilibrium if and only if it is a fixed point of $\phi_\sigma(\cdot)$;
- (iii) As σ goes to 0, $\phi_\sigma(\cdot)$ converges uniformly over any compact set of \mathbb{R}^2 to an increasing mapping $\Phi : \mathbb{R}^2 \mapsto \mathbb{R}^2$ defined by

$$\Phi(V_i, V_{-i}) = \begin{pmatrix} \mathbb{E}_w [(g^i(w) + \beta V_i) \mathbf{1}_{w > x^{RD}(V_i, V_{-i})} + W_{22}^i(w) \mathbf{1}_{w < x^{RD}(V_i, V_{-i})}] \\ \mathbb{E}_w [(g^{-i}(w) + \beta V_{-i}) \mathbf{1}_{w > x^{RD}(V_i, V_{-i})} + W_{22}^{-i}(w) \mathbf{1}_{w < x^{RD}(V_i, V_{-i})}] \end{pmatrix}$$

where $x^{RD}(V_i, V_{-i})$ is the risk-dominant threshold of the one-shot game $\Psi_0(V_i, V_{-i})$.

A corollary of Theorem 2 is that whenever Φ has a unique fixed point, the set of rationalizable strategies of game Γ_σ converges to a single pair of strategies as σ goes to 0.¹³ More generally, extreme fixed points of Φ characterize the range of PBEs of game Γ_σ for σ small.

Computing these extreme fixed points is particularly easy in the focal case where the distribution f of the state of the world w is very concentrated around a state w_0 , so that the game exhibits approximately constant payoffs. As before, for any $\mathbf{V} \in \mathbb{R}^2$, we denote by $G(\mathbf{V}, w_0)$ the associated one-shot complete information game augmented with continuation value \mathbf{V} . Denote by $\mathbf{V}^{H,*} \equiv \frac{1}{1-\beta}(g^i(w_0), g^{-i}(w_0))$ the value of staying in every period and $\mathbf{V}^{L,*} \equiv (W_{22}^i(w_0), W_{22}^i(w_0))$ the value of immediate exit, when the state is constant and equal to w_0 . Consider a sequence $\{f_n\}_{n \in \mathbb{N}}$ of distributions for state of the world w such that for all $n \in \mathbb{N}$, Assumptions 1 to 5 are satisfied and f_n converges to δ_{w_0} , the unit mass at w_0 .¹⁴ Let Φ_n denote the value mapping associated with f_n and let \mathbf{V}_n^H and \mathbf{V}_n^L denote the highest and lowest fixed points of Φ_n . The following holds.

Proposition 2 (a robustness criterion)

- (i) If (E, E) is risk-dominant in $G(\mathbf{V}^{H,*}, w_0)$, $\lim_{n \rightarrow \infty} \mathbf{V}_n^H = \lim_{n \rightarrow \infty} \mathbf{V}_n^L = \mathbf{V}^{L,*}$;
- (ii) If (S, S) is risk-dominant in $G(\mathbf{V}^{L,*}, w_0)$, $\lim_{n \rightarrow \infty} \mathbf{V}_n^L = \lim_{n \rightarrow \infty} \mathbf{V}_n^H = \mathbf{V}^{H,*}$;
- (iii) If (S, S) is risk-dominant in $G(\mathbf{V}^{H,*}, w_0)$ and (E, E) is risk-dominant in $G(\mathbf{V}^{L,*}, w_0)$, $\lim_{n \rightarrow \infty} \mathbf{V}_n^H = \mathbf{V}^{H,*}$ and $\lim_{n \rightarrow \infty} \mathbf{V}_n^L = \mathbf{V}^{L,*}$.

Because Proposition 2 depends only on payoffs at w_0 , it can be used to define a simple robustness criterion for cooperation in exit games with constant payoffs.¹⁵ Cooperation is robust to the fear of miscoordination if and only if staying is risk-dominant in the one-shot game augmented with the value of playing (S, S) in every period. If instead exiting is risk

¹³A final technical step shows that under generic conditions, fixed points of ϕ_σ converge to the fixed points of Φ . See Appendix A.3 for details.

¹⁴Specifically $\{f_n\}_{n \in \mathbb{N}}$ converges to δ_{w_0} in the sense that $\lim_{n \rightarrow +\infty} \int D|f_n - \delta_{w_0}| = 0$, where D , defined in Assumption 1, is a function that dominates payoff functions.

¹⁵Note that these results correspond to an order of limits where first noise σ goes to 0 and then distribution f goes to δ_{w_0} . It can be shown that this order of limits is the one that corresponds to the most stringent robustness test. Taking the other order of limits brings us back to the complete information case.

dominant in this game, then the only robust equilibrium is for both players to exit in every period. Finally, it may be that staying is risk-dominant in the game where players expect to stay in the future, while exit is risk-dominant in the game where players expect to exit in the future. In that environment, staying always and exiting always are both robust equilibria. As Section 5.2 highlights, this tractable robustness criterion offers a convenient way to explore how global games perturbations may change comparative statics.

4.4 Local dominance solvability

One of the central results of Carlsson and van Damme (1993) is that as the noise term σ becomes small, one-shot global games are dominance solvable. In that sense, selection of the risk-dominant equilibrium is robust to the relaxation of common knowledge of equilibrium strategies, and only relies on common knowledge of rationality. As Section 3 highlights in the context of the partnership game example, the global games perturbation does not necessarily yield unique selection in settings with an infinite horizon, and dominance solvability does not generally hold. This section shows that the global games perturbation gives bite to the weaker concept of local dominance solvability.¹⁶

Nash equilibrium assumes common knowledge of equilibrium strategies. When a game is dominance solvable, as one-shot global games are, common knowledge of the set of all strategies is sufficient to get to equilibrium. A game is locally dominance solvable at an equilibrium \mathbf{s} if common knowledge that strategies belong to some neighborhood of \mathbf{s} yields \mathbf{s} as the only rationalizable strategy profile. This section characterizes the asymptotic local dominance solvability of game Γ_σ around its Markovian equilibria. For this we need to introduce the mapping ξ which maps future cooperation thresholds to current cooperation thresholds.

Definition 6 (the threshold mapping) *For all $x \in \mathbb{R}$, let us denote by $BRV_{i,0}(x)$ player i 's value for best-replying to a strategy of threshold x in the complete information game Γ_0 .*

¹⁶See for instance Guesnerie (2002) or Moulin (1984).

The threshold mapping $\xi : \mathbb{R} \rightarrow \mathbb{R}$ is defined by

$$\forall x \in \mathbb{R}, \quad \xi(x) = x^{RD}(BRV_{i,0}(x), BRV_{-i,0}(x)).$$

Note that by Assumptions 3 and 5, ξ is weakly increasing.

For any x , $\xi(x)$ is the risk-dominance threshold of $\Psi_0(BRV_{i,0}(x), BRV_{-i,0}(x))$, the one-shot game augmented with the value of best-replying against x in the future. If a pair of values \mathbf{V} is a fixed point of Φ , then $x^{RD}(\mathbf{V})$ is a fixed point of ξ . Conversely, if x is a fixed point of ξ , then $(BRV_{i,0}(x), BRV_{-i,0}(x))$ is a fixed point of Φ . The principal reason for introducing ξ is that it is a mapping from \mathbb{R} to \mathbb{R} while Φ is a mapping from \mathbb{R}^2 to \mathbb{R}^2 . This facilitates the study of its fixed points.

Given order \preceq , for any real numbers $y < z$, we can define $[s_z, s_y]$ as the interval of strategies greater than s_z and smaller than s_y .¹⁷ The following result holds.

Theorem 3 (asymptotic local dominance solvability) *Whenever x is a stable fixed point of ξ and $y < z$ belong to the basin of attraction of x with respect to ξ then*

$$\lim_{\sigma \rightarrow 0} \lim_{n \rightarrow +\infty} [BR_{i,\sigma}^\Delta \circ BR_{-i,\sigma}^\Delta]^n([s_z, s_y]) = \{s_x\}.$$

It follows from Theorem 3 that as σ goes to 0, game Γ_σ is asymptotically locally dominance solvable at any Markovian equilibrium associated with a stable fixed point of ξ .¹⁸ In particular, extreme Markovian equilibria \mathbf{s}_σ^H and \mathbf{s}_σ^L are asymptotically locally dominance solvable. Theorem 3 also characterizes the basin of attraction of Markovian equilibria with respect to iterated best-reply. This quantifies the extent to which common knowledge of equilibrium strategies can be relaxed. The greater the basin of attraction, the more common knowledge of equilibrium strategies can be relaxed. Finally, Theorem 3 provides additional insight on the structure of equilibria. First, it implies that there can be no equilibrium strictly

¹⁷This interval includes non-Markovian strategies.

¹⁸Note that computations can be simplified by considering the mapping $\zeta : \mathbb{R} \mapsto \mathbb{R}$ defined by, $\zeta(x) = x^{RD}(NV_i(x), NV_{-i}(x))$, where $NV_i(x) \equiv \frac{1}{1-\beta \text{Prob}(w>x)} \mathbb{E}[g^i + (W_{22}^i - g^i)\mathbf{1}_{x>w}]$. Computing ζ is simpler than computing ξ and both functions coincide around their fixed points. However, ζ need not be increasing.

contained within two consecutive Markovian equilibria (Appendix B.2 provides additional results on non-Markovian equilibria). Second, any Markovian equilibrium associated with an unstable fixed point of ξ is unstable with respect to iterated best-reply.

5 Discussion

5.1 Modeling fear of miscoordination

As has been highlighted in Section 3, in the context of this paper, the global games perturbation is best understood as a way to model fear of miscoordination. The idea that players make noisy private assessments of the world, and that this makes coordination difficult, is reasonable. Still, there are other ways to introduce miscoordination in equilibrium. In particular, trembling hand perturbations and quantal response equilibrium both share this feature. However, they correspond to very different models of miscoordination fear.

In a trembling hand approach, for instance, losses upon miscoordination affect the sustainability of cooperation if and only if the likelihood of trembles is high. In that case however, while losses upon miscoordination affect the choices made by the players, realized behavior approaches randomness. A quantal response approach would share the same drawback.¹⁹ This contrasts with the approach developed in this paper, where the ex ante likelihood of miscoordination is vanishing and players are, on average, very good at predicting their opponent's behavior. Here, losses upon miscoordination affects the sustainability of cooperation by restricting the players' ability to select the efficient equilibrium. When losses upon miscoordination increase, joint exit tends to become a focal point.

One can think of the global games approach as endogenizing the likelihood of trembles. In particular, the likelihood of miscoordination depends on both the state of the world and the strategies that players are using. Even as the players' information becomes arbitrarily good, the likelihood of miscoordination remains large around the critical states at which

¹⁹Note that a variant of quantal response in which players obtain precise signals about one another's payoff shocks would generate predictions qualitatively similar to those of this paper.

players change their behavior. This imposes significant constraints on equilibrium strategies and, as the next section discusses, can significantly alter comparative statics.

5.2 Fear of miscoordination and comparative statics

The robustness criterion of Proposition 2 is a useful tool to explore how fear of miscoordination can affect comparative statics. Consider for instance the partnership game of Section 3, in a setting where the state of the world is approximately constant and equal to w_0 , with $\frac{1}{1-\beta}w_0 > V_E$, so that staying permanently is the efficient outcome. Under complete information, staying is a Nash equilibrium if and only if

$$(9) \quad \frac{1}{1-\beta}w_0 > V_E + b.$$

In contrast, Proposition 2 implies that staying is robust to the introduction of small amounts of private information about the state of the world if and only if

$$(10) \quad \frac{1}{1-\beta}w_0 > [V_E + b] + [(1-\beta)V_E + c - w_0].$$

Whenever the opposite inequality holds, permanent exit is the only robust equilibrium. Condition (9) reflects that under complete information, cooperation is sustainable if and only if the value of continued cooperation is greater than the deviation temptation. Condition (10) highlights that when the state of the world is uncertain and players try to second guess each other's actions, then cooperation is sustainable if and only if the value of continued cooperation is greater than the deviation temptation plus a penalty that corresponds, in this symmetric game, to losses upon miscoordination. Whenever a parameter of interest affects the deviation temptation and losses upon miscoordination differently, taking into account fear of miscoordination may significantly change comparative statics. For instance, Chassang and Padro i Miquel (2008) consider a dynamic model of peace and conflict, and show that the impact of weapon stocks on the sustainability of peace depends crucially on whether fear of miscoordination is taken into account.

This being said, the global games perturbation clearly does not overturn all comparative statics. Consider for instance comparative statics with respect to the discount factor β . Let us interpret value V_E in the partnership game as a discounted value $V_E = \frac{1}{1-\beta}w_E$, where $w_E < w_0$ is the players' flow payoff when both choose not to put effort in the partnership. Conditions (9) and (10) both hold for β close enough to 1. In particular, losses upon miscoordination remain bounded while the difference between continued cooperation and the deviation temptation grows arbitrarily large. As a result there exist approximately efficient equilibria as β approaches 1. This property holds more generally for exit games whose payoffs are reduced-forms for trigger strategies in a repeated game. In such games, losses upon miscoordination remain bounded and fear of miscoordination affects predictions only if the discount factor is not arbitrarily close to 1.

Note that this property need not hold for exit games that are not reduced-forms for trigger strategies. Consider the variation on the partnership game, where players get flow payoffs

	<i>S</i>	<i>E</i>
<i>S</i>	w_0	$w_0 - c$
<i>E</i>	$b + V_E$	V_E

where $V_E = \frac{1}{1-\beta}w_E$ with $0 < w_E < w_0$. In this game, if player i stays while her partner exits, she does not get her outside option V_E in the future.²⁰ Hence, as β goes to 1, losses upon miscoordination also grow arbitrarily large. If $w_0 < 2w_E$, increasing the discount factor β makes it more difficult to sustain cooperation robustly, and as β approaches 1, immediate exit is the only robust equilibrium.

6 Conclusion

This paper provides a framework to model fear of miscoordination in dynamic environments. It analyzes the robustness of cooperation to global games perturbations in a class of dynamic games with exit. In equilibrium, this departure from common knowledge generates a fear of

²⁰One interpretation is that she goes bankrupt.

miscoordination that pushes players away from the full information Pareto efficient frontier, even though actual miscoordination happens with a vanishing probability. Payoffs upon miscoordination, which play no role when considering the Pareto efficient frontier under complete information, determine the extent of the efficiency loss.

The first step of the analysis is to show that rationalizable strategies of exit games are bounded by extreme Markovian equilibria. The second step uses the dynamic programming approach of Abreu, Pearce, and Stacchetti (1990) to recursively apply selection results for one-shot global games. As players' signals become increasingly correlated, this yields a fixed point equation characterizing values associated with Markovian equilibria.

Whenever this fixed point equation has a unique solution, the set of rationalizable strategies of the game with perturbed information converges to a singleton as signals become arbitrarily precise. Unlike in one-shot two-by-two games, infinite horizon exit games can admit multiple equilibria under a global games information structure. This implies that the global games perturbation does not necessarily lead to dominance solvability in exit games with infinite horizon. Studying the less stringent notion of local dominance solvability shows that still, the global games perturbation implies a lot of structure on equilibrium strategies. Among other things, Markovian equilibria are typically locally unique. This contrasts with the complete information game which admits a continuum of equilibria.

Finally, by introducing a realistic risk of miscoordination in equilibrium, the global games perturbation places additional intuitive restrictions on sustainable levels of cooperation. In addition to the deviation temptation, losses upon miscoordination become an important determinant of the sustainability of cooperation. Taking into account the impact of fear of miscoordination on cooperation can significantly change comparative statics. With applications in mind, the paper provides a tractable robustness criterion.

A Proofs

A.1 Proofs for Section 4.2

Given a continuation value function V_i , the expected payoffs upon staying and exiting – respectively denoted by $\Pi_S^i(V_i)$ and Π_E^i – are

$$(11) \quad \Pi_S^i(V_i) = \mathbb{E} \left[W_{12}^i(w) + \{g^i(w) + \beta V_i(h_{i,t}, w) - W_{12}^i(w)\} \mathbf{1}_{s_{-i}=S} | h_{i,t}, s_{-i} \right]$$

$$(12) \quad \Pi_E^i = \mathbb{E} \left[W_{22}^i(w) + \{W_{21}^i(w) - W_{22}^i(w)\} \mathbf{1}_{s_{-i}=S} | h_{i,t}, s_{-i} \right].$$

Proof of Lemma 1: We begin with point (i). An action profile a_i belongs to the set of one-shot best-replies $BR_{i,\sigma}(a_{-i}, V_i)$ if and only if a_i prescribes S when $\Pi_S^i(V_i) > \Pi_E^i$ and prescribes E when $\Pi_S^i(V_i) < \Pi_E^i$. Because ties are possible $BR_{i,\sigma}(a_{-i}, V_i)$ need not be a singleton. However, by breaking the ties consistently in favor of either S or E , one can construct strategies a_i^H and a_i^L that are respectively the greatest and smallest elements of $BR_{i,\sigma}(a_{-i}, V_i)$ with respect to \preceq .

The proof of point (ii) goes as follows. Let V_i be the value player i obtains from best-replying to s_{-i} . Since s_{-i} is Markovian, at any history $h_{-i,t}^0$ the conditional strategy $s_{-i} | h_{-i,t}^0$ is identical to s_{-i} , and the value player i expects conditional on $h_{-i,t}^0$ is always V_i . Hence, $s_i \in BR_{i,\sigma}(s_{-i})$ if and only if the one-shot action profile prescribed by s_i at a history $h_{i,t}^0$ belongs to $BR_{i,\sigma}(s_{-i}, V_i)$, where s_{-i} is identified with its one-shot action profile. Since $BR_{i,\sigma}(s_{-i}, V_i)$ admits highest and lowest elements a_i^H and a_i^L , the Markovian strategies s_i^H and s_i^L respectively associated with the one-shot profiles a_i^H and a_i^L are the highest and lowest elements of $BR_{i,\sigma}(s_{-i})$ with respect to \preceq . ■

Proof of Lemma 2: Point (i) is an application of Proposition 1 of Chassang (2008). The proof of (ii) is given for the greatest one-shot best-reply $BR_{i,\sigma}^H$. Player i chooses S over E whenever $\Pi_S^i(V_i) \geq \Pi_E^i$. As equation (11) shows, $\Pi_S^i(V_i)$ is increasing in V_i while Π_E^i does not depend on V_i . This yields that $BR_{i,\sigma}^H(a_{-i}, V) \preceq BR_{i,\sigma}^H(a_{-i}, V')$. The same proof applies for the lowest one-shot best-reply. ■

Proof of Lemma 3: Consider $s' \in BR_{i,\sigma}(s_x)$, and denote by V the value player i expects from best-responding. The one-shot action profile a' induced by s' must belong to $BR_{i,\sigma}(s_x, V)$. Proposition 2 of Chassang (2008) implies that there exists $\bar{\sigma}$ such that for all $\sigma \in (0, \bar{\sigma})$ and all x , there is a unique such one-shot best-reply. It takes a threshold form $a_{x'}$

and the threshold x' is continuous in both x and V . This concludes the proof. ■

Proof of Theorem 1: Given Proposition 1 and Lemma 3, the methodology of Milgrom and Roberts (1990) and Vives (1990) applies directly. Simply note that the strategies corresponding to “always staying” and “always exiting” are Markovian threshold-form strategies and apply the best-reply correspondence iteratively. ■

Proof of Lemma 4: Consider the highest equilibrium s_σ^H . For any rationalizable strategy s_{-i} , $s_{-i} \preceq s_{-i}^H$. Assumption 5 implies that player i gets a higher value from best-replying against $s_{-i,\sigma}^H$ than s_{-i} . Thus $V_i \leq V_{i,\sigma}^H$ in the functional sense. A similar argument yields the other inequality. ■

A.2 Proofs for Section 4.3

The proof follows the structure of the analysis given in Section 3.3. The first lemma establishes that global games selection holds uniformly over families of one-shot global games augmented with continuation values.

Lemma A.1 (uniform selection) *For any compact subset $\mathcal{V} \subset \mathbb{R}^2$, consider the family of one-shot global games $\Psi_\sigma(\mathbf{V})$ indexed by $\mathbf{V} \in \mathcal{V}$. If Assumptions 2 and 3 hold, so that for all $\mathbf{V} \in \mathcal{V}$, the full information one-shot game $G(\mathbf{V}, w)$ has pure equilibria which are all symmetric, and admits dominance regions with respect to w , then*

- (i) *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, all one-shot global games $\Psi_\sigma(\mathbf{V})$, indexed by values $\mathbf{V} \in \mathcal{V}$, have a unique rationalizable equilibrium;*
- (ii) *This equilibrium takes a threshold form with thresholds denoted by $\mathbf{x}_\sigma^*(\mathbf{V}) \in \mathbb{R}^2$. The mapping $\mathbf{x}_\sigma^*(\cdot)$ is continuous over \mathcal{V} ;*
- (iii) *As σ goes to 0, each component of $\mathbf{x}_\sigma^*(\mathbf{V})$ ($\in \mathbb{R}^2$) converges uniformly over $\mathbf{V} \in \mathcal{V}$ to the risk-dominance threshold of $\Psi_0(\mathbf{V})$, denoted by $x^{RD}(\mathbf{V})$ ($\in \mathbb{R}$).*

Proof of Lemma A.1: This is a direct application of Theorems 2, 3 and 4 of Chassang (2008). ■

Proof of Theorem 2: For any fixed σ , any Markovian equilibrium of Γ_σ is associated with a vector of constant continuation values $\mathbf{V}_\sigma = (V_{i,\sigma}, V_{-i,\sigma})$. By continuity of the min-max

values, for any $\nu > 0$, there exists $\bar{\sigma} > 0$, such that for all $\sigma \in (0, \bar{\sigma})$, $V_{i,\sigma} \in [m_i - \nu, M_i]$. Stationarity implies that equilibrium actions at any time t must form a Nash equilibrium of the one-shot game

$$\begin{array}{c|cc} & S & E \\ \hline S & g^i(w_t) + \beta V_{i,\sigma} & W_{12}^i(w_t) \\ E & W_{21}^i(w_t) & W_{22}^i(w_t) \end{array}$$

where i is the row player and players get signals $x_{i,t} = w_t + \sigma \varepsilon_{i,t}$. All such one-shot games $\Psi_\sigma(\mathbf{V})$, indexed by $\mathbf{V} \in [m_i - \nu, M_i] \times [m_{-i} - \nu, M_{-i}]$ and $\sigma > 0$ have a global game structure à la Carlsson and van Damme (1993).

Assumption 4 implies that there exists $\nu > 0$ such that for all $\mathbf{V} \in [m_i - \nu, M_i] \times [m_{-i} - \nu, M_{-i}]$ and all $w \in I$, the one-shot game $G(\mathbf{V}, w)$ admits pure equilibria and they are all symmetric. Hence, Lemma A.1 implies that the following are true

1. There exists $\bar{\sigma}$ such that for all $\sigma \in (0, \bar{\sigma})$ and $V \in [m_i - \nu, M_i] \times [m_{-i} - \nu, M_{-i}]$, the game $\Psi_\sigma(\mathbf{V})$ has a unique pair of rationalizable strategies. These strategies take a threshold-form and the associated pair of thresholds is denoted by $\mathbf{x}_\sigma^*(\mathbf{V})$;
2. The pair of thresholds $\mathbf{x}_\sigma^*(\mathbf{V})$ is continuous in \mathbf{V} ;
3. As σ goes to 0, $\mathbf{x}_\sigma^*(\mathbf{V})$ converges to the risk dominant threshold $x^{RD}(\mathbf{V})$ uniformly over $\mathbf{V} \in [m_i - \nu, M_i] \times [m_{-i} - \nu, M_{-i}]$.

The first result, joint selection, implies that there is a unique vector of expected values from playing game $\Psi_\sigma(\mathbf{V})$, which we denote $\phi_\sigma(\mathbf{V})$. The other two results imply that $\phi_\sigma(\mathbf{V})$ is continuous in \mathbf{V} , and that as σ goes to 0, $\phi_\sigma(\mathbf{V})$ converges uniformly over $\mathbf{V} \in \times_{i \in \{1,2\}} [m_i - \nu, M_i]$ to the vector of values $\Phi(\mathbf{V})$ players expect from using the risk-dominant strategy under full information.

Stationarity implies that the value vector \mathbf{V} of any Markovian equilibrium of Γ_σ must satisfy the fixed point equation $\mathbf{V} = \phi_\sigma(\mathbf{V})$. Conversely, any vector of values \mathbf{V} satisfying $\mathbf{V} = \phi_\sigma(\mathbf{V})$ is supported by the Markovian equilibrium in which players play the unique equilibrium of game $\Psi_\sigma(\mathbf{V})$ each period. This gives us (ii).

Furthermore, we know that the equilibrium strategies of game $\Psi_\sigma(\mathbf{V})$ converge to the risk-dominant strategy as σ goes to 0. This allows us to compute explicitly the limit function Φ . Because the risk-dominance threshold is decreasing in the continuation value, and using Assumption 5, it follows that Φ is increasing in \mathbf{V} . This proves (iii).

Finally, (i) is a straightforward implication of (ii). Values associated with Markovian equilibria of Γ_σ are the fixed points of $\phi_\sigma(\cdot)$. Hence the highest and lowest values associated

with Markovian equilibria are also the highest and lowest fixed points of $\phi_\sigma(\cdot)$. ■

Proof of Proposition 2: First, for any $\mu > 0$, there exists $N > 0$ such that for any $n \geq N$, $\mathbf{V}^{L,*} - \mu \leq \mathbf{V}_n^L \leq \mathbf{V}_n^H \leq \mathbf{V}^{H,*} + \mu$.

Let us show point (i). Assume that exit is risk-dominant in $G(\mathbf{V}^{H,*}, w_0)$. This means that there exists $\tau > 0$ such that $x^{RD}(\mathbf{V}^{H,*}) > w_0 + \tau$. By continuity of x^{RD} this implies that there exists $\mu > 0$ such that for all \mathbf{V} satisfying $\mathbf{V} < \mathbf{V}^{H,*} + \mu$, we have $x^{RD}(\mathbf{V}) > w_0 + \tau/2$. This and the fact that f_n converges to a Dirac mass at w_0 implies that there exists N such that for all $n \geq N$, for all $\mathbf{V} \in [\mathbf{V}^{L,*} - \mu, \mathbf{V}^{H,*} + \mu]$, $\Phi_n(\mathbf{V}) < \mathbf{V}^{L,*} + \mu$. By taking μ arbitrarily small, it follows that \mathbf{V}_n^H converges to $\mathbf{V}^{L,*}$ as n goes to infinity. Similar proofs hold for points (ii) and (iii). ■

A.3 On the convergence of fixed points of ϕ_σ

The uniform convergence of ϕ_σ to Φ is useful only to the extent that it implies that the fixed points of ϕ_σ converge to the fixed points Φ . Proposition A.1 shows that uniform convergence of ϕ_σ to Φ implies that fixed points of ϕ_σ necessarily converge to a subset of fixed points of Φ as σ goes to 0. This corresponds to the upper-hemicontinuity of fixed points of ϕ_σ at $\sigma = 0$. Proposition A.2 shows that under generic conditions, any fixed point of Φ is the limit of a sequence $(\mathbf{V}_\sigma)_{\sigma>0}$ of fixed points of ϕ_σ . This corresponds to the lower-hemicontinuity of fixed points of ϕ_σ as σ goes to 0.

Proposition A.1 (upper-hemicontinuity) *The set of fixed points of ϕ_σ is upper hemicontinuous at $\sigma = 0$. For any sequence of positive numbers $\{\sigma_n\}_{n \in \mathbb{N}}$ converging to 0, if $\{\mathbf{V}_n\}_{n \in \mathbb{N}} \equiv \{(V_{i,n}, V_{-i,n})\}_{n \in \mathbb{N}}$ is a sequence of fixed points of ϕ_{σ_n} converging to a pair of values \mathbf{V} , then \mathbf{V} is a fixed point of Φ .*

Proof of Proposition A.1: Since \mathbf{V}_n converges to \mathbf{V} and Φ is continuous, for all $\tau > 0$, there exists N_1 such that for all $n \geq N_1$

$$\|\Phi(\mathbf{V}) - \mathbf{V}\|_{\text{sup}} \leq \|\Phi(\mathbf{V}_n) - \mathbf{V}_n\|_{\text{sup}} + \tau/2.$$

Since $\phi_{\sigma_n}(\cdot)$ converges uniformly to Φ and \mathbf{V}_n is a fixed point of ϕ_{σ_n} , there exists N_2 such that for all $n \geq N_2$, $\|\Phi(\mathbf{V}_n) - \mathbf{V}_n\|_{\text{sup}} \leq \tau/2$. This yields that $\|\Phi(\mathbf{V}) - \mathbf{V}\|_{\text{sup}} \leq \tau$ for all $\tau > 0$. Hence, \mathbf{V} must be a fixed point of Φ . ■

Let us now turn to the question of whether or not all fixed points of Φ correspond to fixed points of ϕ_σ for σ small. So far Markovian equilibria have been characterized by their values. Now it becomes convenient to characterize Markovian equilibria by their cooperation thresholds. Recall the threshold mapping ξ introduced in Definition 6. A pair \mathbf{V} is a fixed point of Φ if and only if $x^{RD}(\mathbf{V})$ is a fixed point of ξ .

Definition A.1 (non-singular fixed points) *A fixed point x of ξ is non-singular if and only if there exists $\epsilon > 0$ such that either*

$$\begin{aligned} & \forall y \in [x - \epsilon, x), \xi(y) < y \quad \text{and} \quad \forall y \in (x, x + \epsilon], \xi(y) > y \\ \text{or} \quad & \forall y \in [x - \epsilon, x), \xi(y) > y \quad \text{and} \quad \forall y \in (x, x + \epsilon], \xi(y) < y. \end{aligned}$$

In other terms, x is non-singular whenever ξ cuts strictly through the 45° line at x .

Proposition A.2 (lower hemicontinuity) *Consider x , a non-singular fixed point of ξ . For any $\sigma > 0$ small enough, there exists a threshold-form Markovian equilibrium of Γ_σ with threshold \mathbf{x}_σ such that $(\mathbf{x}_\sigma)_{\sigma > 0}$ converges to (x, x) as σ goes to 0.*

Hence, generically, all fixed points of ξ and Φ are associated with Markovian equilibria of Γ_σ for σ small.

Proof of Proposition A.2: The proof uses Theorem 3 on local dominance solvability, proven in Section 4.4. For any $x \in \mathbb{R}$, by Lemma 3, $BR_{i,\sigma} \circ BR_{-i,\sigma}(s_x)$ takes a threshold form, $s_{x'}$. Define $\chi_\sigma(\cdot)$ by $\chi_\sigma(x) = x'$. For σ small enough, Lemma 3 and Proposition 1 imply that χ_σ is continuous and increasing. By definition of χ_σ , s_x is a threshold form Markovian equilibrium of Γ_σ if and only if $\chi_\sigma(x) = x$. Consider a non-singular fixed point of ξ denoted by x . There are two cases: x is either a stable or an unstable fixed point of ξ .

Assume that x is a stable fixed point – i.e. ξ cuts the 45° line from above – then Theorem 3 implies that, for all $\bar{\eta} > 0$, there exists $\bar{\sigma} > 0$ and $\eta \in (0, \bar{\eta})$ such that for all $\sigma \in (0, \bar{\sigma})$, the interval $[x - \eta, x + \eta]$ is stable by χ_σ . Since χ_σ is continuous and increasing, this implies that it has a fixed point belonging to $[x - \eta, x + \eta]$. This proves the lower hemicontinuity of stable fixed points of ξ .

Assume that x is unstable. Then for any $\bar{\eta} > 0$, there exists $\eta \in (0, \bar{\eta})$ such that $x - \eta$ and $x + \eta$ respectively belong to the basins of attraction of a lower and a higher fixed point of ξ . Lemma A.3 implies that there exist η' and η'' in $(0, \eta)$ such that $\chi_\sigma(x - \eta') < x - \eta'$ and $\chi_\sigma(x + \eta'') > x + \eta''$. Since χ_σ is continuous, this implies that it admits a fixed point

within $[x - \eta', x + \eta'']$. This proves the lower hemicontinuity of unstable non-singular fixed points of ξ . ■

A.4 Proofs for Section 4.4

The proof of Theorem 3 is broken down in multiple steps. Lemma A.2 shows that the best-reply correspondence does not deviate from the identity mapping around fixed points of ξ . Lemma A.3 is the main step of the proof. It shows that whenever x is a stable fixed point of ξ , then for σ small enough, the first step of iterated best-response shrinks neighborhoods of s_x .

Lemma A.2 *Consider x , a fixed point of ξ . Then there exists $\eta > 0$ and $\bar{\sigma} > 0$ such that for all $\sigma \in [0, \bar{\sigma})$, $x' \in [x - \eta, x + \eta]$, and $i \in \{1, 2\}$ there exists $x'' \in \mathbb{R}$ such that $BR_{i,\sigma}(s'_x) = \{s_{x''}\}$ and $|x'' - x'| < 2\sigma$.*

Proof of Lemma A.2: Since x is a fixed point of ξ , it must be that x is the risk-dominant threshold of the augmented one-shot game $G(BRV_{i,0}(x), BRV_{-i,0}(x), w)$. Hence, at $w = x$, both (E, E) and (S, S) are strict Nash equilibria of this one-shot game. Since $BRV_{i,\sigma}(x')$ is continuous in σ and x' , and payoffs are continuous in w , there exist $\eta > 0$ and $\bar{\sigma} < \eta/4$ such that for all $\sigma \in (0, \bar{\sigma})$ and $x' \in [x - \eta, x + \eta]$, then for all $w \in [x' - \sigma, x' + \sigma]$, both (E, E) and (S, S) are strict Nash equilibria of $G(BRV_{i,\sigma}(x'), BRV_{-i,\sigma}(x'), w)$.

For any $\sigma \in (0, \bar{\sigma})$ and $x' \in [x - \eta/2, x + \eta/2]$, the best-reply to a threshold-form strategy is also a threshold-form strategy. This implies that indeed $BR_{i,\sigma}(x')$ takes the form $s_{x''}$. Let us show that $|x'' - x'| < 2\sigma$. When she gets a signal $x_{i,t} < x' - 2\sigma$, player i knows for sure that player $-i$ will be playing E . From the definition of η , we know that (E, E) is an equilibrium of $G(BRV_{i,\sigma}(x'), BRV_{-i,\sigma}(x'), w)$ for all values of w consistent with a signal value $x_{i,t}$. Thus, it must be that player i 's best-reply is to play E as well. Inversely, when she gets a signal $x_{i,t} > x' + 2\sigma$, player i knows that player $-i$ will play S , and her best-reply is to Stay as well. This implies that $|x'' - x'| < 2\sigma$. ■

Lemma A.3 *Consider a stable fixed point x of ξ and y in the basin of attraction of x . If $y < x$, then there exists $x' \leq y$ and $\bar{\sigma} > 0$ such that x' belongs to the basin of attraction of x and, for all $\sigma \in (0, \bar{\sigma})$ and $i \in \{1, 2\}$, we have $BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x'}) \preceq s_{x'}$.²¹*

²¹Recall that if a and b are thresholds such that $a > b$ then the corresponding strategies satisfy $s_a \preceq s_b$.

Similarly, if $y > x$, there exists $x'' \geq y$ and $\bar{\sigma}$ such that x'' belongs to the basin of attraction of x and for all $\sigma \in (0, \bar{\sigma})$ and $i \in \{1, 2\}$, $s_{x''} \preceq BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x''})$.

Proof of Lemma A.3: Let us prove the first part of the lemma. Define $ba_-(x) = \inf\{\tilde{x} < x \mid \forall y \in [\tilde{x}, x], \xi(y) > y\}$, the infimum of the basin of attraction of x . Because x is stable, its basin of attraction is nonempty and $ba_-(x)$ is well-defined, although it may take value $-\infty$. We distinguish two cases, either $ba_-(x) = -\infty$ or $ba_-(x) \in \mathbb{R}$.

If $ba_-(x) = -\infty$, any $x' < x$ belongs to the basin of attraction of x . Assumption 2 implies that there exists \underline{x} such that for all $\sigma < 1$, $BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{-\infty}) \preceq s_{\underline{x}}$. Pick any $x' < \min\{x, \underline{x}\}$. Using the monotonicity implied by Proposition 1, we conclude that there exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, $BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x'}) \preceq BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{-\infty}) \preceq s_{x'}$.

Consider now the case where $ba_-(x) \in \mathbb{R}$. By continuity of ξ , we have that $\xi(ba_-(x)) = ba_-(x)$. From Lemma A.2 we know that there exist $\eta > 0$ and $\bar{\sigma}$ such that for all $x' \in [ba_-(x) - \eta, ba_-(x) + \eta]$, and $i \in \{1, 2\}$, $BR_{i,\sigma}(s_{x'}) = s_{x''_i}$ with $|x''_i - x'| < 2\sigma$. By definition, we must have $y > ba_-(x)$. Thus we can pick $x' \in (ba_-(x), ba_-(x) + \eta)$ such that $x' < \min\{x, y\}$. We have that $\xi(x') > x'$. By continuity of ξ there exists \tilde{x}' such that $\tilde{x}' < x'$ and $\xi(\tilde{x}') > x'$. To reduce confusion, we temporarily use the notation $BR_{i,\sigma}^{os}(a, V)$ to denote the best-reply of player i to a one shot action profile a and continuation value V . Using the fact that one-shot action profiles are identical to Markovian strategies, we obtain,

$$(13) \quad BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x'}) = BR_{i,\sigma}^{os}(BR_{-i,\sigma}^{os}(s_{x'}, BRV_{-i,\sigma}(x')) , BRV_{i,\sigma}(BR_{-i,\sigma}(s_{x'})))$$

We know that $|x''_i - x'| \leq 2\sigma$. Thus there exists $\bar{\sigma}$ small enough such that $BR_{-i,\sigma}(s_{x'}) \preceq s_{\tilde{x}'}$. Joint with Assumption 5, this implies that, $BRV_{i,\sigma}(BR_{-i,\sigma}(s_{x'})) \leq BRV_{i,\sigma}(\tilde{x}')$. Furthermore, $\tilde{x}' < x'$ implies that $BRV_{i,\sigma}(x') \leq BRV_{i,\sigma}(\tilde{x}')$. Hence, using inequality (13), and the fact that for $i \in \{1, 2\}$, $BR_{i,\sigma}^{os}(a, V)$ is increasing in a and V with respect to \preceq , we obtain

$$(14) \quad \begin{aligned} BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x'}) &\preceq BR_{i,\sigma}^{os}(BR_{-i,\sigma}^{os}(s_{x'}, BRV_{-i,\sigma}(\tilde{x}')), BRV_{i,\sigma}(\tilde{x}')) \\ &\preceq BR_{i,\sigma}^{os}(\cdot, BRV_{i,\sigma}(\tilde{x}')) \circ BR_{-i,\sigma}^{os}(\cdot, BRV_{-i,\sigma}(\tilde{x}')) (s_{x'}) \end{aligned}$$

Lemma A.1 implies that there exists $\bar{\sigma}$ small enough such that for all $\sigma \in (0, \bar{\sigma})$ and all $(V_i, V_{-i}) \in [m_i, M_i] \times [m_{-i}, M_{-i}]$, the game $\Psi_\sigma(V_i, V_{-i})$ has a unique rationalizable pair of strategies $\mathbf{x}_\sigma^*(V_i, V_{-i})$. Lemma A.1 also implies that $\mathbf{x}_\sigma^*(V_i, V_{-i})$ converges uniformly to $x^{RD}(V_i, V_{-i})$ as σ goes to 0. This implies that $\mathbf{x}_\sigma^*(BRV_{i,\sigma}(\tilde{x}'), BRV_{-i,\sigma}(\tilde{x}'))$ converges to $(\xi(x'), \xi(x'))$ as σ goes to 0. Since $x' < \xi(\tilde{x}')$, it implies there exists $\bar{\sigma}$ such that for all

$\sigma \in (0, \bar{\sigma})$, $x' < x_\sigma^*(BRV_{i,\sigma}(\tilde{x}'), BRV_{-i,\sigma}(\tilde{x}'))$.

The fact that $\Psi_\sigma(BRV_{i,\sigma}(\tilde{x}'), BRV_{-i,\sigma}(\tilde{x}'))$ has a unique rationalizable strategy and the monotonicity property of Proposition 2 imply that the sequence of threshold-form strategies

$$(BR_{i,\sigma}^{os}(\cdot, BRV_{i,\sigma}(\tilde{x}')) \circ BR_{-i,\sigma}^{os}(\cdot, BRV_{-i,\sigma}(\tilde{x}')))^n(s_{x'}), \quad \text{for } n \in \mathbb{N},$$

converges monotonically to the Markovian equilibrium of threshold $x_\sigma^*(BRV_{i,\sigma}(x'), BRV_{-i,\sigma}(x'))$. Since $x' < x_\sigma^*(BRV_{i,\sigma}(\tilde{x}'), BRV_{-i,\sigma}(\tilde{x}'))$, the sequence must be decreasing with respect to the order \preceq on strategies. Thus $BR_{i,\sigma}^{os}(\cdot, BRV_{i,\sigma}(\tilde{x}')) \circ BR_{-i,\sigma}^{os}(\cdot, BRV_{-i,\sigma}(\tilde{x}'))(s_{x'}) \preceq s_{x'}$. Using inequality (14), this yields that indeed $BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x'}) \preceq s_{x'}$.

The second part of the lemma results from a symmetric reasoning, switching all inequalities. ■

Proof of Theorem 3: Using Lemma A.3, we know there exist $\bar{\sigma}$, $x_- \leq y$ and $x_+ \geq z$, with $x \in (x_-, x_+)$ and $[x_-, x_+]$ included in the basin of attraction of x , such that for all $\sigma \in (0, \bar{\sigma})$, and $i \in \{1, 2\}$,

$$BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x-}) \preceq s_{x-} \quad \text{and} \quad s_{x+} \preceq BR_{i,\sigma} \circ BR_{-i,\sigma}(s_{x+}).$$

These inequalities and Proposition 1 imply by iteration that for all $n \in \mathbb{N}$,

$$\begin{aligned} (BR_{i,\sigma}^\Delta \circ BR_{-i,\sigma}^\Delta)^n([s_{x+}, s_{x-}]) &\subset [(BR_{i,\sigma} \circ BR_{-i,\sigma})^n(s_{x+}), (BR_{i,\sigma} \circ BR_{-i,\sigma})^n(s_{x-})] \\ &\subset [(BR_{i,\sigma} \circ BR_{-i,\sigma})^{n-1}(s_{x+}), (BR_{i,\sigma} \circ BR_{-i,\sigma})^{n-1}(s_{x-})] \\ &\subset \cdots \subset [s_{x+}, s_{x-}]. \end{aligned}$$

Consider the decreasing sequence $\{(BR_{i,\sigma} \circ BR_{-i,\sigma})^n(s_{x-})\}_{n \in \mathbb{N}}$. As n goes to ∞ , it must converge to a threshold form strategy with threshold $x_{i,\sigma}^- \in [x_-, x_+]$. Moreover $(s_{x_{i,\sigma}^-}, BR_{-i,\sigma}(s_{x_{i,\sigma}^-}))$ must be a Markovian threshold-form equilibrium of Γ_σ . Lemma A.1 implies that as σ goes to 0, any converging subsequence of $\{(x_{i,\sigma}^-, x_{-i,\sigma}^-)\}_{\sigma > 0}$ must converge to a symmetric pair (\tilde{x}, \tilde{x}) such that \tilde{x} is a fixed point of ξ and $\tilde{x} \in [x_-, x_+]$. The only fixed point of ξ in $[x_-, x_+]$ is x . This implies that as σ goes to 0, $x_{i,\sigma}^-$ must converge to x . Similarly, as n goes to ∞ , the sequence $(BR_{i,\sigma} \circ BR_{-i,\sigma})^n(s_{x+})$ converges to a threshold strategy with a threshold $x_{i,\sigma}^+$ that converges to x as σ goes to 0. This concludes the proof. ■

References

- [1] Abreu, Dilip, David Pearce, and Ennio Stacchetti, *Optimal Cartel Equilibria with Imperfect Monitoring*, Journal of Economic Theory, Vol. 39, No. 1, pp. 251-269, 1986.
- [2] Abreu, Dilip, David Pearce and Ennio Stacchetti, *Toward a Theory of Discounted Repeated Games with Imperfect Monitoring*, Econometrica, Vol. 58, No. 5., pp. 1041-1063. 1990.
- [3] Angeletos, Marios, Christian Hellwig and Alessandro Pavan, *Dynamic Global Games of Regime Change: Learning, Multiplicity, and Timing of Attacks*, Econometrica, Vol. 75, No. 3, pp. 711-756, 2007.
- [4] Athey, Susan, *Monotone Comparative Statics Under Uncertainty*, Quarterly Journal of Economics, Vol. 117, No. 1, pp. 187-223, 2002.
- [5] Baliga, Sandeep and Stephen Morris, *Coordination, Spillovers, and Cheap Talk*, Journal of Economic Theory, Vol. 105, No. 2, pp. 450-468, 2002.
- [6] Burdzy, Krzysztof, David Frankel, and Ady Pauzner, *Fast Equilibrium Selection by Rational Players Living in a Changing World*, Econometrica, Vol. 69, No. 1, pp. 163-189, 2001.
- [7] Carlsson, Hans and Eric van Damme, *Global Games and Equilibrium Selection*, Econometrica, Vol. 61, No. 5, pp. 989-1018, 1993.
- [8] Chamley, Christophe, *Coordinating Regime Switches*, Quarterly Journal of Economics, Vol. 114, No. 3, pp. 869-905, 1999.
- [9] Chassang, Sylvain, *Uniform Selection in Global Games*, Journal of Economic Theory, Vol. 139, No. 1, pp. 222-241, 2008
- [10] Chassang, Sylvain, and Gerard Padro i Miquel, *Conflict and Deterrence under Strategic Risk*, NBER Working Paper No. W13964, 2008.
- [11] Chassang, Sylvain, and Satoru Takahashi, *Robustness to Incomplete Information in Repeated Games*, Unpublished Manuscript, Princeton University, 2009.
- [12] Echenique, Federico, *Extensive Form Complementarities*, Games and Economic Behavior, Vol. 46, No. 2, pp. 348-364, 2004.

- [13] Frankel, David, Stephen Morris, and Ady Pauzner, *Equilibrium Selection in Global Games with Strategic Complementarities*, Journal of Economic Theory, Vol. 108, No. 1, pp. 1-44, 2003.
- [14] Frankel, David, and Ady Pauzner, *Resolving Indeterminacy in Dynamic Settings: the Role of Shocks*, Quarterly Journal of Economics, No. 115, pp. 283-304, 2000.
- [15] Fudenberg, Drew, David Levine, and Eric Maskin, *The Folk Theorem with Imperfect Public Information*, Econometrica, Vol. 62, No. 5., pp. 997-1039, 1994.
- [16] Green, Edward, and Robert Porter, *Noncooperative Collusion under Imperfect Price Information*, Econometrica, Vol. 52, No. 1, pp. 87-100, 1984.
- [17] Giannitsarou, Chrissy, and Flavio Toxvaerd, *Recursive Global Games*, CEPR Discussion Paper No. DP6470, 2007.
- [18] Guesnerie, Roger, *Anchoring Economic Predictions in Common Knowledge*, Vol. 70, No. 2, pp. 439-480, 2002.
- [19] Harsanyi, John and Reinhard Selten, *A General Theory of Equilibrium Selection in Games*, MIT Press, 1988.
- [20] Kajii, Atsushi and Stephen Morris, *The Robustness of Equilibria to Incomplete Information*, Econometrica, Vol 65, No. 6, pp. 1283-1309, 1997.
- [21] Levin, Jonathan, *A Note on Global Games with Overlapping Generations*, Stanford university Mimeo, 2001.
- [22] Mailath, George, and Stephen Morris, *Repeated Games with Almost Public Monitoring*, Journal of Economic Theory, Vol. 102, pp. 189-228, 2002.
- [23] Mailath, George, and Stephen Morris, *Coordination Failure in Repeated Games with Almost-Public Monitoring*, Theoretical Economics, Vol. 1, No. 3, pp. 311-340, 2006.
- [24] Milgrom, Paul, and John Roberts, *Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities*, Econometrica, Vol. 58, No. 6, pp. 1255-1277, 1990.
- [25] Milgrom, Paul, and Chris Shannon, *Monotone Comparative Statics*, Econometrica, Vol. 62, No. 1, pp. 157-180, 1994.

- [26] Morris, Stephen and Hyun Song Shin, *A Theory of the Onset of Currency Attacks*, in Agenor, Vines and Weber, eds., *Asian Financial Crisis: Causes, Contagion and Consequences*, Cambridge, UK, Cambridge University Press, 1999.
- [27] Moulin, Herve, *Dominance Solvability and Cournot Stability*, Mathematical Social Science, Vol. 7, No. 1, pp. 83-102, 1984.
- [28] Ordoñez, Guillermo, *Fragility of Reputation and Clustering in Risk-Taking*, UCLA mimeo, 2008.
- [29] Radner, Roy, Roger Myerson, and Eric Maskin, *An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria*, Review of Economic Studies, Vol. 53, No. 1, pp. 59-69, 1986.
- [30] Steiner, Jakub, *Coordination of Mobile Labor*, Journal of Economic Theory, Vol. 139, No. 1, pp. 25-46, 2008.
- [31] Vives, Xavier, *Nash Equilibrium with Strategic Complementarities*, Journal of Mathematical Economics, No. 19, pp. 305-321, 1990.

B “Fear of Miscoordination and the Robustness of Cooperation in Dynamic Global Games with Exit”

Supplementary Material: Online Appendix

Abstract

This appendix provides additional results on equilibria of exit game Γ_σ as the noise parameter σ becomes small. Section B.1 shows that all Markovian equilibria will be approximately symmetric for σ small. Section B.2 studies the structure of time-dependent Markovian equilibria. Finally, Section B.3 provides sufficient conditions under which game Γ_σ is asymptotically dominance solvable.

B.1 Asymptotically Symmetric Play in Markovian Equilibria

Theorem 1 establishes that for noise parameter σ small enough, the set of rationalizable strategies of Γ_σ is bounded by a most cooperative and a least cooperative equilibria, that are Markovian and take a threshold form.

The following proposition shows that for σ small enough, *all* Markovian equilibria take a threshold form and are asymptotically symmetric.

Proposition B.1 *There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, every Markovian equilibrium \mathbf{s} of Γ_σ takes a threshold form with thresholds $(x_{i,\sigma}, x_{-i,\sigma})$. Furthermore, we have that $|x_{i,\sigma} - x_{-i,\sigma}| < 2\sigma$.*

Proof: The fact that for σ small enough, all Markovian equilibria take a threshold form is a direct consequence of Lemma A.1, applied to the class of one-shot games augmented with the continuation values associated to Markovian equilibria.

We now show that thresholds $(x_{i,\sigma}, x_{-i,\sigma})$ satisfy $|x_{i,\sigma} - x_{-i,\sigma}|$. This follows from the fact that given an equilibrium threshold $x_{-i,\sigma}$, then player i knows that if $x_{i,t} < x_{-i,\sigma} - \sigma$ player $-i$ will play E , so that player i 's best reply is to play E as well. Similarly, if $x_{i,t} > x_{-i,\sigma} + \sigma$, then player i knows that player $-i$ will play S and her best reply is to play S as well. ■

This result highlights that asymptotically, the likelihood of actual miscoordination is vanishing. Note that since the players' payoffs may be asymmetric, the continuation values associated with approximately symmetric Markovian equilibria may be quite different.

In games where payoffs are symmetric, and error terms ε_i and ε_{-i} have identical distributions, Proposition B.1 can be strengthened to show that all Markovian equilibria are symmetric.

Proposition B.2 *Whenever payoff functions are symmetric, and error terms $\varepsilon_{i,t}$ and $\varepsilon_{-i,t}$ have identical distributions, there exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, every Markovian equilibrium takes a threshold form with thresholds $x_{i,\sigma} = x_{-i,\sigma}$.*

Proof: Consider σ small enough that all Markovian equilibria take a threshold form. Consider such a Markovian equilibrium with thresholds $(x_{i,\sigma}, x_{-i,\sigma})$, associated with values $(V_{i,\sigma}, V_{-i,\sigma})$. The proof proceeds by contradiction. Assume for instance that $x_{i,\sigma} > x_{-i,\sigma}$. Because of Assumption 5, that staying benefits one's partner, it follows that $V_{i,\sigma} > V_{-i,\sigma}$. Given that player i is indifferent between staying and exiting at signal $x_{i,t} = x_{i,\sigma}$ and that $V_{-i,\sigma} < V_{i,\sigma}$, player $-i$ must strictly prefer exiting to staying when observing signal $x_{-i,t} = x_{i,\sigma}$. This contradicts $x_{i,\sigma} > x_{-i,\sigma}$, and implies that $x_{i,\sigma} = x_{-i,\sigma}$. ■

B.2 Time-Dependent Markovian Equilibria

Section 4 used a dynamic programming approach à la Abreu, Pearce and Stacchetti (1990) to characterize Markovian equilibria of Γ_σ . The same approach can be used to characterize time-dependent Markovian equilibria.

Definition B.1 *A strategy s_i is time-dependent Markovian if and only if $s_i(h_{i,t})$ depends only on time t and player i 's current signal $x_{i,t}$.*

For σ small enough, and for any pair of values $\mathbf{V} \in \prod_{i \in \{1,2\}} [m_i, M_i]$, we consider the mappings $\mathbf{x}_\sigma^*(\mathbf{V})$, and $\phi_\sigma(\mathbf{V})$ defined in Appendix A.2. Recall that $\mathbf{x}_\sigma^*(\mathbf{V})$ is the unique equilibrium threshold of the augmented one-shot global game $\Psi_\sigma(\mathbf{V})$ and $\phi_\sigma(\mathbf{V})$ is the value of playing $\Psi_\sigma(\mathbf{V})$ according to its unique equilibrium threshold.

A profile of time-dependent Markovian strategies $\mathbf{s} = (s_i, s_{-i})$ is associated with the sequence of values $(\mathbf{V}_t)_{t \in \mathbb{N}} = (V_{i,t}, V_{-i,t})_{t \in \mathbb{N}}$, where \mathbf{V}_t is the pair of values associated with playing according to strategies \mathbf{s} in the subgame starting at date t . A sequence of values $(\mathbf{V}_t)_{t \in \mathbb{N}}$ is supported by a time-dependent Markovian equilibrium of Γ_σ if and only if, the sequence $(\mathbf{V}_t)_{t \in \mathbb{N}}$ is bounded and satisfies the recurrence equation $\mathbf{V}_t = \phi_\sigma(\mathbf{V}_{t+1})$ for all $t \in \mathbb{N}$. Furthermore such a sequence of continuation values is sustained by a unique perfect Bayesian equilibrium such that players choose to stay in period t according to threshold $\mathbf{x}_\sigma^*(\mathbf{V}_t)$. The proof of these results is straightforward, and essentially identical to that of Theorem 2.

In order to say more about time dependent Markovian equilibria, the rest of the section focuses on symmetric games and symmetric equilibria. Mapping Φ can be reduced to a mapping from \mathbb{R} to \mathbb{R} . Denote by $\mathcal{US}(\Phi)$ the set of unstable fixed points of mapping Φ and $\mathcal{S}(\Phi)$ the set of stable fixed points of mapping Φ . The analysis assumes that all the fixed points of Φ are non-singular.

Consider \mathbf{s} a symmetric time-dependent Markovian equilibrium and $(V_t)_{t \in \mathbb{N}}$ the associated sequence of values. Since V_t must be bounded, we can always extract a converging sequence converging to some value $V_{\sigma, \infty}$.

Proposition B.3 *Pick any $\eta > 0$. There exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$, the following hold*

(i) *If $V_{\sigma, \infty} \in \mathbb{R} \setminus \bigcup_{V \in \mathcal{US}(\Phi)} [V - \eta, V + \eta]$ then there exists $V^* \in \mathcal{S}(\Phi)$ such that for all $t \in \mathbb{N}$, $V_t \in [V^* - \eta, V^* + \eta]$.*

(ii) *If $V_{\sigma, \infty} \in \mathbb{R} \setminus \bigcup_{V \in \mathcal{S}(\Phi)} [V - \eta, V + \eta]$ then there exists $V^* \in \mathcal{US}(\Phi)$ and $T > 0$ such that for all $t \geq T$, $V_t \in [V^* - \eta, V^* + \eta]$.*

Proof: The fixed points of Φ belong to some compact interval $[m, M]$. Since by assumption every fixed point of Φ is non-singular, this means that there are only finitely many of them.

Furthermore, since Φ is increasing and all its fixed points are non singular, then for every $\zeta > 0$, there exists $k \in \mathbb{N}$ and $\nu \in (0, \zeta)$ such that:

- for all $V \in [m, M] \setminus \bigcup_{V \in \mathcal{US}(\Phi)} [V - \zeta, V + \zeta]$, $\Phi^k(V) \in \bigcup_{V \in \mathcal{S}(\Phi)} [V - \zeta + \nu, V + \zeta - \nu]$;

- for all $V^* \in \mathcal{S}(\Phi)$, $\Phi([V^* - \zeta, V^* + \zeta]) \subset [V^* - \zeta + \nu, V^* + \zeta - \nu]$.

Since ϕ_σ converges uniformly to Φ as σ goes to 0, there exists $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$,

(a) for all $V \in [m, M] \setminus \bigcup_{V \in \mathcal{US}(\Phi)} [V - \zeta, V + \zeta]$, $\phi_\sigma^k(V) \in \bigcup_{V \in \mathcal{S}(\Phi)} [V - \zeta, V + \zeta]$;

(b) for all $V^* \in \mathcal{S}(\Phi)$, $\phi_\sigma([V^* - \zeta, V^* + \zeta]) \subset [V^* - \zeta, V^* + \zeta]$.

This implies Proposition B.3 (i). Indeed, pick $\eta > 0$ and apply (a) and (b) above with $\zeta < \eta$. Since $V_{\sigma, \infty} \in [m, M] \setminus \bigcup_{V \in \mathcal{US}(\Phi)} [V - \eta, V + \eta]$, there are infinitely many times $t \in \mathbb{N}$ such that $V_t \in [m, M] \setminus \bigcup_{V \in \mathcal{US}(\Phi)} [V - \zeta, V + \zeta]$. By point (a) above, this implies that there exists $V^* \in \mathcal{S}(\Phi)$ such that there are infinitely many times t at which $V_t \in [V^* - \zeta, V^* + \zeta]$. By point (b) it follows that in every earlier period, and hence in every period s , $V_s \in [V^* - \zeta, V^* + \zeta] \subset [V^* - \eta, V^* + \eta]$. This proves point (i).

We now move to point (ii). Pick the same $\bar{\sigma}$ as in the proof above. The fact that $V_{\sigma, \infty} \in \mathbb{R} \setminus \bigcup_{V \in \mathcal{S}(\Phi)} [V - \eta, V + \eta]$ implies that there exists $T_1 > 0$ large enough such that for all $t > T_1$, $V_t \in \bigcup_{V \in \mathcal{US}(\Phi)} [V - \eta, V + \eta]$. Otherwise we would be in case (i), which implies that $V_{\sigma, \infty}$ should be within a small neighborhood of $\mathcal{S}(\Phi)$. Furthermore, since Φ is increasing, it is not possible to transition from a unstable fixed point of Φ to an other unstable fixed point of Φ without being in the neighborhood of a stable fixed point of Φ . Hence this means that there exists T_2 and $V^* \in \mathcal{US}(\Phi)$ such that for all $t \geq T_2$, $V_t \in [V^* - \eta, V^* + \eta]$. ■

Furthermore, note that if we are in case (ii) of Proposition B.3, then for η small, the continuation equilibrium after time T is in an arbitrarily small neighborhood of a Markovian equilibrium that is asymptotically unstable in the sense developed in Section 4.4. This follows from the fact that unstable fixed points of ϕ are associated to unstable fixed points of ξ .

Altogether this means that a time-dependent Markovian equilibrium is either very close to an asymptotically stable Markovian equilibrium, or is arbitrarily close to an asymptotically unstable Markovian equilibrium sufficiently far away in the future.

B.3 Sufficient Conditions for Uniqueness

Theorem 2 implies that whenever the mapping Φ has a unique fixed point, then the set of rationalizable strategies of Γ_σ converges to a singleton as σ goes to 0. The following proposition provides sufficient conditions under which mapping Φ has a unique fixed point.

Proposition B.4 (uniqueness) *Pick K a compact of \mathbb{R}^2 . There exists a constant $\eta > 0$, depending only on payoff functions and K , such that whenever*

(i) *players have individually rational values for playing game Γ_0 that belong to K ,*

(ii) *the distribution of states of the world f satisfies $\max f < \eta$*

mapping Φ admits a unique fixed point and the set of rationalizable strategies of Γ_σ converges to a singleton as σ goes to 0.

Proof: Let $\|\cdot\|_1$ denote the norm on \mathbb{R}^2 defined by $\|\mathbf{V}\|_1 = |V_i| + |V_{-i}|$, and let $\|\cdot\|_\infty$ denote the sup norm. It results from Theorem 2 (iii) that

$$\|\Phi(\mathbf{V}) - \Phi(\mathbf{V}')\|_1 \leq \beta \|\mathbf{V} - \mathbf{V}'\|_1 + \|f\|_\infty \sum_{i \in \{1,2\}} \|g_{11}^i + \beta V_i - W_{22}^i\|_\infty \left\| \frac{\partial x^{RD}}{\partial V_i} + \frac{\partial x^{RD}}{\partial V_{-i}} \right\|_\infty \|\mathbf{V} - \mathbf{V}'\|_1$$

Since $\sum_{i \in \{1,2\}} \|g_{11}^i + \beta V_i - W_{22}^i\|_\infty \left\| \frac{\partial x^{RD}}{\partial V_i} + \frac{\partial x^{RD}}{\partial V_{-i}} \right\|_\infty$ is finite, for any $\delta \in (\beta, 1)$, there exists $\|f\|_\infty$ small enough such that $\|\Phi(\mathbf{V}) - \Phi(\mathbf{V}')\|_1 \leq \delta \|\mathbf{V} - \mathbf{V}'\|_1$. Hence Φ is a contraction mapping, which concludes the proof. ■

Intuitively, Proposition B.4 implies that when the state of the world w_t has sufficient variance, then game Γ_σ is asymptotically dominance solvable. Indeed, when the density of distribution f becomes arbitrarily small, a given change in cooperation levels induces an arbitrarily small change in continuation values, which is not enough to make the original change in cooperation levels self-sustaining.