Robustness to Incomplete Information in Repeated and Dynamic Games

Sylvain Chassang* Satoru Takahashi†
Princeton University Princeton University

May 30, 2008

Abstract

This paper extends Kajii and Morris (1997)'s notion of robustness to incomplete information to repeated and more generally dynamic games. We argue that in dynamic games, the requirements of robustness should be strengthened to allow for small payoff deviations with large probability. Under this strengthening, we show that dynamically robust equilibria can be characterized by applying a one-shot robustness principle that extends the one-shot deviation principle. For repeated games, this implies a factorization result analogous to that of Abreu, Pearce and Stacchetti (1990).

We then apply these results to characterize explicitly the set of robust equilibria in the repeated Prisoners' Dilemma and show that cooperation is robustly sustainable if and only if both (*Cooperate*, *Cooperate*) and (*Cooperate*, *Defect*) are enforceable under complete information. Robustness considerations also nuance our analysis of the repeated Prisoners' Dilemma in interesting ways. In particular, grim-trigger strategies may be less robust than asymmetric strategies which do not punish both players upon unilateral deviation.

---

*chassang@princeton.edu
†satorut@princeton.edu

1 Introduction

When we use game-theoretic models to analyze real life situations, it is important to assess the robustness of our predictions to small misspecifications of the strategic environment. Kajii and Morris (1997, henceforth KM) consider the problem of an analyst who believes a strategic situation is well approximated by a complete information game with high probability, but knows that players may still face some unspecified uncertainty about each other's payoffs. In that setting, KM make the point that the analyst can be justified in using the baseline complete-information game to make his predictions if and only if it is not significantly different from an equilibrium of the actual incomplete-information game that players are really facing. This defines a notion of robustness of equilibrium to incomplete information.[1] While the question of robustness issues is relevant in many settings, KM's analysis is limited to static games and the question of robustness in dynamic games remains somewhat unstudied. The present paper attempts to fill this gap by providing a variation on KM's notion of robustness that is appropriate for repeated games and extends well to general multistage games with discounting.

Given a complete information game G, KM study the robustness of its equilibria by studying equilibria of incomplete information games U that are elaborations of G in the sense that with high probability, payoffs of U are exactly those of G. Our approach to robustness in repeated games is similar. Given a repeated game Γ_G with complete information stage game G, we study the properties of dynamic games $\Gamma_{\mathbf{U}}$ characterized by a sequence $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ of independent incomplete information stage games, all of which are elaborations of G. However, robustness in dynamic games differs from robustness in static games because perturbations that occur in future periods with some probability, change current discounted expected payoffs with high probability. This leads us to consider a stronger notion of robustness that only requires elaborations U to have payoffs close (rather than identical)

---

[1] A closely related approach is that of global games, which considers a specific perturbation of the original complete-information game in which agents get noisy but very precise signals of a payoff relevant state of the world (Carlsson and van Damme (1993)). See Morris and Shin (2003) for a survey on global games and other related topics.

to those of the original game G with high probability.

Our main theoretical results relate the dynamic robustness of equilibria in repeated games to the robustness of equilibrium action profiles in appropriate static games augmented with continuation values. For repeated games these results extend the dynamic programming approach of Abreu, Pearce and Stacchetti (1990, henceforth APS). In particular, we prove a factorization result for robust subgame-perfect equilibria. More generally, we prove a one-shot robustness principle for multi-stage games with discounting which is analogous to the one-shot deviation principle of Fudenberg and Levine (1983). This allows us to characterize dynamically robust equilibria by considering only one-shot elaborations rather than all dynamic elaborations.

Our main applied results highlight the practical value of these characterizations by explicitly computing the set of robust equilibrium values in the repeated Prisoners' Dilemma. We show that whenever outcome $(\mathit{Cooperate}, \mathit{Defect})$ can be enforced in equilibrium under complete information, the set of robust equilibrium values is essentially equal to the set of equilibrium values under complete information. Inversely, whenever $(\mathit{Cooperate}, \mathit{Defect})$ is not enforceable in equilibrium, the set of robust equilibria shrinks to permanent defection, the stage Nash equilibrium. In addition, we highlight that carefully considering the issue of robustness nuances our intuitions about which equilibria are most suited to sustain cooperation. Specifically, asymmetric equilibria that punish only deviators upon unilateral deviations are more effective than grim-trigger strategies at sustaining robust cooperation. We also use the factorization result to prove a folk theorem in robust equilibria.

The question of robustness in dynamic games has been the focus of much of the refinement literature and it is impossible to fully summarize all the relevant work in this field. Closely related to the approach of this paper is Fudenberg, Kreps and Levine (1988), who ask whether a given equilibrium of an extensive-form game can be approximated by a sequence of strict equilibria of elaborations. Dekel and Fudenberg (1990) extends this question to iterative elimination of weakly dominated strategies. Characterizations in these papers are somewhat permissive, however, partly because they require only the *existence* of an approximating sequence of elaborations, and partly because their definition of elaborations allows for intertemporal and interpersonal correlation of payoff perturbations.

More recently, Bhaskar, Mailath and Morris (2007) study the dynamic robustness of specific equilibria in the repeated Prisoners' Dilemma. They focus on the mixed-strategy equilibrium constructed by Ely and Välimäki (2002), and show that the Ely-Välimäki equilibrium cannot be approximated by one period memory equilibria of generic perturbed games, but can be approximated by infinite memory equilibria. One important difference is that they follow the purification literature à la Harsanyi (1973) and perturb payoffs in stage games independently across players and periods. In contrast, we follow KM and add payoff shocks that are independent across periods but not necessarily across players.

Perhaps most closely related to this paper are Giannitsarou and Toxvaerd (2007) and Chassang (2007), who extend global game approaches to dynamic settings by using dynamic programming methods. Giannitsarou and Toxvaerd (2007) show that a finite-horizon game with strategic complementarities and small payoff perturbations has an essentially unique equilibrium. Chassang (2007) considers infinite-horizon dynamic cooperation games with exit and shows that even though the global games perturbation does not yield uniqueness in such settings, it still selects a subset of equilibria whose qualitative properties are driven by risk-dominance considerations. The main difference between these results and ours is that they consider robustness to a specific information perturbation, whereas we demand robustness to any sequence of independent elaborations. This makes our robustness results stronger, and our non-robustness results weaker.

The rest of the paper is structured as follows. Section 2 describes the framework and defines an extension of KM's notion of robustness. Section 3 defines robustness to incomplete information for repeated games and provides a factorization result that characterizes robust subgame-perfect equilibria. Section 4 applies the results of Section 3 to characterize robust equilibrium strategies in the Repeated Prisoners' dilemma and prove a folk theorem in robust equilibria. Section 5 extends the analysis to multistage games with discounting. Section 6 concludes.

2 Robustness in Static Games

In this section, we define a notion of robustness to incomplete information that strengthens the original notion of robustness of KM. We depart from KM in anticipation of specific difficulties that arise when we analyze robustness in repeated games. Indeed, given a complete information game, KM consider elaborations such that payoffs are identical to those of the complete information game with high probability. In repeated games, the fact that payoffs can perturbed with some small probability in future periods implies that current expected continuation values can be slightly different from original continuation values with large probability. As a result, our notion of robustness allows for elaborations that have payoffs close (instead of identical) to the payoffs from the complete-information game with a large probability. In Section 5 we show that unless we impose such a strengthening of robustness, the dynamic robustness of an equilibrium is not implied by the robustness of each one-shot action profile in the appropriate stage game augmented with continuation values.

2.1 Definition

For any finite set X, let $\Delta(X)$ denote the set of probability distributions over X. For x and $y \in \mathbb{R}^n$, we write $x \gg y$ if and only if $x_i > y_i$ for every $i = 1, \dots, n$.

Consider a complete-information game $G = (N, (A_i, g_i)_{i \in N})$ consisting of a finite set of players $N = \{1, \dots, n\}$. To each player $i \in N$ we associate a finite action set A_i, a payoff function $g_i \colon A \to \mathbb{R}$, where $A = \prod_{i \in N} A_i$ is the set of action profiles. Let $a_{-i} \in A_{-i} = \prod_{j \in N \setminus \{i\}} A_j$ denote an action profile of player i's opponents. We use the max norm: $|g_i| = \max_{a \in A} |g_i(a)|$ and $|g| = \max_{i \in N} |g_i|$. A pure-action profile $a^* = (a_i^*)_{i \in N} \in A$ is: a *Nash equilibrium* if $g_i(a^*) \geq g_i(a_i, a_{-i}^*)$ for every $i \in N$ and $a_i \in A_i$; a *d-strict equilibrium* for $d \geq 0$ if and only if $g_i(a^*) \geq g_i(a_i, a_{-i}^*) + d$ for every $i \in N$ and $a_i \in A_i \setminus \{a_i^*\}$; a *strict equilibrium* if it is a d-strict equilibrium for some $d > 0$.

We also consider an incomplete-information game $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$, where Ω is a countable state space, P is a probability distribution on Ω, and for each player $i \in N$,

$u_i\colon A\times\Omega\to\mathbb{R}$ is his state-dependent payoff function and Q_i is his information partition over Ω. We assume that payoffs are bounded by some large but finite number M. The domain of u_i extends to mixed or correlated strategies in the usual way. We say that U *embeds* G if U has the same sets of players and of actions as G does. Mixed strategies of player i are Q_i-measurable mappings $\alpha_i\colon\Omega\to\Delta(A_i)$. A profile of mixed-strategies $\alpha=(\alpha_i)_{i\in N}$ induces a distribution $P^\alpha\in\Delta(A)$ over action profiles defined by

$$\forall a\in A,\quad P^\alpha(a)=\sum_{\omega\in\Omega}\left(\prod_{i\in N}\alpha_i(\omega)(a_i)\right)P(\omega).$$

A mixed-strategy profile α^* is a *Bayesian-Nash equilibrium* if

$$\sum_{\omega\in\Omega}u_i(\alpha^*(\omega),\omega)P(\omega)\geq\sum_{\omega\in\Omega}u_i(\alpha_i(\omega),\alpha^*_{-i}(\omega),\omega)P(\omega)$$

for all $i\in N$ and all Q_i-measurable strategies α_i.

For $\varepsilon>0$ and $d\geq 0$, we say that U is an (ε,d)*-elaboration* if every player in U knows that his payoff function is within distance d from his payoff function in G with probability larger than $1-\varepsilon$, i.e.,

$$P(\{\omega\in\Omega:|u_i(\cdot,\omega')-g_i|\leq d\text{ for all }i\in N\text{ and }\omega'\in Q_i(\omega)\})>1-\varepsilon.$$

Definition 1 (d-robustness)**.** For $d\geq 0$, a Nash equilibrium a^* of G is *d-robust (to incomplete information)* if, for every $\eta>0$, there exists $\varepsilon>0$ such that every (ε,d)-elaboration U of G has a mixed-strategy Bayesian-Nash equilibrium α^* such that $P^{\alpha^*}(a^*)>1-\eta$.

We say that an equilibrium a^* of G is *strongly robust* if it is d-robust for some $d>0$.

Our interpretation of d-robustness is essentially identical to KM's interpretation of robustness. Consider an analyst who thinks that the true but complicated incomplete-information game U is approximately described by a simple complete-information whose payoffs are close to those of U with high probability. In that setting the analyst can make a prediction about the observed action profile without knowing the details of U if and only if her prediction is a robust equilibrium of G.

Note that, for $d = 0$, d-robustness corresponds to robustness in the sense of KM. For $d > 0$ however, d-robustness is strictly stronger than robustness à la KM. In particular, as the following lemma highlights, strong robustness implies strictness of equilibrium.

Lemma 1 (strictness)**.** *If a^* is a d-robust equilibrium, then it is a $2d$-strict equilibrium.*

Proof. Consider the game $G' = (N, (A_i, g'_i)_{i \in N})$ such that, for every $i \in N$, $g'_i(a) = g_i(a) + d$ for $a \neq a^*$ and $g'_i(a^*) = g_i(a^*) - d$. Since G' is an (ε, d)-elaboration of G for every $\varepsilon > 0$, G admits an equilibrium arbitrarily close to a^*. This implies that a^* is also an equilibrium of G', thus a $2d$-strict equilibrium of G. □

This lemma implies that robustness in the sense of KM is strictly weaker than strong robustness. Consider for instance the following game:

	L	R
T	$0,0$	$0,0$
B	$0,0$	$0,0$

Every action profile is 0-robust but Lemma 1 implies that none of them is d-robust for $d > 0$ and hence, no equilibrium is strongly robust.

Still, provided that a^* is a strict Nash equilibrium, Section 2.2 shows that typical sufficient conditions that ensure robustness of a^* in the sense of KM also imply strong robustness.

2.2 Sufficient Conditions for Strong Robustness

In general, it is difficult to check robustness of a given equilibrium a^* of G to all possible nearby elaborations. KM however provide sufficient conditions under which a^* is robust: when it is the unique correlated equilibrium of G, and when it is a **p**-dominant equilibrium with $\sum_i p_i < 1$. Here we show that, whenever a^* is a strict Nash equilibrium, the same conditions imply that a^* is d-robust in G for some $d > 0$. We give proofs for the second and the third classes in the Appendix, which are modifications of KM (Propositions 3.2 and 5.3). We begin with the case where a^* is the unique correlated equilibrium of G.

Proposition 1 (strong robustness of unique correlated equilibria)**.** *If a^* is the unique correlated equilibrium and is a strict Nash equilibrium, then a^* is strongly robust.*

The proof, given in the Appendix, is an extension of the proof of Proposition 3.2 in KM. We also exploit the fact that when a^* is the unique correlated equilibrium of G and is strict Nash, then it is also the unique correlated equilibrium of all sufficiently close complete information games. A useful special case is the one where a^* is the only equilibrium surviving elimination of strictly dominated strategies.

Definition 2 (iteratively d-strict dominant equilibrium)**.** Consider a^* an iteratively strictly dominant equilibrium. At any stage of elimination T let us denote i_T the considered player and A^T_{-i} the set of surviving actions player i may face. We say that a^* is iteratively d-strict dominant in G if and only if at every stage of elimination T, for all $a_i \in A_i$ and all $a_{-i} \in A^T_{-i}$,

$$g_i(a_i^*, a_{-i}) \geq g_i(a_i, a_{-i}) + d.$$

Proposition 2 (strong robustness of iteratively d-strict dominant equilibrium)**.** *Whenever a^* is iteratively d-strict dominant, a^* is $d/2$-robust.*

KM's second sufficient condition—which is particularly useful in applied settings—establishes that, whenever a^* is $\mathbf{p}$-dominant equilibria with $\sum_i p_i < 1$, it is robust in G. We now study how this result extends to d-robustness. We begin by recalling the definition of strict $\mathbf{p}$-dominance.

Definition 3 (strict $\mathbf{p}$-dominance)**.** For $\mathbf{p} = (p_1, \ldots, p_n) \in (0, 1]^n$, an action profile a^* is a *strictly* $\mathbf{p}$-*dominant equilibrium* of G if, for all $i \in N$, $a_i \in A_i \setminus \{a_i^*\}$, and $\lambda \in \Delta(A_{-i})$ such that $\lambda(a_{-i}^*) > p_i$,

$$\sum_{a_{-i}\in A_{-i}} \lambda(a_{-i}) g_i(a_i^*, a_{-i}) > \sum_{a_{-i}\in A_{-i}} \lambda(a_{-i}) g_i(a_i, a_{-i}).$$

Given this definition, the following result holds.

Proposition 3 (strong robustness of strictly **p**-dominant equilibria)**.** *If action profile a^* is strictly* **p***-dominant with $\sum_{i\in N} p_i < 1$, then a^* is strongly robust.*

Since we know from KM Lemma 5.5 that, if a game has a strictly **p**-dominant equilibrium with $\sum_i p_i \leq 1$, then no other action profile is 0-robust, this implies that, if a game has a strictly **p**-dominant equilibrium with $\sum_i p_i < 1$, it is the only strongly robust equilibrium of G.

Proposition 3 implies that, in 2 × 2 coordination games, a strictly risk-dominant equilibrium will also be strongly robust.

3 Robustness in repeated games

The main purpose of this section is to formulate a notion of robustness to incomplete information that's appropriate for repeated games. There are many potential definitions of robustness, each one corresponding to a different class of perturbed games used to check robustness. Here we choose to respect the original repeated-game structure and check robustness only to additively separable and intertemporally independent payoff shocks. We show in Section 3.2.2 that this notion of robustness allows for a convenient recursive representation of dynamically robust equilibria.[2]

---

[2] Another possible approach could be to consider the normal-form representation of a repeated game and check robustness to any sort of payoff perturbation associated with each strategy profile. This however corresponds to a very large class of perturbed games that would satisfy neither additive separability, intertemporal independence, nor even continuity at infinity.

3.1 Definition

Consider a complete-information game $G = (N, (A_i, g_i)_{i \in N})$, and define Γ_G the repeated game of stage game G, with discount factor $\delta < 1$, and without public randomization.[3] Let $H_t = A^t$ be the set of histories (sequences of realized action profiles) of length t and $H = \bigcup_{t \geq 0} H_t$ be the set of all finite histories. A strategy of player i is described by $s_i \colon H \to A_i$. For each history $h \in H$, s_i induces continuation strategy $s_i|h$ after history h by $(s_i|h)(h') = s_i(h, h')$ for each $h' \in H$. A strategy profile s induces a sequence of action profiles, $\{a_t\}$, which gives each player i the total payoff

$$v_i(s) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} g_i(a_t).$$

A strategy profile s^* is a *subgame-perfect equilibrium* (*SPE*) if and only if $v_i(s^*|h) \geq v_i((s_i, s^*_{-i})|h)$ for every $h \in H$, $i \in N$, and every strategy s_i. We denote by $\mathcal{V}^{\text{SPE}}$ the set of SPE values of Γ_G.

We define perturbations of game Γ_G as follows. Consider a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete-information games $U_t = (N, \Omega_t, P_t, (A_i, u_{it}, Q_{it})_{i \in N})$ that embed G. We assume that $\{u_t\}_{t \in \mathbb{N}}$ is bounded uniformly by some M such that $M > |g|$. We also assume that the sequence $\{\omega_t\}_{t \in \mathbb{N}}$ of states is a sequence of independent random variables. At the beginning of period t, state $\omega_t \in \Omega_t$ is generated according to P_t. Given such a sequence U we denote by $\Gamma_{\mathbf{U}}$ the infinite horizon game with observable actions in which players have a common discount factor $\delta < 1$ and play game U_t at every time $t \in \mathbb{N}$. Public histories $h_t \in H$ are sequences of observed action profiles $\{a_1, \cdots, a_t\}$. A public strategy of player i is a mapping $\sigma_i(h_{t-1}, \omega_t)$ such that, at every history h_{t-1}, $\sigma_i(h_t, \cdot)$ is a mapping from Ω_t to $\Delta(A_i)$ that is Q_{it}-measurable. Conditional on a history $h_{t-1} \in H$, a public strategy profile σ induces a probability distribution over sequences of action profiles and states, which allows us to define

---

[3] We incorporate public randomization in Section 3.3.

continuation payoffs

$$\forall i \in N,\ \forall h_t \in H, \quad v_i(\sigma|h_{t-1}) = \mathbb{E}\left[(1-\delta)\sum_{\tau=1}^{\infty} \delta^{\tau-1} u_{i,t+\tau-1}(a_{t+\tau-1}, \omega_{t+\tau-1})\right]$$

The assumption that $\{u_t\}_{t\in\mathbb{N}}$ is bounded uniformly by some M implies that the infinite sum is well defined. A public strategy profile σ^* is a *perfect public equilibrium* (*PPE*) if and only if $v_i(\sigma^*|h_{t-1}) \geq v_i((\sigma_i, \sigma^*_{-i})|h_{t-1})$ for every $h_{t-1} \in H$, $i \in N$, and every public strategy σ_i.

Definition 4 (dynamic d-robustness)**.** For $d \geq 0$, an SPE s^* of Γ_G is *d-robust* if, for every $\eta > 0$, there exists $\varepsilon > 0$ such that, for every sequence $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ of independent (ε, d)-elaborations of G, game $\Gamma_{\mathbf{U}}$ has a PPE σ^* such that $P^{\sigma^*(h_{t-1},\cdot)}(s^*(h_{t-1})) > 1-\eta$ for every $h_{t-1} \in H$.

We say that an SPE s^* of Γ_G is *strongly robust* if it is d-robust for some $d > 0$.

Let $\mathcal{V}^{\text{rob}}$ be the set of strongly robust SPE equilibrium payoff profiles of the repeated game of G. Section 3.2.2 provides a recursive characterization of this set.

Note that our definition of dynamic strong robustness considers only perturbations $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ that are uniformly close to G. If $\mathbf{U}$ approached G only pointwisely, then the robustness criterion would become too stringent. For example, if we allowed for such perturbations, whenever the stage game G had a unique Nash equilibrium a^*, the only strongly robust equilibrium of Γ_G would be the repetition of a^*. Indeed, consider the perturbation $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ such that U_t is identical to G for $t \leq T$ and $u_{it}(a, \omega_t) = 0$ for every $i \in N$, $a \in A$, $t > T$, and $\omega_t \in \Omega_t$, then the game $\Gamma_{\mathbf{U}}$ has an finite effective horizon, and by backward induction, players must play a^* for the first T periods. This is why we consider uniformly close perturbations.

3.2 A factorization result

In this section we show that the set strongly robust equilibrium values admits a recursive characterization analogous to that of Abreu, Pearce and Stacchetti (1990). The construction builds on a one-shot robustness principle that is the equivalent of the one-shot deviation principle for dynamically robust equilibria. As Section 5.2 will highlight it is important to use d-robustness with $d > 0$ rather than robustness in the sense of KM why for this one-shot robustness principle to obtain. Indeed, there can be SPEs whose one-shot action profiles are robust in all the appropriate augmented games but are not dynamically robust under any acceptable definition.

3.2.1 Augmented Games

For $G = (N, (A_i, g_i)_{i \in N})$ and $w \colon A \to \mathbb{R}^n$, let $G(w)$ be the complete-information game augmented with continuation payoff profile w contingent on the current action profile, i.e., $G(w) = (N, (A_i, g'_i)_{i \in N})$ such that $g'_i(a) = (1-\delta) g_i(a) + \delta w_i(a)$ for every $i \in N$ and $a \in A$. It follows from the one-shot deviation principle that s^* is an SPE of repeated game Γ_G if and only if $s^*(h)$ is a Nash equilibrium of $G(v(s^*|(h,\cdot)))$ for all $h \in H$, where $v(s^*|(h,\cdot))$ is a contingent payoff profile that maps a to $v(s^*|(h,a))$.

Similarly, given an incomplete information game $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ and $w \colon A \to \mathbb{R}^n$, let $U(w)$ be the incomplete-information game $(N, \Omega, P, (A_i, u'_i, Q_i)_{i \in N})$ such that $u'_i(a, \omega) = (1-\delta) u_i(a, \omega) + \delta w_i(a)$ for every $i \in N$, $a \in A$, and $\omega \in \Omega$. Similarly, given a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of independent incomplete information games, a strategy profile σ^* is a PPE of $\Gamma_{\mathbf{U}}$ if and only if $\sigma^*(h_{t-1}, \cdot)$ is a static Bayesian-Nash equilibrium of $U_t(v(\sigma^*|(h_{t-1}, \cdot)))$ for all $h_{t-1} \in H$.

The results of Section 3.2.2 relate the d-robustness of an SPE s^* to the d-robustness of $s^*(h)$ in $G(v(s^*|(h,\cdot)))$ for all $h \in H$.

3.2.2 Self-generation and factorization

This section relates the d-robustness of an SPE s^* to the d-robustness of $s^*(h)$ in $G(v(s^*|(h,\cdot)))$ for all $h \in H$ and characterizes strongly robust SPEs of Γ_G. More precisely we prove self-generation and factorizations results analogous to those of APS. We begin by showing that, whenever s^* is strongly robust, $s^*(h)$ is a strict Nash equilibrium of $G(v(s^*|(h,\cdot)))$ for all

$h \in H$.

Lemma 2 (strict enforcement)**.** *If* s^* *is a* d*-robust SPE of the repeated game* Γ_G *with* $d \leq M - |g|$*, then* $s^*(h)$ *is a* $2(1-\delta)d$*-strict equilibrium of* $G(v(s^*|(h,\cdot)))$ *for all* $h \in H$*.*

Proof. Fix any $t^0 \geq 1$ and $h^0 \in H_{t^0-1}$. Consider $\mathbf{U} = \{U_t\}$ such that U_t is identical to G for $t \neq t^0$ and U_{t^0} is identical to a complete information game $G' = (N, (A_i, g'_i)_{i \in N})$ such that, for every $i \in N$, $g'_i(a) = g_i(a) + d$ for $a \neq s^*(h^0)$ and $g'_i(s^*(h^0)) = g_i(s^*(h^0)) - d$. Since every U_t is an (ε, d)-elaboration of G for any $\varepsilon > 0$ with payoffs bounded by M, $\Gamma_{\mathbf{U}}$ admits a PPE arbitrarily close to s^*. This implies that $s^*(h^0)$ is an equilibrium of $G'(v(s^*|(h^0,\cdot)))$, thus a $2(1-\delta)d$-strict equilibrium of $G(v(s^*|(h^0,\cdot)))$. □

Definition 5 (robust enforcement and robust decomposability)**.** For $a \in A$, $w \colon A \to \mathbb{R}^n$, and $d > 0$, w *enforces* a *d-robustly* if a is a d-robust equilibrium of $G(w)$.

For $v \in \mathbb{R}^n$, $\mathcal{V} \subseteq \mathbb{R}^n$, and $d > 0$, v is *d-robustly decomposable on* $\mathcal{V}$ if there exist $a \in A$ and $w \colon A \to \mathcal{V}$ such that w enforces a d-robustly and $v = (1-\delta)g(a) + \delta w(a)$.

Let $B^d(\mathcal{V})$ be the set of all payoff profiles that are d-robustly decomposable on $\mathcal{V}$. B^d is a generating function that maps subsets of $\mathbb{R}^n$ to subsets of $\mathbb{R}^n$.

Lemma 3 (monotonicity of B^d)**.** *If* $\mathcal{V} \subseteq \mathcal{V}' \subseteq \operatorname{co} g(A)$*, then* $B^d(\mathcal{V}) \subseteq B^d(\mathcal{V}') \subseteq \operatorname{co} g(A)$*.*

Proof. Monotonicity follows directly from the definition of B^d. □

Thus, by Tarski's fixed point theorem, B^d admits the largest fixed point among all subsets of $\operatorname{co} g(A)$. We denote it by $\mathcal{V}^d$. We say that $\mathcal{V}$ is *self-generating with respect to* B^d if $\mathcal{V} \subseteq B^d(\mathcal{V})$.

Lemma 4 (self-generation)**.** *If* $\mathcal{V}$ *is a subset of* $\operatorname{co} g(A)$ *and is self-generating, then* $\mathcal{V} \subseteq \mathcal{V}^d$*.*

Proof. Lemma 4 follows from Tarski's fixed point theorem. □

The next proposition characterizes the set of all strongly robust SPE payoff profiles in terms of fixed points of B^d.

Theorem 1 (characterization of $\mathcal{V}^{\text{rob}}$)**.** *Whenever* $\mathcal{V}$ *is any compact set that contains* $\mathcal{V}^{\text{rob}}$*,*

$$\mathcal{V}^{\text{rob}} = \bigcup_{d>0} \mathcal{V}^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (B^d)^k(\mathcal{V}).$$

Theorem 1 corresponds to APS's self-generation, factorization, and algorithm results (Abreu et. al., 1990, Theorems 1, 2, and 5). APS define $B(\mathcal{V})$ as the set of all payoff profiles v such that there exist $a \in A$ and $w\colon A \to \mathcal{V}$ such that a is a Nash equilibrium of $G(w)$. APS show that $\mathcal{V}^{\text{SPE}}$ is the largest bounded fixed point of B and computed by iterating the B mapping to a sufficiently large compact set infinitely many times. The main difference is that we require strong robustness of SPE and thus take into account the robustness criterion in each B^d mapping. Note that $\operatorname{co} g(A)$ is a compact set that contains $\mathcal{V}^{\text{rob}}$.

The proof of Theorem 1 is based on the following proposition, whose proof is given in the Appendix.

Proposition 4 (one-shot robustness principle)**.**

(i) If $0 < d \leq M - |g|$ *and* s^* *is a* d*-robust SPE of the repeated game of* G*, then* $s^*(h)$ *is a* $(1-\delta)d$*-robust equilibrium of* $G(v(s^*|(h,\cdot)))$ *for every* $h \in H$*.*

(ii) If $d > 0$ *and* $s^*(h)$ *is a* d*-robust equilibrium of* $G(v(s^*|(h,\cdot)))$ *for every* $h \in H$*, then* s^* *is a* d'*-robust SPE of the repeated game of* G *for every* $d' < d$*.*

Proof of Theorem 1. For each $v \in \mathcal{V}^{\text{rob}}$, let s^* be a d-robust SPE of Γ_G, with $0 < d \leq M - |g|$, that generates v . Then, by Proposition 4, $\mathcal{V}^* = \{v(s^*|h) \in \mathbb{R}^n \mid h \in H\}$ is self-generating with respect to $B^{(1-\delta)d}$. By Lemma 4, $v \in \mathcal{V}^* \subseteq \mathcal{V}^{(1-\delta)d}$. Thus $\mathcal{V}^{\text{rob}} \subseteq \bigcup_{d>0} \mathcal{V}^{(1-\delta)d} = \bigcup_{d>0} \mathcal{V}^d$.

Since $\mathcal{V}^d$ is self-generating with respect to B^d, for each $v \in \mathcal{V}^d$ there exist $a(v) \in A$ and $w(v,\cdot)\colon A \to \mathcal{V}^d$ such that $w(v,\cdot)$ enforces $a(v)$ d-robustly and $v = (1-\delta)g(a(v)) + \delta w(v, a(v))$. Pick any $v \in \mathcal{V}^d$ with $d > 0$. We construct s^* recursively as follows: $s^*(\emptyset) = a(v)$, $s^*(a_1) = a(w(v,a_1))$, $s^*(a_1,a_2) = a(w(w(v,a_1),a_2))$, and so on... By construction, for every

$h \in H$, $s^*(h)$ is d-robust in $G(v(s^*|(h, \cdot)))$ for any d. Then, by Proposition 4, s^* is $d/2$-robust in the repeated game of G and thus $v \in \mathcal{V}^{\mathrm{rob}}$. Thus $\mathcal{V}^d \subseteq \mathcal{V}^{\mathrm{rob}}$.

To prove the analogue of APS's algorithm result, we define $C^d(\mathcal{V})$ by the closure of $B^d(\mathcal{V})$. Fix a compact set $\mathcal{V}$ that contains $\mathcal{V}^{\mathrm{rob}}$. We will show that $\mathcal{V}^{\mathrm{rob}} = \bigcup_{d>0}(B^d)^\infty(\mathcal{V}) = \bigcup_{d>0}(C^d)^\infty(\mathcal{V})$, where $f^\infty(\mathcal{V}) = \bigcap_{k=0}^\infty f^k(\mathcal{V})$ for $f = B^d$ or C^d. Since $\mathcal{V}^d \subset \mathcal{V}^{\mathrm{rob}} \subseteq \mathcal{V}$ and by monotonicity of B^d and C^d, we have $\mathcal{V}^d \subseteq (B^d)^\infty(\mathcal{V}) \subseteq (C^d)^\infty(\mathcal{V})$ for each $d > 0$. Thus, by Lemma 4, it is enough to show that $(C^d)^\infty(\mathcal{V})$ is self-generating with respect to $B^{d'}$ for some $d' < d$ and hence $(C^d)^\infty(\mathcal{V}) \subseteq \mathcal{V}^{d'}$.

Fix any $d' < d$. Pick any $v \in (C^d)^\infty(\mathcal{V})$. For each $k \geq 1$, since we have $v \in (C^d)^\infty(\mathcal{V}) \subseteq (C^d)^k(\mathcal{V})$, there exist $a^k \in A$ and $w^k \colon A \to (C^d)^{k-1}(\mathcal{V})$ such that a^k is d-robust in $G(w^k)$. By taking a subsequence, we can assume without loss of generality that $a^k = a^*$ and $w^k(a) \to w^*(a)$ as $k \to \infty$ for each $a \in A$. This means there exists $k \geq 1$ such that for all $l \geq k$, $|w^l - w^*| \leq (d - d')/\delta$. Since a^* is d-robust in $G(w^l)$, a^* is d'-robust in $G(w^*)$. Moreover, for each $k \geq 1$, since $w^l \in (C^d)^{k-1}(\mathcal{V})$ for any $l \geq k$ and $(C^d)^{k-1}(\mathcal{V})$ is compact, $w^* \in (C^d)^{k-1}(\mathcal{V})$. Since this holds for every $k \geq 1$, $w^* \in (C^d)^\infty(\mathcal{V})$. Thus we have $v \in B^{d'}((C^d)^\infty(\mathcal{V}))$. $\square$

An immediate corollary of Proposition 4 is that a finite-automaton-strategy equilibrium s^* is strongly robust if and only if $s^*(h)$ is strongly robust in $G(v(s^*|(h, \cdot)))$ for every $h \in H$. As Section 4.1 shows, Theorem 1 and Proposition 4 are useful in characterizing both robust SPEs and robust equilibrium values of repeated games of interest, such as the repeated Prisoners' Dilemma. Note that the one-shot robustness principle of Proposition 4 would not hold if we considered robustness in the sense of KM rather than d-robustness. As a counter example, Section 5.2 describes a mixed-strategy equilibrium of an infinite horizon dynamic game Γ^* such that, for every history, the action profile prescribed by the equilibrium is robust in the sense of KM in the appropriate augmented game, but the whole equilibrium is not dynamically robust in any reasonable sense.

3.3 Public Randomization

This subsection extends our framework to allow for public randomization. Given a complete information G, we denote by $\tilde{\Gamma}_G$ the repeated game of stage game G with public randomization, in which at the beginning of each period t, players observe a common signal θ_t uniformly distributed on $[0, 1)$. We assume that the sequence $\{\theta_t\}_{t\in\mathbb{N}}$ is independent across time. We write $\theta^t = (\theta_1, \ldots, \theta_t) \in [0, 1)^t$. A pure strategy of player i is described by

$$\tilde{s}_i \colon \bigcup_{t\geq 1}(H_{t-1} \times [0, 1)^t) \to A_i$$

such that, for each $t \geq 1$ and $h_{t-1} \in H_{t-1}$, $\tilde{s}_i(h_{t-1}, \cdot)$ is Borel-measurable on $[0, 1)^t$. For each history $(h_{t-1}, \theta^t) \in H_{t-1} \times [0, 1)^t$, $\tilde{s}_i$ induces continuation strategy $\tilde{s}_i|(h_{t-1}, \theta^t)$. Conditional on each realization θ, a strategy profile $\tilde{s}$ induces a distribution over the set of sequences of action profiles $\{a_t\}$, which gives each player i payoff

$$\tilde{v}_i(\tilde{s}|\theta) = \mathbb{E}\left[(1-\delta)\sum_{t=1}^{\infty}\delta^t g_i(a_t)\right].$$

A strategy profile $\tilde{s}^*$ is an SPE if $\tilde{v}_i(s^*|(h_{t-1}, \theta^t)) \geq \tilde{v}_i(\tilde{s}_i, \tilde{s}^*_{-i}|(h_{t-1}, \theta^t))$ for every $t \geq 1$, $(h_{t-1}, \theta^t) \in H_{t-1} \times [0, 1)^t$, $i \in N$, and $\tilde{s}_i$.

Similarly, given a sequence $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ of independent incomplete-information games, we consider the corresponding dynamic game $\tilde{\Gamma}_{\mathbf{U}}$ with public randomization, in which at the beginning of period t players observe a common signal θ_t, that is uniformly distributed on $[0, 1)$. A public strategy of player i is described by

$$\tilde{\sigma}_i \colon \bigcup_{t\geq 1}(H_{t-1} \times [0, 1)^t \times \Omega_t) \to \Delta(A_i)$$

such that, for every $t \geq 1$ and $h_{t-1} \in H_{t-1}$, $\tilde{\sigma}_i(h_{t-1}, \theta^t, \cdot)$ is Q_{it}-measurable for almost every $\theta^t \in [0, 1)^t$ and $\tilde{\sigma}_i(h_{t-1}, \cdot, \omega_t)$ is Borel-measurable for every $\omega_t \in \Omega_t$. Continuation payoffs $\tilde{v}_{it}(\tilde{\sigma}|(h_{t-1}, \theta^t))$ and PPE are defined as in Section 3.1.

We define robustness in repeated games with public randomization as follows.

Definition 6. An SPE $\tilde{s}^*$ of $\tilde{\Gamma}_G$ is *d-robust* for $d > 0$ if, for all $\eta > 0$, there exists $\varepsilon > 0$ such that for every sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of independent (ε, d)-elaboration of G, game $\tilde{\Gamma}_{\mathbf{U}}$ has a PPE $\tilde{\sigma}^*$ such that $P_t^{\tilde{\sigma}^*(h_{t-1},\theta^t,\cdot)}(\tilde{s}^*(h_{t-1},\theta^t)) > 1 - \eta$ for every $t \geq 1$, and $h_{t-1} \in H_{t-1}$, and $\theta^t \in [0,1)^t$. An SPE s^* is *strongly robust* if it is d-robust for some $d > 0$.

Let $\tilde{\mathcal{V}}^{\text{rob}}$ denote the set of strongly robust SPE payoff profiles in $\tilde{\Gamma}_G$. For each $d > 0$ and $W \subseteq \mathbb{R}^n$, let

$$\tilde{B}^d(W) = \operatorname{co} B^d(W).$$

Mapping $\tilde{B}^d$ satisfies the monotonicity property described in Lemma 3, and hence $\tilde{B}^d$ admits a largest fixed point among subsets of $\operatorname{co} g(A)$. We denote it by $\tilde{\mathcal{V}}^d$. We have the following characterization, similar to Theorem 1. The proof is relegated to the Appendix.

Proposition 5 (characterization of $\tilde{\mathcal{V}}^{\text{rob}}$)**.** *Whenever $\mathcal{V}$ is any compact set that contains $\mathcal{V}^{\text{rob}}$,*

$$\tilde{\mathcal{V}}^{\text{rob}} = \bigcup_{d>0} \tilde{\mathcal{V}}^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (\tilde{B}^d)^k(\mathcal{V}).$$

4 Applications

In this section, we give two applications of the main characterization results of Section 3 (Theorem 1, Proposition 4 and Proposition 5).

In our main application, we characterize robust equilibrium values and robust strategies in the repeated Prisoners' Dilemma. We show that whenever outcome (*Cooperate*, *Defect*) can be enforced in equilibrium under complete information, the set of robust equilibrium values is essentially equal to the set of equilibrium values under complete information. Inversely, whenever (*Cooperate*, *Defect*) is not enforceable in equilibrium, the set of robust equilibria shrinks to permanent defection. We also show that robustness considerations nuance our intuitions about which equilibria best sustain cooperation. More precisely asymmetric equilibria in which only the deviator is punished upon unilateral deviation are more effective than grim-trigger strategies that punish everybody.

In our second application, we use the factorization result to prove a folk theorem in robust equilibria. Because a folk theorem holds, this means that for discount factor close to one, robustness does not restrict the set of sustainable equilibrium values. Still robustness may imply significant restrictions given a fixed discount factor.

4.1 Robustness in the Repeated Prisoners' Dilemma

Consider the one-shot two-player Prisoners' Dilemma (PD) with flow payoffs

	C	D
C	$1, 1$	$-c, b$
D	$b, -c$	$0, 0$

where $b > 1$, $c > 0$ and $b - c < 2$. We are interested in the repeated game $\tilde{\Gamma}_G$ in which players observe a public randomization device. In what follows, we characterize the set of robust equilibrium payoffs of this game and in particular we determine when cooperating is a robust equilibrium outcome.

4.1.1 Equilibrium payoffs under complete information

In this section, we summarize the results of Stahl (1991), which characterize the set of equilibrium payoffs of the complete information repeated Prisoners' Dilemma as a function of its parameters b, c, and δ.

Definition 7 (auxiliary parameters)**.** Given (b, c, δ), we define the following parameters

$$\begin{aligned} p &\equiv \frac{b+c}{1+c} \\ h &\equiv \frac{(b-1)(5b-1)}{4b} \\ \delta^* &\equiv \frac{(b-1)^2 - 2(1+c) + 2\sqrt{(1+c)^2 - (b-1)^2}}{(b-1)^2} \\ q &\equiv \max\left\{1, \frac{1+\delta+(1-\delta)b+\sqrt{[1+\delta+(1-\delta)b]^2 - 4(1-\delta)(b+c)}}{2}\right\}. \end{aligned}$$

Definition 8 (value sets)**.** Let us denote

(i) $\mathcal{V}_0$ the set of feasible, individually rational values of the one shot PD: $\mathcal{V}_0 = \mathrm{co}\{(0,0),(1,1),(0,p),(p,0)\}$.

(ii) $\mathcal{V}_Q$ the set of values defined by $\mathcal{V}_Q = \mathrm{co}\{(0,0),(1,1),(0,q),(q,0)\}$.

(iii) $\mathcal{V}_T$ the set of values defined by $\mathcal{V}_T = \mathrm{co}\{(0,0),(0,b-c),(b-c,0)\}$.

(iv) $\mathcal{V}_D$ the set of values defined by $\mathcal{V}_D = \mathrm{co}\{(0,0),(1,1)\}$.

Proposition 6 (Stahl's characterization)**.** *Let us denote by $\mathcal{V}^{\mathrm{SPE}}$ the set of subgame-perfect equilibrium values of the repeated PD. We have the following.*

(i) If $\delta \geq \max\{(b-1)/b, c/(c+1)\}$ then $\mathcal{V}^{\mathrm{SPE}} = \mathcal{V}_0$.

(ii) If $b-1 \leq c \leq h$ and $(b-1)/b \leq \delta < c/(c+1)$, or $c > h$ and $\delta^* \leq \delta < c/(c+1)$, then $\mathcal{V}^{\mathrm{SPE}} = \mathcal{V}_Q$.

(iii) If $c < b-1$ and $\delta \in [c/b, (b-1)/b]$ then $\mathcal{V}^{\mathrm{SPE}} = \mathcal{V}_T$.

(iv) If $c > h$ and $\delta \in [(b-1)/b, \delta^*)$, then $\mathcal{V}^{\mathrm{SPE}} = \mathcal{V}_D$.

(v) If $\delta < \min\{c/b, (b-1)/b\}$, then $\mathcal{V}^{\mathrm{SPE}} = \{(0,0)\}$.

4.1.2 Robust SPEs

Using Stahl's characterization, we can classify PD games as follows.

Definition 9 (classification of Prisoners' Dilemma games)**.** The repeated PD $\tilde{\Gamma}_G$ is characterized by parameters (b, c, δ). We define the three classes of games $\mathcal{G}_{DC/CC}$, $\mathcal{G}_{DC}$ and $\mathcal{G}_{CC}$ as follows:

(i) $\mathcal{G}_{DC/CC}$ is the class of PDs such that both DC and CC are enforceable action profiles in equilibrium.

(ii) $\mathcal{G}_{DC}$ is the class of PDs such that DC is an enforceable action profile in equilibrium, but CC is not.

(iii) $\mathcal{G}_{CC}$ is the class of PDs such that CC is an enforceable action profile in equilibrium, but DC is not.

The class of games $\mathcal{G}_{DC/CC}$ corresponds to cases (i) and (ii) of Stahl's characterization, while $\mathcal{G}_{DC}$ and $\mathcal{G}_{CC}$ respectively correspond to cases (iii) and (iv). Whenever action profile DC is enforceable in equilibrium, the following remarkable property holds.

Lemma 5 (stage dominant implementation)**.** *Consider a Prisoners' Dilemma game belonging to* $\operatorname{int}\mathcal{G}_{DC/CC} \cup \operatorname{int}\mathcal{G}_{DC}$*. Then for any pair of SPE values* $v \in \{(0,0),(1,1)\} \cup \operatorname{int}\mathcal{V}^{\mathrm{SPE}}$*, there exists* $d > 0$ *and a perfect public equilibrium* σ *implementing these values such that at every history* h*, the equilibrium action profile* $\sigma(h)$ *is iteratively* d*-strict dominant in the augmented game* $G(v(\sigma|(h,\cdot)))$*.*

We call such SPEs stage dominant equilibria. Lemma 5 yields the following robustness result.

Proposition 7 (robust equilibria)**.** *Whenever* $\tilde{\Gamma}_G \in \operatorname{int}\mathcal{G}_{DC,CC} \cup \operatorname{int}\mathcal{G}_{DC}$*,*

$$\{(0,0),(1,1)\} \cup \operatorname{int}\mathcal{V}^{\mathrm{SPE}} \subset \mathcal{V}^{\mathrm{rob}} \subset \mathcal{V}^{\mathrm{SPE}}.$$

Proof. This is a corollary of Theorem 1 and Lemma 5. $\square$

Hence, whenever DC is enforceable in equilibrium, essentially any SPE value is robustly implementable. We now show that, whenever DC is not enforceable, the only strongly robust SPE is permanent defection.

Proposition 8 (fragile equilibria)**.** *Whenever* $\tilde{\Gamma}_G \in \mathcal{G}_{CC}$*, we have* $\mathcal{V}^{\text{rob}} = \{(0,0)\}$ *and the only strongly robust SPE is permanent defection.*

Proof. The proof is by contradiction. Assume there exists a strongly robust subgame-perfect equilibrium of $\tilde{\Gamma}_G$ generating value different from $(0,0)$. This implies that at some equilibrium history, action profile CC must be taken. Hence there must exist a strongly robust equilibrium s^* of $\tilde{\Gamma}_G$ such that action profile CC is taken at the initial history $h = \emptyset$. Because $\tilde{\Gamma}_G \in \mathcal{G}_{CC}$, we have that, at all histories h, $s^*(h) \in \{CC, DD\}$, and hence we have that for all $h \in H$, $v_1(h) = v_2(h)$. Furthermore, we must have that $\delta/(1-\delta) < c$; otherwise, DC would be enforceable.

The augmented game at history $h = \emptyset$ is symmetric and takes the form

	C	D
C	$1-\delta+\delta v_{CC}, 1-\delta+\delta v_{CC}$	$-c(1-\delta)+\delta v_{CD}, b(1-\delta)+\delta v_{CD}$
D	$b(1-\delta)+\delta v_{DC}, -c(1-\delta)+\delta v_{DC}$	$\delta v_{DD}, \delta v_{DD}$

where v_{CC}, v_{CD}, v_{DC} and v_{DD} are in $[0,1]$.

We now show that DD is the risk-dominant equilibrium of this augmented game, i.e. that

$$(1+\lambda v_{CC} - b - \lambda v_{CD})(1+\lambda v_{CC} - b - \lambda v_{DC}) < (\lambda v_{DD} + c - \lambda v_{CD})(\lambda v_{DD} + c - \lambda v_{DC})$$

where $\lambda = \delta/(1-\delta)$. Since $c > \delta/(1-\delta)$, we have that $\delta v_{DD} + c - \delta v_{CD} > 0$ and that

$\delta v_{DD} + c - \delta v_{DC} > 0$. This means we can define

$$R \equiv \frac{(\pi + \lambda v_{CC} - b - \lambda v_{CD})(\pi + \lambda v_{CC} - b - \lambda v_{DC})}{(\lambda v_{DD} + c - \lambda v_{CD})(\lambda v_{DD} + c - \lambda v_{DC})}.$$

As a function of v_{CC}, v_{CD}, v_{DC} and v_{DD}, the ratio R is maximized for $v_{CC} = 1$, $v_{DD} = 0$ and $v_{DC} = v_{CD} = v$. Since $c > \delta/(1-\delta)$ and $b > 1$, for these vs we have

$$\begin{aligned} & c > 1 + \frac{\delta}{1-\delta} - b \\ \implies\ & \lambda v_{DD} + c - \lambda v > 1 + \lambda v_{CC} - b - \lambda v \\ \implies\ & (\lambda v_{DD} + c - \lambda v)^2 > (1 + \lambda v_{CC} - b - \delta v)^2 \\ \implies\ & R < 1. \end{aligned}$$

This implies that indeed DD is risk dominant in $G(v(s|(\emptyset, \cdot)))$. Hence CC cannot be robust which contradicts the fact that s is strongly robust. This concludes the proof. The only strongly robust equilibrium of $\tilde{\Gamma}_G$ is for both players to defect at every history. $\square$

A corollary of this result is that, whenever cooperation is robustly sustainable in grim-trigger strategies, DC is enforceable in equilibrium and cooperation is in fact sustainable in stage dominant strategies. Inversely, it is possible for cooperation to be robustly sustainable in stage dominant strategies but not in grim-trigger strategies. This suggests that grim-trigger strategies are not the most efficient way to robustly implement cooperation. Indeed, while grim-trigger strategies punish both players when a unilateral deviation occurs, stage dominant strategies only punish the deviator, while rewarding the other player. This enhances their robustness to incomplete information.

4.2 A Folk Theorem in Robust strategies

In this Section, we prove a folk theorem for strongly robust SPEs. Namely, we show that every feasible and strictly individually rational payoff profile can be sustained as a strongly robust SPE payoff profile for δ sufficiently close to 1. It implies that, in the limit as $\delta \to 1$, requiring strong robustness does not impose any essential restriction on the set of equilibrium payoffs (up to weakly individually rational payoff profiles). This does not mean there is not much to learn from robustness analysis. As Section 4.1 shows, robustness considerations can lead to significant restriction of the equilibrium value set for fixed discount factors. Moreover, to prove the folk theorem under robustness constraints we must consider classes of strategies that are intuitively appealing and differ from those used in Fudenberg and Maskin (1986) or Abreu, Dutta and Smith (1994). More precisely to implement values robustly, we must take care not only of continuation values following unilateral deviations, but also of continuation values following the simultaneous deviation of multiple player. This leads us to construct SPEs in which each one-shot action profile is iteratively strictly dominant in the appropriate augmented game.

We say that two payoff functions g_i and g_j are *equivalent* if g_i is a positive affine transformation of g_j. G satisfies the *nonequivalent utility* (*NEU*) *condition* if no pair of players have equivalent payoff functions in G (Abreu, Dutta, and Smith (1994)). Note that the full-dimensionality condition in Fudenberg and Maskin (1986) implies the NEU condition.

For each $i \in N$, let

$$m^i_{-i} \in \arg\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} g_i(a),$$

$$m^i_i \in \arg\max_{a_i \in A_i} g_i(a_i, m^i_{-i})$$

be an action profile that minimaxes player i's payoff in pure actions and a best response to m^i_{-i}. (If there are multiple minimax action profiles, pick any such one.) Without loss of generality, we assume that $g_i(m^i) = 0$.

Proposition 9 (folk theorem)**.** *Assume that G satisfies the NEU condition and public randomization is available.*[4] *If $v \in \operatorname{co} g(A)$ with $v \gg 0$, then there exists $\underline{\delta} < 1$ such that, for*

---

[4] We could extend the results to games without public randomization by using nonstationary sequences of action profiles as in Fudenberg and Maskin (1991).

every $\delta > \underline{\delta}$, there exists a strongly robust SPE of $\tilde{\Gamma}_G$ with payoff profile v.

The idea of the proof is to show that for each $v \in \operatorname{co} g(A)$ with $v \gg 0$, we can construct s^* with payoff profile v such that s^* is a profile of finite-state automaton strategies, and, for large δ, $s^*(h)$ survives iterative elimination of strictly dominated actions in $G(v(s^*|(h, \cdot)))$ for every history $h \in H$. Since there are finitely many states, we can pick $d > 0$ uniformly in $h \in H$ such that $s^*(h)$ is d-robust in $G(v(s^*|(h, \cdot)))$, and thus, by Proposition 4, s^* is strongly robust.

Our construction of s^* is very similar to the standard ones such as those in Fudenberg and Maskin (1986) and in Abreu, Dutta, and Smith (1994), but differs with respect to state transition after simultaneous deviation from the equilibrium path by multiple players. Indeed, the standard construction usually ignores any simultaneous deviation or picks any arbitrary continuation strategy profile, for specifying off-path play after simultaneous deviation does not affect any of the players' incentives. In contrast, robustness may be affected by off-path continuation payoff profiles that can realize only after multiple players' simultaneous deviation. In our construction, players punish the deviator with the smallest name index to make sure that the current action profile is not only a Nash equilibrium of the augmented one-shot game, but also the unique action profile that survives iterative elimination of strictly dominated actions. See Appendix A.3.2 for the proof.

5 Extension to Dynamic Games

Our factorization result for repeated games (Theorem 1) relies crucially on a one-shot robustness principles (Proposition 4) which relates the dynamic strong robustness of an SPE s to the static strong robustness of action profiles $s(h)$ in $G(v(s|(h, \cdot)))$ for all h. In this section, we show that this one shot robustness principle extends to a general class of dynamic games with discounted stage payoffs. We also show by way of a counter example that a one-shot robustness principle does not hold anymore if we use robustness in the sense of KM.

5.1 Definition

We consider the following class of infinite-horizon dynamic games with observable actions. Let N be the set of players. Let H_t denote the set of all histories of length $t \geq 0$, which is defined recursively as follows. $H_0 = \{\emptyset\}$ and at each history $h \in H_{t-1}$ with $t \geq 1$, players choose actions $a_h = (a_{ih})_{i \in N} \in A_h = \prod_{i \in N} A_{ih}$ simultaneously. At the end of period t, players observe action profile a_h, obtain stage-game payoffs $g_{ih}(a_h)$, and move to the next history $(h, a_h) \in H_t$. Thus, we have $H_t = \{(h, a_h) \mid h \in H_{t-1},\ a_h \in A_h\}$. We write $H = \bigcup_{t \geq 0} H_t$. The total payoff for each player i is the sum of the stage-game payoffs discounted by δ and normalized by $1 - \delta$. We denote by $G_h = (N, (A_{ih}, g_{ih})_{i \in N})$ the stage game at history h and by $\Gamma = (N, H, (A_{ih}, g_{ih})_{i \in N, h \in H}, \delta)$ the entire dynamic game.

We allow for infinitely many players, but we assume that the number of available action profiles at each stage game G_h is bounded uniformly in h. That is, the number of "active" players who have multiple available actions at history h is bounded uniformly in h, and the number of player i's available actions at history h is bounded uniformly in i and h. Thus our class of dynamic games includes repeated games with finitely many players as well as overlapping generation games with finitely many "active" players at each period. Also, we assume that $|g| = \sup_{i,h,a_h} |g_{ih}(a_h)| < \infty$.

A strategy of player i is described by $s_i(h) \in A_{ih}$ for each $h \in H$. For each history $h \in H$, a strategy profile s induces continuation payoff $v_i(s|h)$ for player i. s^* is an SPE if $v_i(s^*|h) \geq v_i(s_i, s^*_{-i}|h)$ for every $h \in H$, $i \in N$, and s_i.

We perturb Γ and consider a collection $\mathbf{U} = \{U_h\}_{h \in H}$ of independent incomplete information games such that, for each $h \in H$, $U_h = (N, \Omega_h, P_h, (A_i, u_{ih}, Q_{ih})_{i \in N})$ embeds G_h, and $|u_{ih}(a)|$ is bounded by some $M > |g|$. At each history $h \in H$, $\omega_h \in \Omega_h$ is generated according to P_h. Each player i chooses $\sigma_i(h, \omega_h) \in \Delta(A_{ih})$ such that $\sigma_i(h, \cdot)$ is Q_{ih}-measurable. The total payoff is given by the discounted sum of stage-game payoffs $u_{ih}(a_h, \omega_h)$. We denote this game by $\Gamma_{\mathbf{U}}$.

A public strategy profile σ induces continuation payoff $v_i(\sigma|h)$ for player i. A public strategy profile σ^* is a *perfect public equilibrium* (*PPE*) if $v_i(\sigma^*|h) \geq v_i((\sigma_i, \sigma^*_{-i})|h)$ for every $h \in H$, $i \in N$, and σ_i.

Definition 10 (strong robustness for dynamic games)**.** An SPE s^* of Γ is *d-robust* for $d > 0$ if, for every $\eta > 0$, there exists $\varepsilon > 0$ such that, for every $\mathbf{U} = \{U_h\}$ of independent (ε, d)-elaborations of G_h, $\Gamma_{\mathbf{U}}$ has a PPE σ^* such that $P_h^{\sigma^*(h,\cdot)}(s^*(h)) > 1 - \eta$ for every $h \in H$.

An SPE s^* of Γ is *strongly robust* if it is d-robust for some $d > 0$.

Note that, in the above definition, we use not only time-dependent but also history-dependent elaborations even if Γ is a repeated game, a wider class of perturbations than what we use in Definition 4. Nevertheless, the one-shot robustness principle holds for dynamic games as below, which implies that the set of strongly robust equilibrium values for repeated games under Definition 10 is characterized in the same way as in Theorem 1.

Note also that our notion of dynamic robustness depends on how to represent total payoffs in terms of stage-game payoffs and the discount factor. Consider two dynamic games $\Gamma = (N, H, (A_{ih}, g_{ih})_{i \in N, h \in H}, \delta)$ and $\Gamma' = (N, H, (A_{ih}, g'_{ih})_{i \in N, h \in H}, \delta')$ such that $g'_{ih}(a) = (\delta/\delta')^{t-1} g_{ih}(a)$ and $\delta' > \delta$. Then, since period-t payoffs in Γ' converges to 0 as $t \to \infty$, our d-robustness notion for Γ' checks robustness to infinitely large payoff perturbations relative to the size of stage-game payoffs in future periods.

In this setup, we can state the one-shot robustness principle as follows.

Proposition 10 (one-shot robustness principle for dynamic games)**.**

(i) If $0 < d \leq M - |g|$ and s^* is a d-robust SPE of Γ, then $s^*(h)$ is a $(1-\delta)d$-robust equilibrium of $G_h(v(s^*|(h, \cdot)))$ for every $h \in H$.

(ii) If $d > 0$ and $s^*(h)$ is a d-robust equilibrium of $G_h(v(s^*|(h, \cdot)))$ for every $h \in H$, then s^* is a d'-robust SPE of Γ for every $d' < d$.

The proof is essentially identical to that of Proposition 4.

5.2 Failure of the one-shot robustness principle under weak robustness

The notion of strong robustness we develop in Section 2 rules out weak Nash equilibria. In particular it rules out unique correlated equilibria in mixed strategies, which are robust in the sense of KM. In this section we provide an example highlighting why this strengthening of robustness is necessary once we look at dynamic, infinite horizon games. More specifically we describe a game and an equilibrium of that game that is not dynamically robust in any reasonable sense, although its one-shot action profiles are weakly robust in all appropriately augmented stage games.

For any $T \in \mathbb{N}$ we consider the finite horizon overlapping generations game Γ_T defined as follows:

- Time is discrete, with $t \in \{1, \cdots, T\}$.
- At each period t, there are two active players denoted X_t, Y_t. Player X_t takes a decision $x_t \in \{0, 1\}$. Player Y_t takes a decision $y_t \in \{0, 1\}$.
- For every $t \in \{0, \cdots, T-1\}$, the payoffs to players X_t and Y_t are given by

	$y_t = 1$	$y_t = 0$
$x_t = 1$	$a - b\mathbb{E}y_{t+1}; 0$	$0; 1$
$x_t = 0$	$0; 1$	$1; 0$

  where $b > a$.

- At $t = T$, the payoffs to players X_T and Y_T are given by

	$y_T = 1$	$y_t = 0$
$x_T = 1$	$a - b\lambda_T; 0$	$0; 1$
$x_T = 0$	$0; 1$	$1; 0$

where $\lambda_T \in [0,1]$ is some parameter of the game that will be specified later.

Player Y_t gets payoffs only in period t while player X_t gets payoffs in period t and $t+1$. Such an overlapping generations game is described by parameters a, b and λ_T. Let us define the function f such that

$$\forall \lambda \in [0,1], \quad f(\lambda) = \min\left\{1, \frac{1}{1+a-\lambda b}\right\}.$$

Since $b > a$, the function f has three fixed points λ_L, λ_M and λ_H, that satisfy

$$0 < \lambda_L < \lambda_M < a/b < \lambda_H = 1.$$

Note that λ_M is unstable. Let us denote $G(\lambda)$ the stage game

	$y_t = 1$	$y_t = 0$
$x_t = 1$	$a - b\lambda; 0$	$0; 1$
$x_t = 0$	$0; 1$	$1; 0$

The following result holds.

Lemma 6 (SPEs of Γ_T).

(i) In every SPE of Γ_T player Y_t mixes and chooses $y_t = 1$ with probability $\lambda_t = f^{T-t}(\lambda_T)$.

(ii) If $b > a$ and $\lambda_T = \lambda_M$, then game Γ_T has a unique SPE such that for all t, player Y_t plays $y_t = 1$ with probability $\lambda_t = \lambda_M$ and player X_t plays $x_t = 1$ with probability $1/2$. This one-shot mixed strategy profile is the unique correlated equilibrium of the augmented stage game $G(\lambda_M)$.

Proof. The proof of the first result is by induction. Assume that at stage $t+1$, player Y_{t+1} chooses $y_{t+1} = 1$ with probability λ_{t+1}. Then players X_t and Y_t are playing the stage game

$G(\lambda_{t+1})$. If $\lambda_{t+1} \geq a/b$, then playing $x_t = 0$ is weakly dominant for player X_t and hence, player Y_t must be playing $y_t = 1$ with probability $1 = f(\lambda_{t+1})$. If $\lambda_{t+1} < a/b$ then game $G(\lambda_{t+1})$ has a unique mixed equilibrium in which player Y_t plays $y_t = 1$ with probability $f(\lambda_{t+1})$. The second result follows from the fact that λ_M is a fixed-point of f that's strictly below 1. ☐

We use game Γ_T to illustrate the fact that when an equilibrium involves a long sequence of mixed action profiles, then small elaborations on the game payoffs can have far reaching consequences. In particular, given the game Γ_T described by a, b and λ, we consider the elaboration Γ_T^ε in which player X_{T-1} expects that his period T payoff will be equal to $-bx_{T-1}y_{T-1}y_T$ with probability $1-\varepsilon$ and equal to 1 with probability ε.

Proposition 11 (sensitivity of mixed equilibria to perturbations)**.** *Consider game* Γ_T *with* $b > a$ *and* $\lambda_T = \lambda_M$. *The following results hold.*

(i) For any ε, *game* Γ_T^ε *has a unique subgame-perfect equilibrium in which at any time* $t < T$ *player* Y_t *plays* $y_t = 1$ *with probability* $\lambda_t^\varepsilon = f^{T-t}(\lambda_M^\varepsilon)$, *where* $\lambda_M^\varepsilon = (1-\varepsilon)\lambda_M - \varepsilon/b$.

(ii) There exists T *large enough such that* $|\lambda_1^\varepsilon - \lambda_1| > |\lambda_M - \lambda_L|/2 > 0$.

Proof. The proof of point (i) is identical to that of Lemma 6. Point (iv) is a consequence of the fact that λ_M is an unstable fixed point of f. Hence, since $\lambda_M^\varepsilon < \lambda_M$, there exists T such that $|f^{T-1}(\lambda_M^\varepsilon) - \lambda_L| < |\lambda_M - \lambda_L|/2$. Hence, we get that $|\lambda_1^\varepsilon - \lambda_1| > |\lambda_M - \lambda_L|/2 > 0$. ☐

By point (ii) of Proposition 6, game Γ^* is such that at every stage, players play the unique correlated equilibrium of the augmented stage game. However, by point (ii) of Proposition 11, the unique SPE of game Γ^* is not dynamically robust. Thus the one-shot robustness principle does not hold under KM's definition of robustness. This shows that when an equilibrium involves a long sequence of mixed action profiles, small differences in future payoffs can be greatly magnified over time. This does not happen when an equilibrium is in uniformly strict strategies.

Note that while point (*ii*) of Proposition 11 applies to families of games with increasing length T, sequences of games Γ_T can simply be regrouped in a single infinite horizon game Γ^* in which players play game Γ_T in the interval of time $\{T(T-1)/2+1,\dots,T(T+1)/2\}$.

6 Conclusion

In this paper we introduced d-robustness and strong robustness as notions of robustness to incomplete information that are well adapted to the analysis of repeated and dynamic games. These notions of robustness are more restrictive than, but closely related to robustness in the sense of KM. Specifically, they only depart from KM's notion of robustness by allowing perturbed payoffs to be slightly different (rather than identical) from those of the original game with large probability.

Our main theoretical results show that the dynamic strong robustness of SPEs can be related to the static strong robustness of their implied one-shot action profiles in appropriately augmented stage games. More specifically, we prove a factorization result for strongly robust SPEs in repeated games that parallels the approach of APS (1990). We also prove a one-shot robustness principle analogous to the one-shot deviation principle for a general class of dynamic games with discounted flow payoffs. We also show that this one-shot robustness principle would not hold if we used KM's weaker notion of robustness.

These characterizations of strongly robust SPEs are useful for applied work. We illustrate this by computing explicitly the set of strongly robust SPEs in the repeated Prisoners' Dilemma. We show that cooperation can be robustly sustained if and only if the outcome (*Cooperate*, *Defect*) is enforceable in equilibrium under complete information. In that setting, we highlight that grim-trigger strategies is not the most robust way to sustain cooperation, and that stage dominant strategies that punish only the deviator upon unilateral deviation may be more effective. We also show that a folk theorem in robust strategies holds, which means that robustness considerations only restrict equilibrium values for fixed discount factors.

We believe that taking into account robustness issues in dynamic games can yield useful insights in many applied settings. In particular, when robustness is carefully considered, a player's payoff when others deviate will affect the robustness of strategies significantly, which can lead to richer analysis and comparative statics (see, for instance, Chassang and Padro i Miquel (2008) for an application).

With respect to applications, one advantage of our approach is that it is quite general and the main characterization results are fairly intuitive. To some extent, this is due to the specific class of elaborations against which we test robustness. In particular, considering more general perturbations (for instance, allowing for correlation between signals, or doing away with the common prior assumption) would make the robustness results stronger, while considering more specific perturbations such as global games would make the fragility results stronger. While we think that the set of perturbations we consider is natural, in particular for repeated games, we also believe that it would be a worthwhile exercise to explore robustness against such classes of elaborations.

Appendix

A.1 Proofs for Section 2

A.1.1 Proof of Proposition 1

The proof is by contradiction, and follows the structure of KM (Proposition 3.2). It uses Lemmas 7 and 8, which are of independent interest and given below.

Definition 11 (canonic normalization)**.** Consider an incomplete information game $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ and an strategy profile α^* of U. We call $\tilde{U} = (N, \tilde{\Omega}, \tilde{P}, (A_i, \tilde{u}_i, \tilde{Q}_i)_{i \in N})$ the *canonic normalization of* U *with respect to* α^* if

(i) $\tilde{\Omega} = A$

(ii) For $\tilde{\omega} = a \in \tilde{\Omega}$, $\tilde{P}(\tilde{\omega}) = P^{\alpha^*}(a)$.

(iii) $\tilde{Q}_i = \{\{a_i\} \times A_{-i} \mid a_i \in A_i\}$

(iv) If $\tilde{\omega} \in \{a_i\} \times A_{-i}$, then

$$\tilde{u}_i(a_i', a_{-i}, \tilde{\omega}) = \frac{1}{\sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega)} \sum_{\omega \in \Omega} u_i(a_i', a_{-i}, \omega)\alpha_i^*(\omega)(a_i)P(\omega)$$

whenever the denominator on the right-hand side is nonzero.[5]

We say that $\tilde{\alpha}_i^*$ is the *truthtelling strategy* in $\tilde{U}$ if $\tilde{\alpha}_i^*(\tilde{\omega})(a_i) = 1$ whenever $\tilde{\omega} \in \{a_i\} \times A_{-i}$.

Lemma 7 (canonic normalization with respect to a Bayesian-Nash equilibrium)**.** *Let $\tilde{U}$ be the canonic normalization of U with respect to α^*, and let $\tilde{\alpha}^*$ be the truthtelling strategy profile in $\tilde{U}$. Then we have the following.*

(i) α^* and $\tilde{\alpha}^*$ induce the same distribution on A: $P^{\alpha^*} = \tilde{P}^{\tilde{\alpha}^*}$.

(ii) If α^* is a Bayesian-Nash equilibrium of U, then $\tilde{\alpha}^*$ is a Bayesian-Nash equilibrium of $\tilde{U}$.

(iii) If U is an (ε, d)-elaboration of G, then $\tilde{U}$ is an $(n\varepsilon^{1/2}, d_\varepsilon)$-elaboration of G, where $d_\varepsilon = d + \varepsilon^{1/2}(|g| + M)$.

Proof. (i) and (ii) follow directly from the definition of the canonic normalization.

For (iii), let

$$\Omega_d = \{\omega \in \Omega : |u_i(\cdot, \omega') - g_i| \le d \text{ for all } i \in N \text{ and } \omega' \in Q_i(\omega)\}.$$

Since U is an (ε, d)-elaboration, $P(\Omega_d) > 1 - \varepsilon$. Let B_i be the set of actions $a_i \in A_i$ such that

$$\sum_{\omega \in \Omega_d} \alpha_i^*(\omega)(a_i)P(\omega) \ge (1 - \varepsilon^{1/2}) \sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega),$$

---

[5] The denominator is nonzero $\tilde{P}$-almost surely.

and let $B = \prod_{i \in N} B_i$. We will show that, in $\tilde{U}$, every player i knows that $\tilde{u}_i$ is close enough to g_i on the event of B and $\tilde{P}(B)$ is high enough.

If $\tilde{\omega} = a \in B$, then, for each $i \in N$ and each $\tilde{\omega}' \in \{a_i\} \times A_{-i}$,

$$\begin{aligned}|\tilde{u}_i(\cdot, \tilde{\omega}') - g_i| &\le \frac{1}{\sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega)} \sum_{\omega \in \Omega} |u_i(\cdot, \omega) - g_i| \alpha_i^*(\omega)(a_i)P(\omega) \\ &\le d + \frac{1}{\sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega)} \sum_{\omega \in \Omega \setminus \Omega_d} |u_i(\cdot, \omega) - g_i| \alpha_i^*(\omega)(a_i)P(\omega) \\ &\le d + \varepsilon^{1/2}(|g| + M).\end{aligned}$$

For each $a_i \in A_i \setminus B_i$, we have

$$\sum_{\omega \in \Omega \setminus \Omega_d} \alpha_i^*(\omega)(a_i)P(\omega) > \varepsilon^{1/2} \sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega).$$

Adding the both sides for all $a_i \in A_i \setminus B_i$, we have

$$\begin{aligned}\varepsilon > P(\Omega \setminus \Omega_d) &\ge \sum_{\omega \in \Omega \setminus \Omega_d} \sum_{a_i \in A_i \setminus B_i} \alpha_i^*(\omega)(a_i)P(\omega) \\ &> \varepsilon^{1/2} \sum_{\omega \in \Omega} \sum_{a_i \in A_i \setminus B_i} \alpha_i^*(\omega)(a_i)P(\omega) = \varepsilon^{1/2} \tilde{P}((A_i \setminus B_i) \times A_{-i}),\end{aligned}$$

thus $\tilde{P}((A_i \setminus B_i) \times A_{-i}) < \varepsilon^{1/2}$. Thus, $\tilde{P}(B) \ge 1 - \sum_i \tilde{P}((A_i \setminus B_i) \times A_{-i}) > 1 - n\varepsilon^{1/2}$. □

The point of canonic normalizations is that given a set of players and an action space they form a finite dimensional class of games.

Lemma 8 (locally unique equilibrium)**.** *If a^* is the unique correlated equilibrium of G and a^* is a strict Nash equilibrium, then there exists $d > 0$ such that a^* is the unique Bayesian-Nash equilibrium of any $(0, d)$-elaboration of G.*

Proof. The proof is by contradiction. Assume that for any $d > 0$ there exists a $(0, d)$-elaboration $U_d = (N, \Omega_d, P_d, (A_i, u_{id}, Q_{id})_{i \in N})$ of G such that U_d has a Bayesian-Nash equilibrium α_d such that $P_d^{\alpha_d}(a^*) \neq 1$. Since the canonic form of a $(0,d)$-elaboration is also a $(0,d)$-elaboration by Lemma 7, without loss of generality, we can assume that U_d takes a canonic form with respect to α_d, and that α_d is the truthfull telling stratregy profile.

Since $P_d^{\alpha_d}(a^*) \neq 1$ the conditional distribution over $A \setminus \{a^*\}$ is well defined by:

$$\forall a \in A \setminus \{a^*\}, \quad \mu_d(a) = \frac{P_d^{\alpha_d}(a)}{P_d^{\alpha_d}(A \setminus \{a^*\})}.$$

Since α_d is a Bayesian-Nash equilibrium of U_d, we must have that for all $a_i \in A_i \setminus \{a_i^*\}$ and all $a'_i \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} \mu_d(a_i, a_{-i})[u_{id}(a_i, a_{-i}, \omega) - u_{id}(a'_i, a_{-i}, \omega)] \geq 0 \tag{1}$$

whenever $\omega \in \{a_i\} \times A_{-i}$. As d goes to 0, payoffs u_d converge to the payoff g of game G. Since $\mu_d \in \Delta(A \setminus \{a^*\})$, which is compact, as d goes to 0, we can extract a sequence of probability distributions μ_d over $A \setminus \{a^*\}$ that all satisfy family of inequalities (1) for some $(0,d)$-elaboration and converge to a distribution $\mu_0 \in \Delta(A \setminus \{a^*\})$. By continuity, we have that, for all $a_i \in A_i \setminus \{a_i^*\}$ and all $a'_i \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} \mu_0(a_i, a_{-i})[g_i(a_i, a_{-i}) - g_i(a'_i, a_{-i})] \geq 0. \tag{2}$$

We now use distribution μ_0 to build a correlated equilibrium of G distinct from a^*. For $\lambda > 0$ define $\mu \in \Delta(A)$ by $\mu(a^*) = (1 - \lambda)$ and $\mu(a) = \lambda \mu_0(a)$. It follows from the family of inequalities (2) and the fact that a^* is a strict Nash equilibrium of G that for $\lambda > 0$ small enough, μ is a correlated equilibrium of G. Indeed, since a^* is a strict Nash equilibrium, for λ small enough, a player getting message a_i^* will find it profitable to play a_i^*. Similarly, by (2), a player that gets message $a_i \neq a_i^*$ also finds it profitable to obey.

This contradicts the fact that a^* is the unique correlated equilibrium of G. Hence, there must exist $d > 0$ such that a^* is the unique Bayesian-Nash equilibrium of all $(0,d)$-elaborations of G. □

Proof of Proposition 1. By Lemma 8, we know that there exists $d > 0$ such that a^* is the unique Bayesian-Nash equilibrium of any $(0, d)$-elaboration of G. Assume that there exists $\eta > 0$ such that for all $\varepsilon > 0$, there exists an (ε, d)-elaboration $U_\varepsilon = (N, \Omega_\varepsilon, P_\varepsilon, (A_i, u_{i\varepsilon}, Q_{i\varepsilon})_{i\in N})$ of G such that any Baysian-Nash equilibrium α_ε of U_ε satisfies $P_\varepsilon^{\alpha_\varepsilon}(a^*) \leq 1 - \eta$. Pick such an equilibrium α_ε. Let $\tilde{U}_\varepsilon$ be the canonic normalization of U_ε with respect to α_ε. By Lemma 7, the truthtelling strategy profile α^* is a Bayesian-Nash equilibrium of $\tilde{U}_\varepsilon$, $\tilde{P}_\varepsilon^{\alpha^*}(a^*) \leq 1 - \eta$, and $\tilde{U}_\varepsilon$ is an $(n\varepsilon^{1/2}, d_\varepsilon)$-elaboration of G, where $d_\varepsilon = d + \varepsilon^{1/2}(|g| + M)$.

By an argument identical to KM (Lemma 3.4), we have that the truthtelling strategy profile α^* is an $Mn\varepsilon^{1/2}$-equilibrium[6] of the game $\hat{U}_\varepsilon$ identical to $\tilde{U}_\varepsilon$ except that $\hat{u}_\varepsilon(\cdot, \omega) = g$ whenever $|\tilde{u}_\varepsilon(\cdot, \omega) - g(\cdot)| > d_\varepsilon$. Note that game $\hat{U}_\varepsilon$ is a $(0, d_\varepsilon)$-elaboration of G with state space A. Now take ε to 0. Because the set of incomplete-information games with state space A and uniformly bounded payoff functions is compact, we can extract a converging sequence of $(0, d_\varepsilon)$-elaborations $\hat{U}_\varepsilon$ such that $\hat{P}_\varepsilon^{\alpha^*}(a^*) \leq 1 - \eta$. Denote $\hat{U}_0$ the limit of the sequence.

By continuity, $\hat{U}_0$ is a $(0, d)$-elaboration of G, the truthtelling strategy profile is a Bayesian-Nash equilibrium of $\hat{U}_0$, and $\hat{P}_0^{\alpha^*}(a^*) \leq 1 - \eta$. This contradicts the fact that a^* is the unique Bayesian-Nash equilibrium of all $(0, d)$-elaborations. Hence for any $\eta > 0$, there exists $\varepsilon > 0$ such that every (ε, d)-elaboration of G admits an equilibrium that induces a^* with probability larger than $1 - \eta$. □

A.1.2 Proof of Proposition 3

The proof is in two steps. The first is to show that, whenever a^* is strictly $\mathbf{p}$-dominant in G, there exists $d > 0$ small enough that a^* is approximately $\mathbf{p}$-dominant in all $(0, d)$-elaborations of G. The second step exploits the critical path result of KM (Proposition 4.2) to establish that (ε, d)-elaborations admit equilibria uniformly close to a^*.

Pick $\mathbf{p}' \gg \mathbf{p}$ such that $\sum_{i\in I} p'_i < 1$. Since a^* is strictly $\mathbf{p}$-dominant, for all $i \in N$,

---

[6] Where ε-equilibrium is defined in the usual way, i.e. players best reply to the strategies of others up to a loss of ε.

$a_i \in A_i \setminus \{a_i^*\}$ and $\lambda \in \Delta(A_{-i})$ such that $\lambda(a_{-i}^*) \geq p_i'$,

$$\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i(a_i^*, a_{-i}) > \sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i(a_i, a_{-i}).$$

Because the set of probability distributions $\lambda \in \Delta(A_{-i})$ such that $\lambda(a_{-i}^*) \geq p_i'$ is compact and because the functions that associate λ to numbers $\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i})[g_i(a_i^*, a_{-i}) - g_i(a_i, a_{-i})]$ are continuous and there is a finite number of them, it follows that there exists $d > 0$ such that for all $i \in N$, $a_i \in A_i \setminus a_i^*$ and $\lambda \in \Delta(A_{-i})$ with $\lambda(a_{-i}^*) \geq p_i'$,

$$\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i(a_i, a_{-i}) + 2d.$$

This implies that, for all for all $i \in N$, $a_i \in A_i \setminus \{a_i^*\}$, $\lambda \in \Delta(A_{-i})$ s.t. $\lambda(a_{-i}^*) \geq p_i$,

$$\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i, a_{-i})$$

whenever $|g' - g| \leq d$.

For any (ε, d)-elaboration $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ of G, let us define

$$\Omega_d = \{\omega \in \Omega : |u_i(\cdot, \omega') - g_i| \leq d \text{ for all } i \in N \text{ and } \omega' \in Q_i(\omega)\}.$$

By definition of (ε, d)-elaborations, we have that $P(\Omega_d) > 1 - \varepsilon$. As in KM we are now interested in the set of states $C^{\mathbf{p}'}(\Omega_d)$, i.e. the set of states where event Ω_d is common $\mathbf{p}'$-belief. Proposition 4.2 (the critical path result) of KM implies that

$$P(C^{\mathbf{p}'}(\Omega_d)) \geq 1 - (1 - P(\Omega_d)) \frac{1 - \min_{i \in N} p_i'}{1 - \sum_{i \in N} p_i'}. \tag{3}$$

Since $\sum_{i \in N} p_i' < 1$, for any $\eta > 0$ there exists ε small enough such that for any (ε, d)-elaboration U, $P(C^{\mathbf{p}'}(\Omega_d)) > 1 - \eta$. Whenever the state of the world ω belongs to $C^{\mathbf{p}'}(\Omega_d)$, there is common $\mathbf{p}'$-belief that payoffs are those of a $(0, d)$-elaboration. Since a^* is $\mathbf{p}'$-

dominant in any $(0, d)$-elaboration, Lemma 5.2 of KM implies that U admits an equilibrium α^* such that $\alpha_i^*(\omega)(a_i^*) = 1$ for all $\omega \in C^{\mathbf{p}'}(\Omega_d)$. Equilibrium α^* satisfies $P^{\alpha^*}(a^*) \geq P(C^{\mathbf{p}'}(\Omega_d)) > 1 - \eta$, which concludes the proof.

A.2 Proofs for Section 3

A.2.1 Proof of Proposition 4

For (i), by Lemma 2, $s^*(h)$ is a $2(1-\delta)d$-strict equilibrium of $G(v(s^*|(h, \cdot)))$ for every $h \in H$.

Pick any $t^0 \geq 1$ and $h^0 \in H_{t^0-1}$. Let $a^0 = s^*(h^0)$, $v^0 = v(s^*|h^0)$, and $w^0 = v(s^*|(h^0, \cdot))$. We want to show that a^0 is $(1 - \delta)d$-robust in $G(w^0)$. That is, for every $\eta > 0$, there exists $\varepsilon > 0$ such that every $(\varepsilon, (1 - \delta)d)$-elaboration of $G(w^0)$ has a Bayesian-Nash equilibrium such that $P_{t^0}^{\sigma^*(h^0,\cdot)}(a^0) > 1 - \eta$. Fix any $\eta > 0$. Without loss of generality, we can assume η to be small enough so that

- for any $t^1 > t^0$, $h^1 \in H_{t^1-1}$, $\mathbf{U} = \{U_t\}$ such that U_t is identical to G for all $t > t^1$, and strategy profile σ of $\Gamma_{\mathbf{U}}$, if $P_t^{\sigma^*(h,\cdot)}(s^*(h)) > 1 - \eta$ for all $h \in H_{t-1}$ with $t > t^1$, then

$$|v_{t^1+1}(\sigma|(h^1, \cdot)) - v(s^*|(h^1, \cdot))| < (1 - \delta)d,$$

- if a^* is a $2(1 - \delta)d$-strict equilibrium of G, then G does not have any other equilibrium in the η-neighborhood of a^*.

We say that $=\{U_t\}$ is a "one-shot" sequence if U_t is identical to G for all $t \neq t^0$. Since s^* is d-robust, there exists $\varepsilon > 0$ such that, for every "one-shot" sequence $\mathbf{U}$ with an arbitrary (ε, d)-elaboration U_{t^0}, $\Gamma_{\mathbf{U}}$ has a PPE σ^* such that $P_t^{\sigma^*(h,\cdot)}(s^*(h)) > 1 - \eta$ for every $t \geq 1$ and $h \in H_{t-1}$.

For any $t > t_0$ and $h \in H_{t-1}$, we show that $P_t^{\sigma^*(h,\cdot)}(s^*(h)) = 1$. Since $\sigma^*(h', 0)$ puts probability $1-\eta$ on $s^*(h')$ for all $h' \in H_{t'-1}$ with $t' > t$, we have $|v_{t+1}(\sigma^*|(h, \cdot)) - v(s^*|(h, \cdot))| < (1-\delta)d$. Then, since $s^*(h)$ is a $2(1-\delta)d$-strict equilibrium in $G(v(s^*|(h, \cdot)))$, it is a $2(1-\delta)^2d$-

strict equilibrium in $G(v_{t+1}(\sigma^*|(h,\cdot)))$. Thus $U_t(v_{t+1}(\sigma^*|(h,\cdot)))$ does not have any other equilibrium in the η-neighborhood of $s^*(h)$. Thus $P_t^{\sigma^*(h,\cdot)}(s^*(h)) = 1$.

Then we have $v_{t^0+1}(\sigma^*|(h^0,\cdot)) = v(s^*|(h^0,\cdot)) = w^0$ and $\sigma^*(h^0,\cdot)$ is a Bayesian-Nash equilibrium of an $(\varepsilon, (1-\delta)d)$-elaboration of $G(w^0)$ that puts probability more than $1-\eta$ on a^0.

For (ii), we will show that, for every $\eta > 0$, there exists $\varepsilon > 0$ such that, for every sequence $\mathbf{U} = \{U_t\}$ of (ε, d')-elaborations of G, $\Gamma_{\mathbf{U}}$ has a PPE σ^* such that $P_t^{\sigma^*(h,\cdot)}(s^*(h)) \geq 1-\eta$ for every $h \in H$.

Pick $\bar{\varepsilon} > 0$ and $\bar{\eta} > 0$ such that, for every $t \geq 1$, $h \in H_{t-1}$, $\mathbf{U} = \{U_t\}$ of (ε, d')-elaborations of G, and strategy profile σ of $\Gamma_{\mathbf{U}}$, if $P_{t'}^{\sigma(h',\cdot)}(s^*(h')) > 1-\bar{\eta}$ for all $h \in H_{t'-1}$ with $t' > t$, then $|v_{t+1}(\sigma|(h,\cdot)) - v(s^*|(h,\cdot))| < d'$. Pick $d'' > 0$ such that $d' + \delta d'' < d$.

Fix any $\eta > 0$. We can assume without loss of generality that $\eta < \bar{\eta}$.

For each $a \in A$, let $\mathbf{v}(a) \subseteq \mathbb{R}^{n|A|}$ be the set of contingent continuation payoff profiles $v(s^*|(h,\cdot))$ such that $s^*(h) = a$. Since $\mathbf{v}(a)$ is bounded, there exists a finite set of histories, $H(a)$, such that $s^*(h) = a$ for any $h \in H(a)$ and any point in $\mathbf{v}(a)$ is within d'' distance of $v(s^*|(h,\cdot))$ for some $h \in H(a)$.

For each $a \in A$ and $h \in H(a)$, since a is d-robust in $G(v(s^*|(h,\cdot)))$, there exists $\varepsilon_h > 0$ such that every (ε_h, d)-elaboration of $G(v(s^*|(h,\cdot)))$ has a Bayesian-Nash equilibrium that puts probability larger than $1-\eta$ on a. Let $\varepsilon = \min(\bar{\varepsilon}, \min_{a\in A}\min_{h\in H(a)} \varepsilon_h)$. Then, for every $h \in H$, every (ε, d')-elaboration of $G(v(s^*|(h,\cdot)))$ has a Bayesian-Nash equilibrium that puts probability larger than $1-\eta$ on $s^*(h)$. Note that ε is chosen uniformly in $h \in H$.

Fix any sequence $\mathbf{U} = \{U_t\}_{t\in\mathbb{N}}$ of (ε, d')-elaborations of G. Now we will construct a PPE σ^* of $\Gamma_{\mathbf{U}}$ as follows.

For each $T < \infty$, consider the "truncated" sequence $\mathbf{U}^T = \{U_t^T\}_{t\in\mathbb{N}}$ of elaborations, i.e., $U_t^T = U_t$ for $t \leq T$ and U_t^T is identical to G for all $t > T$. We can backwardly construct a PPE σ^T of $\Gamma_{\mathbf{U}^T}$ as follows.

- For $h \in H_{t-1}$ with $t > T$, we set $\sigma_i^T(h,0)(s_i^*(h)) = 1$ for every $i \in N$.

- For $h \in H_{t-1}$ with $t \leq T$, we set $\sigma^T(h, \cdot)$ equal to a Bayesian-Nash equilibrium of $U_t(v_{t+1}(\sigma^T|(h, \cdot)))$ that puts probability larger than $1-\eta$ on $s^*(h)$. Such a Bayesian-Nash equilibrium exists because $\sigma^T(h', \cdot)$ puts probability larger than $1-\eta$ on $s^*(h')$ for all $h' \in H_{t'-1}$ with $t' > t$ and thus $|v_{t+1}(\sigma^T|(h, \cdot)) - v(s^*|(h, \cdot))| < d'$. Thus $U_t(v_{t+1}(\sigma^T|(h, \cdot)))$ is an (ε, d')-elaboration of $G(v(s^*|(h, \cdot)))$.

Since the set of all public strategy profiles is a compact metrizable space in the product topology, let σ^* be the limit of $\{\sigma^T\}_{T\in\mathbb{N}}$ (take a subsequence if necessary). That is, $\sigma^T(h, \omega_t) \to \sigma^*(h, \omega_t)$ as $T \to \infty$ pointwisely for all $t \geq 1$, $h \in H_{t-1}$, and $\omega_t \in \Omega_t$. By the upper hemicontinuity of PPEs with respect to payoff perturbations, σ^* is a PPE of $\Gamma_{\mathbf{U}}$. Also, $\sigma^*(h, \cdot)$ puts probability at least $1-\eta$ on $s^*(h)$ for every $h \in H$.

A.2.2 Proof of Proposition 5

Our proof of Proposition 5 is based on Lemmas 9 and 10.

Lemma 9. *If* $0 < d \leq M - |g|$ *and* $\tilde{s}^*$ *is a d-robust SPE of* $\tilde{\Gamma}_G$ *with public randomization with* $0 < d \leq M - |g|$, *then* $\tilde{s}^*(h_{t-1}, \theta^t)$ *is a* $(1-\delta)d$*-robust equilibrium of* $G(\tilde{v}(s^*|(h_{t-1}, \theta^t, \cdot)))$ *for every* $t \geq 1$, $h_{t-1} \in H_{t-1}$, *and* $\theta^t \in [0, 1)^t$.

We omit the proof as it is essentially the same as the proof of Proposition 4 (i).

For a sequence $\{R_t\}$ of partitions on $[0, 1)$, we say that $\tilde{s}$ is *adapted to* $\{R_t\}$ if $\tilde{s}(h, \cdot)$ is $R_1 \otimes \cdots \otimes R_t$-measurable for every $t \geq 1$ and $h_{t-1} \in H_{t-1}$, i.e., $\tilde{s}(h_{t-1}, \theta^t) = \tilde{s}(h_{t-1}, \phi^t)$ whenever $\phi_\tau \in R_\tau(\theta_\tau)$ for all $\tau \leq t$. A public strategy profile $\tilde{s}$ is *simple* if it is adapted to some sequence of finite Borel-measurable partitions on $[0, 1)$.

Lemma 10. *If* $d > 0$, $\tilde{s}^*$ *is simple, and* $\tilde{s}^*(h_{t-1}, \theta^t)$ *is a* $(1-\delta)d$*-robust equilibrium of* $G(\tilde{v}(s^*|(h_{t-1}, \theta^t, \cdot)))$ *for every* $t \geq 1$, $h_{t-1} \in H_{t-1}$, *and* $\theta^t \in [0, 1)^t$, *then* $\tilde{s}^*$ *is a* d'*-robust SPE of* $\tilde{\Gamma}_G$ *for every* $d' < d$.

Again, the proof is very similar to the proof of Proposition 4 (ii). One difference is in the last step, where we construct a PPE of the "truncated" game, and then take the limit of these PPEs to obtain a PPE of the original sequence of elaborations. Here, because $\tilde{s}^*$ is assumed to be adapted to some sequence $\{R_t\}$ of finite Borel-measurable partitions on $[0,1)$, we can construct a PPE $\tilde{\sigma}^T$ of $\tilde{\Gamma}_{\mathbf{U}^T}$ truncated at period T such that $\tilde{\sigma}^T(h,\cdot,\omega_t)$ is $R_1 \otimes \cdots \otimes R_t$-measurable for every $t \geq 1$, $h_{t-1} \in H_{t-1}$, and $\omega_t \in \Omega_t$. Since the set of all $\{R_t\}$-adapted public strategy profiles is a compact metrizable space in the product topology, there exists $\tilde{\sigma}^*$ such that $\tilde{\sigma}^T(h_{t-1},\theta^t,\omega_t) \to \tilde{\sigma}^*(h_{t-1},\theta^t,\omega_t)$ as $T \to \infty$ pointwisely for all $t \geq 1$, $h_{t-1} \in H_{t-1}$, $\theta^t \in [0,1)^t$, $\omega_t \in \Omega_t$, and uniformly in θ^t on each cell of $R_1 \otimes \cdots \otimes R_t$ (take a subsequence if necessary). Then σ^* is a PPE of $\tilde{\Gamma}_{\mathbf{U}}$.

Proof of Proposition 5. For each $v \in \tilde{\mathcal{V}}^{\mathrm{rob}}$, let $\tilde{s}^*$ be a d-robust SPE of the repeated game of G that attains v with $d \leq M - |g|$. Then, by Lemma 9, $\mathcal{V}^* = \{\tilde{v}(\tilde{s}^*|h_{t-1},\theta^t) \in \mathbb{R}^n \mid t \geq 1, h_{t-1} \in H_{t-1}, \theta^t \in [0,1)^t\}$ satisfies $\mathcal{V}^* \subseteq B^{(1-\delta)d}(\operatorname{co}\mathcal{V}^*)$, and thus $\operatorname{co}\mathcal{V}^*$ is self-generating with respect to $\tilde{B}^{(1-\delta)d}$. So we have $v \in \operatorname{co}\mathcal{V}^* \subseteq \tilde{\mathcal{V}}^{(1-\delta)d}$. Thus $\tilde{\mathcal{V}}^{\mathrm{rob}} \subseteq \bigcup_{d>0} \tilde{\mathcal{V}}^{(1-\delta)d} = \bigcup_{d>0} \tilde{\mathcal{V}}^d$.

For each $v \in \tilde{\mathcal{V}}^d$, since $\tilde{\mathcal{V}}^d$ is self-generating with respect to $\tilde{B}^d$, there exist $\lambda(v,k) \geq 0$, $a(v,k) \in A$, and $w(v,k,\cdot)\colon A \to \tilde{\mathcal{V}}^d$ for $k = 1,\ldots,K(v)$ such that $\sum_{k=1}^{K(v)} \lambda(v,k) = 1$, $w(v,k,\cdot)$ enforces $a(v,k)$ d-robustly for every k, and $v = \sum_k \lambda(v,k)[(1-\delta)g(a(v,k)) + \delta w(v,k,a(v,k))]$. For each $\theta \in [0,1)$, let $k(v,\theta)$ be k such that $\sum_{l=1}^{k-1} \lambda(v,l) \leq \theta < \sum_{l=1}^{k} \lambda(v,l)$. Pick any $v \in \tilde{\mathcal{V}}^d$ and we construct $\tilde{s}^*$ recursively as follows: For each θ_1, let $\tilde{s}^*(\theta_1) = a(v,k(v,\theta_1))$. For each $a_1 \in A$, and $(\theta_1,\theta_2) \in [0,1)^2$, let $\tilde{s}^*(a_1,\theta_1,\theta_2) = a(w(v,k(v,\theta_1),a_1),k(w(v,k(v,\theta_1),\theta_2))$, and so on... By construction, $\tilde{s}^*$ is simple and, for every $t \geq 1$, and $h_{t-1} \in H_{t-1}$, and $\theta^t \in [0,1)^t$, $\tilde{s}^*(h_{t-1},\theta^t)$ is d-robust in $G(v(s^*|(h_{t-1},\theta^t,\cdot)))$. Then, by Lemma 10, s^* is $d/2$-robust in the repeated game of G and thus $v \in \tilde{\mathcal{V}}^{\mathrm{rob}}$. Thus $\tilde{\mathcal{V}}^d \subseteq \tilde{\mathcal{V}}^{\mathrm{rob}}$.

The proof of the algorithm part is the same as the proof of Theorem 1 for repeated games without public randomization. $\square$

A.3 Proofs for Section 4

A.3.1 Proof of Lemma 5

We begin with the following lemma.

Lemma 11 (equilibrium Pareto frontier of games in $\mathcal{G}_{DC/CC}$). *Pick a game G in $\mathcal{G}_{DC/CC}$. The following hold.*

(i) The Pareto frontier of G is self-enforcing.

(ii) Outcome DD is never used on the equilibrium path of a Pareto efficient PPE.

(iii) Pareto efficient equilibria can be chosen such that outcome CC is used permanently once it is used.

(iv) The Pareto frontier is self-enforced by equilibria that never use DD in or out of equilibrium.

Proof. From Stahl's characterization, we know that the set of equilibrium values $\mathcal{V}^{\text{SPE}}$ of $\tilde{\Gamma}_G$ takes the form $\mathcal{V}^{\text{SPE}} = \text{co}\{(0,0),(1,1),(0,\rho),(\rho,0)\}$ where $\rho \geq 1$. We begin with point (i). Pick a Pareto efficient equilibrium s^* and denote H_{s^*} the corresponding set of equilibrium histories. Since s^* is Pareto efficient, it must be that for all $h \in H_{s^*}$, the continuation equilibrium is Pareto efficient (otherwise replacing the continuation by a dominating equilibrium would improve on s^*). In what follows, we modify s^* so that even off of the equilibrium path only Pareto efficient continuations are used. This is possible because points $(0,\rho)$ and $(\rho,0)$ belong to the equilibrium Pareto frontier. Consider the equilibrium $\hat{s}^*$ that coincides with s^* over H_{s^*}, but such that, whenever player 1 deviates, continuation values are $(0,\rho)$, and whenever player 2 deviates alone, continuation values are $(\rho,0)$. Since 0 is the minmax value, the fact that s^* is an equilibrium implies that $\hat{s}^*$ is also an equilibrium. This shows that the equilibrium Pareto frontier of $\mathcal{G}_{DC/CC}$ is self-enforcing.

Let us turn to point (ii). Consider a Pareto efficient equilibrium s^*. If there is an equilibrium history h such that $s^*(h) = DD$, then, the strategy profile $\hat{s}^*$ obtained by skipping history h and instead playing as if the next period had already been reached is also an equilibrium and dominates s^*. Hence on the equilibrium path, action DD should never be taken.

We now proceed with point (iii). From point (i), we know that the equilibrium Pareto frontier of $\mathcal{G}_{DC/CC}$ is self enforcing. Since we have public randomization, this means that the set of Pareto efficient values can be generated using only extreme points of the Pareto frontier. This is the bang-bang property of Abreu, Pearce and Staccheti (1990). There are three such extreme points, $(0, \rho)$, $(\rho, 0)$ and $(1, 1)$. Because $(1, 1)$ is not the weighted sum of action profiles other than CC, this means that in the equilibrium that sustains value $(1, 1)$, outcome CC is played permanently on the equilibrium path. Inversely, when values $(0, \rho)$ must be delivered, then the current action profile must be DC, (otherwise player 1 would get strictly positive value). Inversely, when values $(\rho, 0)$ must be delivered, then the current action profile must be CD. This means that Pareto efficient equilibria taking a bang-bang form are such that once action CC is taken, then it is taken forever. This yields point (iii).

Finally we note that point (iv) is a direct implication of points (i) and (ii). This concludes the proof. $\square$

We now prove Lemma 5. Let us consider parameters $(b, c, \delta) \in \operatorname{int} \mathcal{G}_{DC/CC}$. Denote by G the corresponding PD. For any $d > 0$ let us denote by G_d the game

$$\begin{array}{c|cc} & C & D \\ \hline C & 1,1 & -c,b \\ D & b,-c & d,d \end{array}\ .$$

By substracting d to all payoffs and dividing them by $1-d$, we obtain that G_d is strategically equivalent to the repeated game G'_d with flow payoffs

$$\begin{array}{c|cc} & C & D \\ \hline C & 1,1 & \frac{-c-d}{1-d}, \frac{b-d}{1-d} \\ D & \frac{b-d}{1-d}, \frac{-c-d}{1-d} & 0,0 \end{array}.$$

Since $\tilde{\Gamma}_G \in \operatorname{int} \mathcal{G}_{DC/CC}$, for $d > 0$ small enough, we have that $\tilde{\Gamma}_{G'_d} \in \mathcal{G}_{DC/CC}$. This means that the set of SPE values of $\tilde{\Gamma}_{G'_d}$ is a quadrangle $\mathcal{V}'_d = \text{co}\{(0,0),(1,1),(0,\rho),(\rho,0)\}$ where $\rho \geq 1$. Note that since DC is enforceable, we must have $\frac{\delta}{1-\delta}\rho \geq \frac{c+d}{1-d}$. By Lemma 11, we know that the Pareto frontier of equilibrium payoffs is self enforced by a class of equilibrium strategies such that 1) action profile DD is never used on or off the equilibrium path, 2) once action CC is taken, it is taken forever. Let us denote by $\mathcal{E}$ this class of strategy profiles.

Since game G'_d is strategically equivalent to game G_d we obtain that strategy profiles in $\mathcal{E}$ are also equilibria of $\tilde{\Gamma}_{G_d}$ and generate its equilibrium Pareto frontier. The equilibrium Pareto frontier of $\tilde{\Gamma}_{G_d}$ is obtained by multiplying equilibrium values of $\tilde{\Gamma}_{G'_d}$ by $1-d$ and adding d. Its extreme points are $\{(d,\tilde{\rho});(1,1);(\tilde{\rho},d)\}$ where $\tilde{\rho} = (1-d)\rho + d \geq c + 2d$. Note also that in $\tilde{\Gamma}_{G_d}$, the values generated by these equilibria are such that at every history h and for all $i \in \{1,2\}$, $v_i(h) \geq d$.

Let us now show that strategy profiles in $\mathcal{E}$ are also equilibria of G. This occurs because G differs from G_d only in that the payoff from DD is $(0,0)$ rather than (d,d). Since strategy profiles in $\mathcal{E}$ never use outcome DD and $d > 0$, this implies that, whenever the one-shot incentive compatibility holds in G_d, it also holds in G. Hence strategy profiles in $\mathcal{E}$ are equilibria of G. Since payoffs upon CD, DC and CC are the same in G and G_d, equilibria $\mathcal{E}$ of G support the self-generating set of values with extreme points $\{(d,\tilde{\rho});(1,1);(\tilde{\rho},d)\}$.

We now reach the final step of the proof. First note that the equilibrium in which players both play D at every history generates payoffs $(0,0)$ and that DD is iteratively d-strict dominant in game $G((0,0))$ for any $d < \min\{b-1;c\}$. Pick $d \in (0, \min\{c, b-1, 1/(1-\delta)-b\})$ and consider an equilibrium s in $\mathcal{E}$. The set of values generated by $\mathcal{E}$ is self enforcing. Let us show that there exists $d > 0$ and an equilibrium $\hat{s}$ generating the same value as s, and s.t. at every history h, action $\hat{s}(h)$ is an iteratively d-strict dominant equilibrium of $G(v(\hat{s}|(h,\cdot)))$. Define H_s the set of histories on the equilibrium path of s. At any history $h \in H_s$ we modify s off of the equilibrium path as follows.

- At any history h such that $s(h) = CD$ use off path continuation values $v'_{DC} = v'_{DD} = v'_{CC} = (0,0)$, where $(0,0)$ is generated by the equilibrium in which players defect always. Since $s \in \mathcal{E}$, we must have that the value from playing CD at h is greater than d. This yields that at h, CD is iteratively d-strict dominant in $G(v(\hat{s}|(h,\cdot)))$. If $s(h) = DC$, a symmetric change makes DC iteratively d-strict dominant in a game $G(v(\hat{s}|(h,\cdot)))$ where off path continuation values have been set to $(0,0)$ while on path continuation values have not changed.
- At a history h such that $s(h) = CC$, we must have that $v_{CC} = (1,1)$ since equilibria in $\mathcal{E}$ are such that once action profile CC occurs, it occurs forever. We change off path values by setting $v_{DD} = 0$, $v_{DC} = (d, \tilde{\rho})$ and $v_{CD} = (\tilde{\rho}, d)$. Since $\frac{\delta}{1-\delta}\tilde{\rho} > c + 2d$, and $1/(1-\delta) > b + d$, we obtain that CC is iteratively strictly dominant in $G(v(\hat{s}|(h,\cdot)))$.

It results from this that $\text{co}\{(0,0);(\tilde{\rho},d);(d,\tilde{\rho});(1,1)\}$ is self enforced by equilibria s such that at every history h, action profile $s(h)$ is iteratively d-strict dominant in $G(v(s|(h,\cdot)))$. Because h is continuous in d, by taking d to 0, we obtain that indeed, whenever $v \in \text{int}\,\mathcal{V}^{\text{SPE}} \cup \{(0,0);(1,1)\}$, there exist $d > 0$ and a subgame-perfect equilibrium s implementing these values such that at every history h, the equilibrium action profile $s(h)$ is iteratively d-strictly dominant in the augmented game $G(v(s|(h,\cdot)))$.

This concludes the proof when $G \in \text{int}\,\mathcal{G}_{DC/CC}$. A similar proof holds when $G \in \text{int}\,\mathcal{G}_{DC}$.

A.3.2 Proof of Proposition 9

For each $v \in \text{co}\, g(A)$ with $v \gg 0$, we will construct $\tilde{s}^*$ with payoff profile v such that $\tilde{s}^*$ is a profile of finite-state automaton strategies, $\tilde{s}^*$ is simple, and $\tilde{s}^*(h, \theta^t)$ survives iterative elimination of strictly dominated actions in $G(\tilde{\mathcal{V}}(s^*|(h, \theta^t, \cdot)))$ for every $t \geq 1$, $h \in H_{t-1}$,

$\theta^t \in [0,1)^t$, and large δ. For each such δ, since there are finitely many states, we can pick $d > 0$ uniformly such that $\tilde{s}^*(h, \theta^t)$ is d-robust in $G(\tilde{v}(s^*|(h, \theta^t, \cdot)))$ for every $t \geq 1$, $h \in H_{t-1}$, and $\theta^t \in [0,1)$, and thus, by Lemma 10, $\tilde{s}^*$ is strongly robust.

By the NEU condition, there exist $x^1, \ldots, x^n \in \mathbb{R}^n$ such that, for every $i \in N$, we have $x^i \gg 0$, $x^i_i < v_i$, and $x^i_i < x^j_i$ for every $j \neq i$ (Abreu, Dutta, and Smith (1994, Theorem 1)).

Let $a \in A$ be an action profile with $g(a) = v$, and, for each $i \in N$, let $a^i \in A$ be an action profile with $g(a^i) = x^i$. (Use public randomization if there are no such pure action profiles.) Let $\tilde{s}^*$ be the following Markov strategy profile starting with "state v."

- In state v, play a. If the observed action profile differs from a,then go to "state $\underline{v}^i$," where i is the smallest name index among all the deviators.

- In state $\underline{v}^i$, play m^i. If the observed action profile differs from m^i,then go to "state $\underline{v}^j$," where j is the smallest name index among all the deviators. Else, with probability q stay in "state $\underline{v}^i$", while with probability $1 - q$ proceed to "state x^i."

- In state x^i, play a^i. If the observed action profile differs from m^i,then go to "state $\underline{v}^j$," where j is the smallest name index among all the deviators. Else, stay in "state x^i."

We pick $q \in (0,1)$ close enough to 1 so that

$$|g| < \frac{2-q}{1-q} \min_i x^i_i.$$

Let $L_j(\underline{v}^i)$ be player j's continuation payoff in state $\underline{v}^i$. We have

$$L_i(\underline{v}^i) = \delta(qL_i(\underline{v}^i) + (1-q)x^i_i),$$

thus

$$L_i(\underline{v}^i) = \frac{\delta(1-q)}{1-\delta q} x^i_i > 0,$$

which converges to x_i^i as $\delta \to 1$. Similarly,

$$L_j(\underline{v}^i) \leq \frac{\delta(1-q)}{1-\delta q} x_j^i - \frac{1-\delta}{1-\delta q}|g|,$$

which converges to x_j^i as $\delta \to 1$.

In what follows, we verify that, at each histoy, the action profile prescribed by $\tilde{s}^*$ survives iterative elimination of strictly dominated actions in the augmented stage game for large δ. Consider state $\underline{v}^i$.

First, consider state $\underline{v}_i$. We will show that, for each j, m_j^i is a dominant action for player j in the augmented stage game for large δ after we eliminate all actions of player k other than m_k^i for all $k < j$. Then, by induction on j, m^i survives iterative elimination of strictly dominated actions in the augmented stage game for large δ.

- Suppose that $i = j$ and the opponents' action profile is m_{-j}^i. Then player j's payoff from m_j^i is $L_j(\underline{v}^i)$, which is larger than his maximal payoff from one-shot deviation $\delta L_j(\underline{v}^i)$.

- Suppose that $i = j$ and the opponents' action profile is not m_{-j}^i. Let $k > j$ be the smallest name index among the deviators. Then player j's payoff from m_j^i is at least $-(1-\delta)|g| + \delta L_j(\underline{v}^k)$, which is larger than his maximal payoff from one-shot deviation $(1-\delta)|g| + \delta L_i(\underline{v}^i)$ for large δ since $x_j^k > x_j^j$.

- Suppose that $i \neq j$ and the opponents' action profile is m_{-j}^i. Then player j's payoff from m_j^i is $L_j(\underline{v}^i)$, which is larger than his maximal payoff from one-shot deviation $(1-\delta)|g| + \delta L_j(\underline{v}^j)$ for large δ since $x_j^i > x_j^j$.

- Suppose that $i \neq j$ and the opponents' action profile is not m_{-j}^i. Let $k > j$ be the smallest name index among the deviators. Then player j's payoff from m_j^i is at least $-(1-\delta)|g| + \delta L_j(\underline{v}^k)$, which is larger than his maximal payoff from one-shot deviation $(1-\delta)|g| + \delta L_j(\underline{v}^j)$ for large δ since $x_j^k > x_j^j$.

Next, consider state x^i. Similarly, we will show that, for each j, a^i_j is a dominant action for player j in the augmented stage game after we eliminate all actions of player k other than a^i_k for all $k < j$.

- Suppose that the opponents' action profile is a^i_{-j}. Then player j's payoffs from a^i_j is x^i_j, whereas his maximal payoff from one-shot deviation is $(1-\delta)|g| + \delta L_j(\underline{v}^j)$. Due to the choice of q and $x^i_j > x^j_j$, we have

  $$x^i_j - [(1-\delta)|g| + \delta L_j(\underline{v}^j)] > (1-\delta)\left(|g| - \frac{1+\delta-\delta q}{1-\delta q} x^j_j\right) > 0$$

  for large δ.

- Suppose that the opponents' action profile is not a^i_{-j}. Then the argument is the same as the one for state $\underline{v}^i$.

Finally, consider state v. Once again, we will show that, for each j, a_j is a dominant action for player j in the augmented stage game after we eliminate all actions of player k other than a_k for all $k < j$.

- Suppose that the opponents' action profile is a_{-j}. Then player j's payoffs from a_j is v_j, which is larger than his maximal payoff from one-shot deviation $(1-\delta)|g| + \delta L_j(\underline{v}^j)$ for large δ since $v_j > x^j_j$.

- Suppose that the opponents' action profile is not a_{-j}. Then the argument is the same as the one for state $\underline{v}^i$.

References

[1] Abreu, Dilip, 1988.

[2] Abreu, Dilip, David Pearce and Ennio Stacchetti 1990.

[3] Bhaskar V., George Mailath and Stephen Morris

[4] Chassang, Sylvain, *Fear of Miscoordination and the Robustness of Cooperation in Dynamic Global Games with Exit*, Princeton University mimeo, 2007.

[5] Chassang, Sylvain and Gerard Padro i Miquel *Conflict and Deterrence under Strategic Risk*, Princeton University mimeo, 2008.

[6] Dekel, Eddie and Drew Fudenberg

[7] Ely, Jeff and Juuso Välimaki

[8] Kajii, Atsushi and Stephen Morris, *The Robustness of Equilibria to Incomplete Information*, Econometrica, 1997.

[9] Fudenberg, Drew, David Kreps and David Levine

[10] Fudenberg, Drew and Eric Maskin

[11] Giannitsarou, Chrissy and Flavio Toxvaerd