

Robustness to Incomplete Information in Repeated Games*

Sylvain Chassang Satoru Takahashi
Princeton University Princeton University

March 19, 2009

Abstract

This paper extends to repeated games Kajii and Morris (1997)'s notion of robustness to incomplete information. Under a mild strengthening of the requirements of robustness, we show that dynamically robust equilibria can be characterized by applying a one-shot robustness principle that extends the one-shot deviation principle. As a corollary, we prove a factorization result analogous to that of Abreu, Pearce and Stacchetti (1990).

Using these results, we characterize explicitly the set of robust equilibria in the repeated Prisoners' Dilemma, and show that robustness requirements imply sharp restrictions on when cooperation can be sustained. We also show that a folk theorem in robust equilibria holds, but requires stronger identifiability conditions than the pairwise full rank condition of Fudenberg, Levine and Maskin (1994).

*We are grateful to Dilip Abreu, Eduardo Faingold, Drew Fudenberg, Johannes Hörner, Stephen Morris, Marcin Peški, Wolfgang Pesendorfer and Yuliy Sannikov for helpful discussions. We thank seminar participants at John Hopkins, Princeton, Penn State, Rochester, Rutgers, Toronto, UT Austin and at the 2008 Stony Brook Game Theory Festival, for useful comments. Chassang: chassang@princeton.edu, Takahashi: satorut@princeton.edu.

1 Introduction

The pioneering work of Rubinstein (1989) and Carlsson and van Damme (1993) has shown that strict equilibria of two-by-two coordination games can be destabilized by arbitrarily small perturbations in the players' information structure. Since then, robustness to incomplete information has received much attention. A variety of applied work uses robustness to incomplete information as a criterion for equilibrium selection.¹ A complementary literature explores robustness to incomplete information to ensure that specific equilibria of interest are robust to reasonable perturbations in the information structure.² In the context of static games, Kajii and Morris (1997, henceforth KM) provide a benchmark for both types of studies by analyzing the robustness of equilibria to all small perturbations in the information structure.³ KM offer sufficient conditions under which equilibria are robust to any small amount of incomplete information and, as a corollary, establish bounds on how much selection can be achieved using perturbations in the information structure. In this paper, we extend both the approach and the results of KM to repeated games with imperfect public monitoring.⁴ Following a dynamic programming approach à la Abreu, Pearce and Stacchetti (1990, henceforth APS), we also relate robustness in dynamic games to robustness in appropriate families of static games augmented with continuation values.

Given a complete-information game G , KM consider incomplete-information games U that are elaborations of G in the sense that, with high probability, payoffs of U are exactly those of G , and the agents know it. An equilibrium of G is robust if it is close to some Bayesian Nash equilibrium of U , for all elaborations U sufficiently close to G . Our approach

¹See for instance Morris and Shin (1998), Chamley (1999), Frankel, Morris and Pauzner (2003), Goldstein and Pauzner (2004) or Argenziano (2008). See Morris and Shin (2003) for an extensive literature review.

²See for instance Bergemann and Morris (2005), Oury and Tercieux (2007), or Aghion, Fudenberg and Holden (2008).

³KM, as well as this paper, maintain the common prior assumption and considers perturbations that are small from an ex ante perspective. Weinstein and Yildiz (2007) consider perturbations that are close from an interim perspective. Oyama and Tercieux (2007) and Izmalkov and Yildiz (2008) consider incomplete information perturbations that do not satisfy the common prior assumption.

⁴Appendix C provides an extension to general multistage games with discounting.

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 sequences $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of independent incomplete-information stage games, all of which are elaborations of G . Because perturbations that occur in future periods with some probability change current discounted expected payoffs with high probability, we consider a more stringent notion of robustness which only requires that elaborations U have payoffs close, rather than identical, 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 one-shot action profiles in appropriate families of static games. In particular, we prove a one-shot robustness principle which shows that an equilibrium of Γ_G is robust if and only if at any history the prescribed action profile is a robust equilibrium in the static game G augmented with continuation values. This allows us to characterize dynamically robust equilibria by considering only one-shot elaborations rather than all sequences of elaborations. Furthermore, this one-shot robustness principle implies a factorization result à la APS. Specifically, equilibrium values sustained by robust equilibria of Γ_G essentially correspond to the largest fixed point of a robust value mapping which associates future continuation values with current values generated by robust equilibria of the corresponding augmented stage game. Our two main applications highlight the practical value of these characterizations.

First, we compute explicitly the set of robust equilibrium values in the repeated Prisoners' Dilemma. We show that whenever outcome $(Defect, Cooperate)$ can be enforced under complete information, the set of robust equilibrium values is essentially equal to the set of equilibrium values under complete information. Inversely, whenever $(Defect, Cooperate)$ is not enforceable under complete information, the set of robust equilibria shrinks to permanent defection. In addition, we highlight that asymmetric equilibria which punish only deviators upon unilateral deviation are more effective than grim-trigger strategies at sustaining robust cooperation.

Second, we show that a folk theorem in robust equilibria holds for repeated games with imperfect public monitoring, but that it requires stronger identifiability conditions than the pairwise full rank conditions of Fudenberg, Levine and Maskin (1994). The reason for this difference is that in order to enforce behavior robustly, one needs to control continuation payoffs upon both unilateral and joint deviations from equilibrium behavior.

Much of the refinement literature is concerned with robustness in dynamic games in some form or another, and we cannot hope to successfully summarize all of the relevant work. 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) extend this question to iterative elimination of weakly dominated strategies. The approach to robustness developed in these papers is very different and less stringent than the one we develop here. In particular, their approach requires only the existence of an approximating sequence of elaborations for which the target equilibrium is strict.

More recently, Bhaskar, Mailath and Morris (2008) study the dynamic robustness of a specific equilibrium in the repeated Prisoners' Dilemma. They focus on the mixed-strategy equilibrium constructed by Ely and Välimäki (2002), and show that, under generic distributions of payoff perturbations, the Ely-Välimäki equilibrium cannot be approximated by equilibria with one-period memory but can be approximated by equilibria with infinite memory. 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.

Also closely related to this paper are Giannitsarou and Toxvaerd (2007) and Chassang (2007), who extend global-game approaches to dynamic settings. Giannitsarou and Toxvaerd (2007) show that in finite-horizon games with strategic complementarities, a global games perturbation à la Carlsson and van Damme (1993) selects a unique equilibrium. Chassang

(2007) considers infinite-horizon cooperation games with exit and shows that, even though the global-game perturbation does not yield uniqueness in such settings, it still selects a subset of equilibria whose qualitative properties are driven by risk-dominance considerations. An important difference between these papers and ours is that they consider robustness to a specific information perturbation whereas we study robustness to any sequence of independent elaborations. This makes our robustness results stronger and our non-robustness results weaker. From a technical perspective, considering robustness to all small perturbations greatly simplifies the analysis and in particular allows us to do away with strategic complementarity assumptions.

Finally this paper can be related to the recent work of Mailath and Morris (2002, 2006), Hörner and Olszewski (2008) and Mailath and Olszewski (2008) on almost-public monitoring. This literature explores the robustness of equilibria in repeated games with public monitoring to small perturbations in the monitoring structure. This departure from public monitoring induces incomplete information perturbations at every history, which depend both on the strategies players use and past histories. To some extent, we consider related perturbations which depend neither on strategies nor past histories.

The paper is structured as follows. Section 2 describes the framework and introduces an extension of KM's notion of robustness in static games. Section 3 defines robustness to incomplete information for repeated games and proves an analogue of the one-shot deviation principle, as well as a factorization result. Section 4 applies the results of Section 3 to characterize robust equilibrium values in the repeated Prisoners' Dilemma. Section 5 proves a folk theorem in robust equilibria for repeated games with imperfect public monitoring. Section 6 discusses alternative approaches to robustness in dynamic games. Section 7 concludes. Proofs are contained in Appendix A unless otherwise noted. Appendix B deals with technical measurability issues that occur when continuous public randomization devices are available. Appendix C extends our approach to multistage games with discounting.

2 Robustness in Static Games

In this section, we define a static notion of robustness to incomplete information that is more stringent than that 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 be 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. To accommodate this feature, our notion of robustness allows for elaborations that have payoffs close (instead of identical) to the payoffs of the complete-information game with a large probability.⁵

2.1 Definitions

Consider a complete-information game $G = (N, (A_i, g_i)_{i \in N})$ with a finite set $N = \{1, \dots, n\}$ of players. Each player $i \in N$ is associated with a finite set A_i of actions and a payoff function $g_i: A \rightarrow \mathbb{R}$, 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 for payoff functions: $|g_i| = \max_{a \in A} |g_i(a)|$ and $|g| = \max_{i \in N} |g_i|$. For $d \geq 0$, an action profile $a^* = (a_i^*)_{i \in N} \in A$ is a d -strict equilibrium if $g_i(a^*) \geq g_i(a_i, a_{-i}^*) + d$ for every $i \in N$ and $a_i \in A_i \setminus \{a_i^*\}$. An action profile a^* is a *Nash equilibrium* if it is a 0-strict equilibrium; a^* is a *strict equilibrium* if it is a d -strict equilibrium for some $d > 0$.

An elaboration U of game G is an incomplete-information game $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$, where Ω is a countable set of states, P is a probability distribution over Ω , and, for each player $i \in N$, $u_i: A \times \Omega \rightarrow \mathbb{R}$ is her bounded state-dependent payoff function and Q_i is her in-

⁵In Appendix C, we show that unless we impose such a strengthening of robustness, the one-shot deviation principle does not have a robust analogue: the dynamic robustness of an equilibrium would not be implied by the robustness of each one-shot action profile in appropriate stage games augmented with continuation values.

formation partition over Ω . Let $|u| = \sup_{\omega \in \Omega} |u(\cdot, \omega)|$. For any finite set X , let $\Delta(X)$ denote the set of probability distributions over X . A mixed strategy of player i is a Q_i -measurable mapping $\alpha_i: \Omega \rightarrow \Delta(A_i)$.⁶ The domain of u_i extends to mixed or correlated strategies in the usual way. Prior P and a profile $\alpha = (\alpha_i)_{i \in N}$ of mixed strategies induce a distribution $P^\alpha \in \Delta(A)$ over action profiles defined by

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

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 every $i \in N$ and every Q_i -measurable strategy α_i of player i .

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

$$P(\{\omega \in \Omega \mid \forall i \in N, \forall \omega' \in Q_i(\omega), |u_i(\cdot, \omega') - g_i| \leq d\}) \geq 1 - \varepsilon,$$

where $Q_i(\omega)$ denotes the element of partition Q_i that contains ω .

Definition 1 (static robustness). For $d \geq 0$, a pure 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 Bayesian-Nash equilibrium α^* such that $P^{\alpha^*}(a^*) \geq 1 - \eta$.⁷

A pure Nash equilibrium a^* of G is *strongly robust* if it is d -robust for some $d > 0$.

⁶With a slight abuse of terminology, we say that α_i is Q_i -measurable if it is measurable with respect to the σ -algebra generated by Q_i .

⁷To avoid unnecessary notation, we do not extend our definition of d -robustness to mixed equilibria of G . If we did, a straightforward extension of Lemma 1 would show that in fact no mixed strategy equilibria are d -robust.

Note that 0-robustness corresponds to robustness in the sense of KM. To some extent, for $d > 0$, one can think of d -robustness as requiring that robustness in the sense of KM hold uniformly over a neighborhood of complete information games. As can be expected, for $d > 0$, 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 d -robust, then it is $2d$ -strict.*

Lemma 1 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

All action profiles are 0-robust, but none of them are strongly robust.

Still, provided that a^* is a strict 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

KM provide the following sufficient conditions for robustness: a^* is robust in G if it is the unique correlated equilibrium of G , or if it is a \mathbf{p} -dominant equilibrium with $\sum_i p_i < 1$.⁸ Here we show that, whenever a^* is a strict equilibrium, the same conditions imply that a^* is strongly robust. 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 a strict equilibrium and the unique correlated equilibrium of G , then it is strongly robust in G .*

The proof, given in Appendix A.2, is an extension of the proof of KM (Proposition 3.2). We also exploit the fact that, when a^* is strict and the unique correlated equilibrium of G ,

⁸See KM, or Definition 3 in this paper for a definition of \mathbf{p} -dominance.

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 actions.

Definition 2 (iterative d -dominance). For $d \geq 0$, we say that an action profile a^* is *iteratively d -dominant in G* if there exists a sequence $\{X_{i,t}\}_{t=0}^T$ of action sets with $A_i = X_{i,0} \supseteq X_{i,1} \supseteq \dots \supseteq X_{i,T} = \{a_i^*\}$ for each $i \in N$ such that, at every stage t of elimination with $1 \leq t \leq T$, for each $i \in N$ and $a_i \in X_{i,t-1} \setminus X_{i,t}$, there exists $a'_i \in X_{i,t-1}$ such that $g_i(a'_i, a_{-i}) > g_i(a_i, a_{-i}) + d$ for all $a_{-i} \in \prod_{j \in N \setminus \{i\}} X_{j,t-1}$.

The next proposition shows that, for any given d , iterative d -dominance implies $d/2$ -robustness. This improves on Proposition 1, which only asserts d -robustness for *some* $d > 0$. The proof is given in Appendix A.3.

Proposition 2 (strong robustness of iteratively d -dominant equilibria). *If a^* is iteratively d -dominant, then it is $d/2$ -robust.*

KM's second sufficient condition is particularly useful in applied settings.

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

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

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 $\lambda(a_{-i}^*) > p_i$,

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

Note that an action profile is a strictly \mathbf{p} -dominant equilibrium if and only if it is a strict and \mathbf{p} -dominant equilibrium. KM establish that, whenever a^* is a \mathbf{p} -dominant equilibrium

of G with $\sum_i p_i < 1$, it is robust in G . This extends to the case of strong robustness as follows.

Proposition 3 (strong robustness of strictly \mathbf{p} -dominant equilibria). *If a^* is a strictly \mathbf{p} -dominant equilibrium of G with $\sum_i p_i < 1$, then it is strongly robust.*

In Appendix A.4, we extend this proposition and obtain a sufficient condition for d -robustness given a fixed d .

We know from KM (Lemma 5.5) that, if a game has a strictly \mathbf{p} -dominant equilibrium with $\sum_i p_i < 1$, then no other action profile is 0-robust. Combined with Proposition 3, this implies that, if a game has a strictly \mathbf{p} -dominant equilibrium with $\sum_i p_i < 1$, it is the unique strongly robust equilibrium. For example, in a two-by-two coordination game, a strict risk-dominant equilibrium is the unique strongly robust equilibrium.

3 Robustness in Repeated Games

In this section, we formulate a notion of robustness to incomplete information that is appropriate for repeated games. We consider payoff shocks that are additively separable and stochastically independent across periods. We show in Sections 3.2 and 3.3 that under this definition, dynamically robust equilibria admit a convenient recursive representation. Section 6 discusses other definitions of dynamic robustness which correspond to different classes of perturbations against which to check robustness.

3.1 Definitions

Consider a complete-information game $G = (N, (A_i, g_i)_{i \in N})$ as well as a public monitoring structure (Y, π) , where Y is a finite set of public outcomes and $\pi : A \rightarrow \Delta(Y)$ maps action profiles to distributions over public outcomes. Keeping fixed the discount factor $\delta \in [0, 1)$, let Γ_G denote the infinitely repeated game with stage game G , discrete time $t \in \{1, 2, 3, \dots\}$,

and monitoring structure (Y, π) .⁹ Given $l \geq 0$, Let $H_l = Y^l$ be the set of public histories of length l and $H = \bigcup_{l \geq 0} H_l$ the set of all finite public histories. H_{t-1} corresponds to possible histories at the beginning of period t . A pure public strategy of player i is a mapping $s_i: H \rightarrow A_i$. Conditional on history $h_{t-1} \in H$, a public strategy profile $s = (s_i)_{i \in N}$ induces a distribution over sequences (a_t, a_{t+1}, \dots) of future action profiles, which in turn induces continuation payoffs $v_i(s|h_{t-1})$ such that

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

A public-strategy profile s^* is a *perfect public equilibrium* (PPE) if $v_i(s^*|h_{t-1}) \geq v_i(s_i, s_{-i}^*|h_{t-1})$ for every $h_{t-1} \in H$, $i \in N$ and public strategy s_i of player i (Fudenberg, Levine and Maskin, 1994). We denote by V^{PPE} the set of PPE payoff profiles of Γ_G . The restriction to public strategies corresponds to the assumption that, although player i observes her own action a_i as well as stage game payoffs $g_i(a)$ (or perhaps some noisy signal of $g_i(a)$), she conditions her behavior only on public outcomes.

We define perturbations of Γ_G as follows. Consider a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete-information elaboration $U_t = (N, \Omega_t, P_t, (A_i, u_{it}, Q_{it})_{i \in N})$ of G . We define the norm $|\mathbf{U}| = \sup_{t \in \mathbb{N}} |u_t|$. Given a sequence \mathbf{U} such that $|\mathbf{U}| < \infty$, we denote by $\Gamma_{\mathbf{U}}$ the following infinite-horizon game with public monitoring. In each period t , state $\omega_t \in \Omega_t$ is generated according to P_t independently of past action profiles, past outcomes and past states. Each player i receives a signal according to her information partition Q_{it} and chooses action $a_{it} \in A_i$. At the end of the period, an outcome $y \in Y$ is drawn according to $\pi(a_t)$ and publicly observed. A public strategy of player i is a mapping $\sigma_i: \bigcup_{t \geq 1} H_{t-1} \times \Omega_t \rightarrow \Delta(A_i)$ such that $\sigma_i(h_{t-1}, \cdot)$ is Q_{it} -measurable for every public history $h_{t-1} \in H$.

Conditional on public history h_{t-1} , a public-strategy profile $\sigma = (\sigma_i)_{i \in N}$ induces a prob-

⁹We omit to index the game by its monitoring structure for conciseness. Note that this class of games includes games with perfect monitoring and games with finite public randomization devices. Appendix B deals with the measurability restrictions necessary to extend our results to continuous public randomizations.

ability distribution over sequences of future action profiles and states, which allows us to define continuation payoffs $v_i(\sigma|h_{t-1})$ such that

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

The assumption of uniformly bounded stage-game payoffs implies that the above infinite sum is well defined. A public-strategy profile σ^* is a *perfect public equilibrium* (PPE) 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 public strategy σ_i of player i .

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

A PPE s^* of Γ_G is *strongly robust* if it is d -robust for some $d > 0$.

Let V^{rob} be the set of all payoff profiles of strongly robust PPEs in Γ_G . Note that our definition of dynamic robustness considers only sequences $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete-information games that are uniformly close to G . If $(U_t)_{t \in \mathbb{N}}$ approached G only pointwise, then the robustness criterion would become too restrictive. For example, if we allowed for such perturbations, whenever the stage game G had a unique Nash equilibrium a^* , the only robust equilibrium of Γ_G would be the repetition of a^* . Indeed, for any large but finite T , consider $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ 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 $\omega_t \in \Omega_t$. Sequence (U_t) converges to G in a pointwise sense. Clearly, the game $\Gamma_{\mathbf{U}}$ has a finite effective horizon, and by backward induction players play a^* for the first T periods in every PPE of $\Gamma_{\mathbf{U}}$. This is why we consider uniformly small perturbations.

We now relate the dynamic robustness of PPEs of Γ_G to the robustness of one-shot action profiles in appropriate static games augmented with continuation values. This yields a one-shot robustness principle analogous to the one-shot deviation principle.

3.2 A One-Shot Robustness Principle

Given a stage game G and a one-period-ahead continuation-payoff profile $w: Y \rightarrow \mathbb{R}^n$ contingent on the realized public outcome, let $G(w)$ be the complete-information game augmented with continuation values w , i.e., $G(w) = (N, (A_i, g'_i)_{i \in N})$ such that $g'_i(a) = (1 - \delta)g_i(a) + \delta \mathbb{E}[w_i(y)|a]$ for every $i \in N$ and $a \in A$. For a strategy profile s of repeated game Γ_G and a history h , let $w_{s,h}$ be the contingent-payoff profile given by $w_{s,h}(y) = (v_i(s|(h, y)))_{i \in N}$ for each $y \in Y$. By the one-shot deviation principle, s^* is a PPE of repeated game Γ_G if and only if $s^*(h)$ is a Nash equilibrium of $G(w_{s^*,h})$ for every $h \in H$ (Fudenberg and Tirole, 1991, Theorem 4.2).

The next lemma extends Lemma 1 and shows that at any history, the one-shot action profile prescribed by a strongly robust PPE is a strict equilibrium of the appropriate augmented stage game.

Lemma 2 (strictness in augmented games). *If s^* is d -robust in Γ_G , then $s^*(h)$ is $2(1 - \delta)d$ -strict in $G(w_{s^*,h})$ for every $h \in H$.*

The following theorem relates strong robustness in Γ_G to strong robustness in all appropriately augmented stage games. This is the analogue for robust PPEs of the one-shot deviation principle. The proof is given in Appendix A.6.

Theorem 1 (one-shot robustness principle). *A strategy profile s^* is a strongly robust PPE of Γ_G if and only if there exists $d > 0$ such that, for every $h \in H$, $s^*(h)$ is a d -robust equilibrium of $G(w_{s^*,h})$.*

An immediate corollary of Theorem 1 is that a finite-automaton PPE s^* is strongly robust if and only if $s^*(h)$ is strongly robust in $G(w_{s^*,h})$ for every $h \in H$.

As Appendix C highlights, this one-shot robustness principle generalizes to a general class of multistage games with discounting. The proof of Theorem 1 exploits heavily the fact that strong robustness is a notion of robustness that holds uniformly over small neighborhoods of

games. In fact, the one-shot robustness principle would not generally hold if we considered robustness in the sense of KM rather than d -robustness. As a counter example, Appendix C describes a mixed-strategy equilibrium of an infinite-horizon dynamic game such that, for every history, the one-shot action profile prescribed by the equilibrium is robust in the sense of KM in the appropriately augmented stage game, but the overall equilibrium should not be considered dynamically robust as small perturbations far away in time can trickle down and cause significant behavior change earlier on.¹⁰

3.3 Factorization

In this section, we use Theorem 1 to obtain a recursive characterization of V^{rob} , the set of strongly robust PPE payoff profiles. More precisely, we prove self-generation and factorization results analogous to those of APS. We begin with a few definitions.

Definition 5 (robust enforcement). For $a \in A$, $v \in \mathbb{R}^n$, $w: Y \rightarrow \mathbb{R}^n$ and $d \geq 0$, w enforces (a, v) d -robustly if a is a d -robust equilibrium of $G(w)$ and $v = (1 - \delta)g(a) + \delta\mathbb{E}[w(y)|a]$.

For $v \in \mathbb{R}^n$, $V \subseteq \mathbb{R}^n$ and $d \geq 0$, v is d -robustly generated by V if there exist $a \in A$ and $w: Y \rightarrow V$ such that w enforces (a, v) d -robustly.

Let $B^d(V)$ be the set of payoff profiles that are d -robustly generated by V . This is the robust analogue of mapping $B(V)$ introduced by APS, where $B(V)$ is the set of all payoff profiles $v = (1 - \delta)g(a) + \delta\mathbb{E}[w(y)|a]$ with $w: Y \rightarrow V$ and a some Nash equilibrium of $G(w)$. We say that V is *self-generating with respect to B^d* if $V \subseteq B^d(V)$.

Lemma 3 (monotonicity).

(i) If $V \subseteq V' \subseteq \text{co } g(A)$, then $B^d(V) \subseteq B^d(V') \subseteq \text{co } g(A)$.

(ii) B^d admits a largest fixed point V^d among all subsets of $\text{co } g(A)$.

¹⁰Because multistage games with discounting allow for greater degrees of freedom than repeated games, it is considerably easier to come up with a counter-example from this larger class of games.

(iii) If $V \subseteq \text{co } g(A)$ and V is self-generating with respect to B^d , then $V \subseteq V^d$.

Proof. (i) follows from the definition of B^d . (ii) and (iii) follow from (i) and Tarski's fixed point theorem. \square

Note also that by definition, $B^d(V)$ is decreasing in d with respect to set inclusion. We characterize V^{rob} in terms of B^d as follows.

Theorem 2 (characterization of V^{rob}). *If V is a compact set such that $V^{\text{rob}} \subseteq V \subseteq \text{co } g(A)$ and $B^d(V) \subseteq V$ for every $d > 0$, then*

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

Proof. We first show that $V^{\text{rob}} = \bigcup_{d>0} V^d$. For each $v \in V^{\text{rob}}$, let s^* be a strongly robust PPE of Γ_G that yields value v . Then, by Theorem 1, there exists $d > 0$ such that $V^* = \{v(s^*|h) \in \mathbb{R}^n \mid h \in H\}$ is self-generating with respect to B^d . By Lemma 3, $v \in V^* \subseteq V^d$. Thus $V^{\text{rob}} \subseteq \bigcup_{d>0} V^d$. Let us turn to the other direction of set inclusion.

For each $d > 0$, since V^d is self-generating with respect to B^d , for each $v \in V^d$, there exist $a(v) \in A$ and $w(v, \cdot): Y \rightarrow V^d$ such that $w(v, \cdot)$ enforces $(a(v), v)$ d -robustly. Pick any $v \in V^d$. We construct s^* recursively as follows: $s^*(\emptyset) = a(v)$, $s^*(y_1) = a(w(v, y_1))$, $s^*(y_1, y_2) = a(w(w(v, y_1), y_2))$, and so on. By construction, $s^*(h)$ is d -robust in $G(w_{s^*, h})$ for every $h \in H$. Therefore, by Theorem 1, s^* is a strongly robust PPE of Γ_G that attains v , and thus $v \in V^{\text{rob}}$. Thus $V^d \subseteq V^{\text{rob}}$ for every $d > 0$.

Let us now show that $\bigcup_{d>0} V^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (B^d)^k(V)$, which corresponds to APS's algorithm result. To this end, we define $\bar{B}^d(V)$ by the closure of $B^d(V)$. Fix a compact set V such that $V^{\text{rob}} \subseteq V \subseteq \text{co } g(A)$ and $B^d(V) \subseteq V$ for every $d > 0$. Denote $f^\infty(V) = \bigcap_{k=0}^{\infty} f^k(V)$ for $f = B^d$ or \bar{B}^d . Since $V^d \subseteq V^{\text{rob}} \subseteq V$, by the monotonicity of B^d and \bar{B}^d , we have $V^d \subseteq (B^d)^\infty(V) \subseteq (\bar{B}^d)^\infty(V)$ for every $d > 0$.

To prove the opposite direction of set inclusion, we show that, for each $d > 0$, $(\bar{B}^d)^\infty(V)$

is self-generating with respect to $B^{d/2}$, which implies that $(\bar{B}^d)^\infty(V) \subseteq V^{d/2}$ by Lemma 3. Pick any $v \in (\bar{B}^d)^\infty(V)$. For each $k \geq 1$, since we have $v \in (\bar{B}^d)^\infty(V) \subseteq (\bar{B}^d)^k(V)$, there exist $a^k \in A$ and $w^k: Y \rightarrow (\bar{B}^d)^{k-1}(V)$ such that w^k enforces (a^k, v) d -robustly. Since A and Y are finite and $(\bar{B}^d)^k(V)$ is compact, by taking a subsequence, we can assume without loss of generality that $a^k = a^*$ and $w^k \rightarrow w^*$ as $k \rightarrow \infty$ for some $a^* \in A$ and $w^*: Y \rightarrow \mathbb{R}^n$. This implies that there exists $k^* \geq 1$ such that $|w^{k^*} - w^*| \leq d/(2\delta)$. Since w^{k^*} enforces (a^*, v) d -robustly, w^* enforces (a^*, v) $d/2$ -robustly. Moreover, for each $k \geq 1$ and $y \in Y$, since $w^l(y) \in (\bar{B}^d)^{l-1}(V) \subseteq (\bar{B}^d)^{k-1}(V)$ for every $l \geq k$ and $(\bar{B}^d)^{k-1}(V)$ is compact, by taking $l \rightarrow \infty$, we have $w^*(y) \in (\bar{B}^d)^{k-1}(V)$. Since this holds for every $k \geq 1$, $w^*(y) \in (\bar{B}^d)^\infty(V)$. Thus $v \in B^{d/2}((\bar{B}^d)^\infty(V))$, and $(\bar{B}^d)^\infty(V)$ is self-generating with respect to $B^{d/2}$. \square

Note that $V = \text{co } g(A)$ satisfies the assumption of Theorem 2. Theorem 2 corresponds to APS's self-generation, factorization and algorithm results (APS, Theorems 1, 2 and 5), which show that V^{PPE} is the largest bounded fixed point of the mapping B and can be computed by iteratively applying B to a sufficiently large compact set. Here, we require robust enforcement at every stage and mapping B is replaced by B^d .

We illustrate the practical value of Theorems 1 and 2 by means of two applications. First, we characterize explicitly the set of strongly robust PPEs in the repeated Prisoners' Dilemma and highlight how it differs from the set of PPEs under complete information. Second, we prove a folk theorem in robust equilibria and show that it requires stronger identifiability conditions than those of Fudenberg, Levine and Maskin (1994).

4 Robustness in the Repeated Prisoners' Dilemma

In this section, we characterize strongly robust PPE payoff profiles in the repeated Prisoners' Dilemma with perfect monitoring.¹¹ We show that, whenever outcome $(\textit{Defect}, \textit{Cooperate})$ can be enforced in equilibrium under complete information, the set of strongly robust PPE

¹¹Note that, under perfect monitoring, PPEs simply correspond to subgame-perfect equilibria.

payoff profiles is essentially equal to the set of complete information PPE payoff profiles. Inversely, whenever $(Defect, Cooperate)$ cannot be enforced under complete information, the set of strongly robust PPEs shrinks to permanent defection. We also show that robustness considerations refine our intuitions about which equilibria best sustain cooperation. More precisely, asymmetric equilibria in which only the deviator is punished upon unilateral deviation are more likely to be robust than grim-trigger strategies that punish everybody.

Throughout this section, let PD denote the two-player Prisoners' Dilemma 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 also allow players to condition their behavior on a continuous public randomization device.¹² We are interested in Γ_{PD} , the repeated PD with public randomization devices and perfect monitoring.

4.1 Robust Cooperation in Trigger Strategies

As an illustration, we begin by studying the robustness of trigger strategies. Consider trigger strategies such that players play C if D has never been played (cooperative state), and players play D if D has been played at some past history (punishment state). Under complete information, trigger strategies form a PPE if and only if $1/(1 - \delta) \geq b$. By Theorem 1, trigger strategies form a strongly robust equilibrium if and only if, CC is strongly robust in

	C	D
C	$1, 1$	$-(1 - \delta)c, (1 - \delta)b$
D	$(1 - \delta)b, -(1 - \delta)c$	$0, 0$

¹²Formally, the framework of Section 3 only covers finite public randomization devices. See Appendix B for a description of the measurability conditions necessary to extend our analysis to continuous public randomizations.

and DD is strongly robust in

	C	D
C	$1 - \delta, 1 - \delta$	$-(1 - \delta)c, (1 - \delta)b$
D	$(1 - \delta)b, -(1 - \delta)c$	$0, 0$

The latter property clearly holds since D is dominant in the augmented stage game at punishment histories. From Proposition 3 and KM (Lemma 5.5), it follows that strategies are strongly robust if and only if CC is strictly risk dominant in the augmented stage game at the cooperative state, i.e., if and only if $1/(1 - \delta) > b + c$. This condition highlights that for cooperation to be robustly sustainable in trigger strategies, the losses c that a player experiences when she cooperates and her opponent defects matter as much as the deviation temptation b .

This last result also suggests that trigger strategies may not be the most efficient way to sustain cooperation. Indeed trigger strategies punish both players upon unilateral deviation, which diminishes the deviation temptation of each player, but also increases a player's losses when she cooperates but her opponent defects. The next sections consider the possibility of asymmetric strategies and show that they can increase the scope for cooperation. We begin by providing a classification of PDs.

4.2 A Classification of PDs

We say that action profile a is *enforceable* in Γ_{PD} if some PPE of Γ_{PD} prescribes a at some history. We classify PDs as follows.

Definition 6 (classification of Prisoners' Dilemma games). Fix δ . We define four classes of PDs, $\mathcal{G}_{DC/CC}$, \mathcal{G}_{DC} , \mathcal{G}_{CC} and \mathcal{G}_\emptyset as follows:

- (i) $\mathcal{G}_{DC/CC}$ is the class of PDs such that DC and CC are enforceable in Γ_{PD} .
- (ii) \mathcal{G}_{DC} is the class of PDs such that DC is enforceable in Γ_{PD} , but CC is not.

(iii) \mathcal{G}_{CC} is the class of PDs such that CC is enforceable in Γ_{PD} , but DC is not.

(iv) \mathcal{G}_\emptyset is the class of PDs such that neither CC , DC or CD are enforceable.

Note that DD is always enforceable. Stahl (1991) characterizes explicitly the set V^{PPE} of PPE payoff profiles under complete information as a function of parameters δ , b and c . Given a fix parameter δ , Appendix A.7 provides simple conditions on b and c characterizing which of these classes a given PD belongs to. See Figure 1 for a representation of classes of PDs as a function of b and c , for δ fixed.

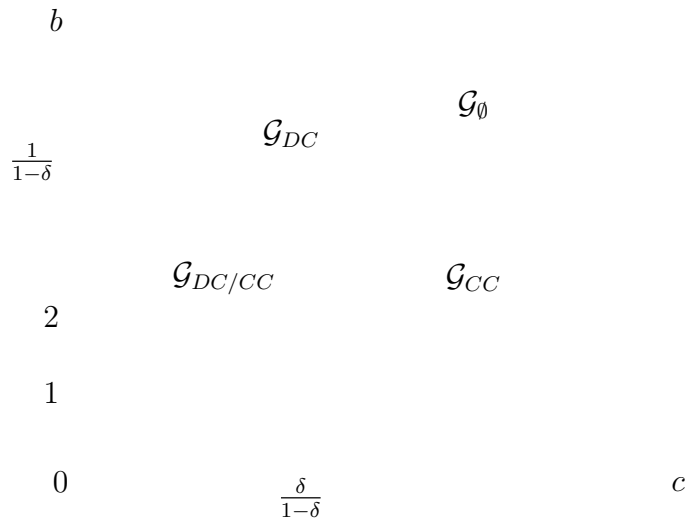


Figure 1: V^{PPE} as a function of b and c

Stahl (1991) also shows that for PDs in $\mathcal{G}_{DC/CC}$, we have $V^{\text{PPE}} = \text{co}\{(0, 0), (1, 1), (0, \phi), (\phi, 0)\}$, with $\phi \geq 1$. This means that for PDs in $\mathcal{G}_{DC/CC}$ it is possible to punish one player while giving the other one her maximum continuation value. For PDs in \mathcal{G}_{DC} , we have $V^{\text{PPE}} = \text{co}\{(0, 0), (0, b - c), (b - c, 0)\}$.¹³ Finally, we have that $V^{\text{PPE}} = \text{co}\{(0, 0), (1, 1)\}$ for PDs in \mathcal{G}_{CC} , and $V^{\text{PPE}} = \{(0, 0)\}$ for PDs in \mathcal{G}_\emptyset .

¹³Note that if $PD \in \mathcal{G}_{DC}$, then $b > c$.

4.3 A Robustness Result

We first show that when DC is enforceable in Γ_{PD} , then the set V^{rob} of strongly robust PPE payoff profiles is essentially equal to V^{PPE} . Indeed, if action profile DC is enforceable in Γ_{PD} , then, essentially every payoff profile $v \in V^{\text{PPE}}$ can be implemented by a PPE satisfying the following remarkable property, which we call *stage dominance*.¹⁴

Lemma 4 (stage dominance in augmented games). *Fix δ . If $PD \in \text{int } \mathcal{G}_{DC/CC}$, then, for any $v \in \{(0, 0), (1, 1)\} \cup \text{int } V^{\text{PPE}}$, there exist $d > 0$ and a PPE s^* of Γ_{PD} with payoff profile v such that, for every public history h , $s^*(h)$ is iteratively d -dominant in $PD(w_{s^*, h})$.*¹⁵

If $PD \in \text{int } \mathcal{G}_{DC}$, then, for any $v \in \{(0, 0)\} \cup \text{int } V^{\text{PPE}}$, there exist $d > 0$ and a PPE s^ of Γ_{PD} with payoff profile v such that, for every public history h , $s^*(h)$ is iteratively d -dominant in $PD(w_{s^*, h})$.*

The detailed proof of Lemma 4, given in Appendix A.8, is lengthy but the main idea of the argument is fairly straightforward. We show that, for every equilibrium, its off-path behavior can be modified so that each equilibrium action profile is iteratively dominant in the appropriate augmented stage game. The proof exploits the fact that values in V^{PPE} allow us to punish one player while giving the other her maximum continuation value.

As an example, consider a game $PD \in \text{int } \mathcal{G}_{DC/CC}$ and consider the profile of grim-trigger strategies. On the equilibrium path CC is a Nash equilibrium of

	C	D
C	1, 1	$-(1 - \delta)c, (1 - \delta)b$
D	$(1 - \delta)b, -(1 - \delta)c$	0, 0

Because DD is also an equilibrium of that game, CC is not iteratively dominant in this

¹⁴This stage dominance property is related to Miller (2007)'s notion of ex post equilibrium in repeated games of adverse selection, but allows for iterated elimination of strictly dominated actions.

¹⁵We identify a PD by its parameters $(b, c) \in \mathbb{R}^2$, so the interior of a class of PDs is derived from the standard topology on \mathbb{R}^2 .

augmented game. This can be resolved by changing continuation strategies upon outcomes CD and DC . By Stahl's characterization, we know that V^{PPE} takes the form $\text{co}\{(0, 0), (1, 1), (0, \phi), (\phi, 0)\}$, where $\phi \geq 1$.

Consider any public history of the form (CC, \dots, CC, CD) .¹⁶ The grim-trigger strategy prescribes permanent defection. We replace this continuation strategy by a PPE s_{CD} that attains $(\phi, 0)$ so that only the deviator is punished upon unilateral deviation. We also replace the continuation strategy after (CC, \dots, CC, DC) by a PPE s_{DC} that attains $(0, \phi)$. Then the augmented game after (CC, \dots, CC) becomes

	C	D
C	1, 1	$-(1 - \delta)c + \delta\phi, (1 - \delta)b$
D	$(1 - \delta)b, -(1 - \delta)c + \delta\phi$	0, 0

By assumption, CD and DC are enforceable in PPE, and we have that $-(1 - \delta)c + \delta\phi \geq 0$. Thus C is weakly dominant for both players in this augmented game. Because PD is in the interior of $\mathcal{G}_{DC/CC}$, it follows that C must in fact be strictly dominant. The difficult part of the proof consists of showing that this can be done iteratively and that strategy profiles s_{CD} and s_{DC} can be modified off of their equilibrium paths so that they become stage-dominant as well. The following robustness result follows directly from Lemma 4.

Proposition 4 (robust equilibria). *Fix δ . If $\text{PD} \in \text{int } \mathcal{G}_{DC/CC}$, then*

$$\{(0, 0), (1, 1)\} \cup \text{int } V^{\text{PPE}} \subseteq V^{\text{rob}} \subseteq V^{\text{PPE}}.$$

If $\text{PD} \in \text{int } \mathcal{G}_{DC}$, then

$$\{(0, 0)\} \cup \text{int } V^{\text{PPE}} \subseteq V^{\text{rob}} \subseteq V^{\text{PPE}}.$$

Proof. This follows from Proposition 2, Theorem 1 and Lemma 4. □

¹⁶We omit public randomizations to simplify notation.

4.4 A Fragility Result

Proposition 4 shows that if DC is enforceable under complete information, then essentially any PPE payoff profile can be sustained by a strongly robust PPE. Conversely, we now show that, if DC is not enforceable in PPE, then the only strongly robust PPE is permanent defection.

Proposition 5 (fragile equilibria). *Fix δ . If $PD \in \mathcal{G}_{CC}$, then the only strongly robust PPE of Γ_{PD} is permanent defection, and $V^{\text{rob}} = \{(0, 0)\}$.*

Proof. The proof is by contradiction. Assume that there exist a strongly robust PPE s^* of Γ_{PD} and a public history h such that $s^*(h) \neq DD$. Since $PD \in \mathcal{G}_{CC}$, s^* is necessarily strongly symmetric, i.e., it prescribes only action profiles CC or DD . This implies that $s^*(h) = CC$ and that, for every action profile a , the continuation-payoff profile after (h, a) is symmetric between the players. Furthermore, we have $c > \delta/(1 - \delta)$; otherwise, DC would be enforceable in PPE.

Because of the symmetry of continuation payoffs, the augmented game $PD(w)$ at history h takes the form

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

where w_{CC}, w_{CD}, w_{DC} and w_{DD} are in $[0, 1]$. Note that CC is a Nash equilibrium of $PD(w)$ since s^* is a PPE of Γ_{PD} . DD is also a Nash equilibrium of $PD(w)$ because $c > \delta/(1 - \delta)$, $w_{DD} - w_{CD} \geq -1$ and $w_{DD} - w_{DC} \geq -1$.

We now show that DD is strictly risk-dominant in $PD(w)$, i.e., that

$$\begin{aligned}
 & [\delta w_{DD} + (1 - \delta)c - \delta w_{CD}][\delta w_{DD} + (1 - \delta)c - \delta w_{DC}] \\
 & > [1 - \delta + \delta w_{CC} - (1 - \delta)b - \delta w_{CD}][1 - \delta + \delta w_{CC} - (1 - \delta)b - \delta w_{DC}]. \quad (1)
 \end{aligned}$$

Note that each square bracket term of (1) is nonnegative because CC and DD are Nash equilibria of $PD(w)$. Also note that

$$\delta w_{DD} + (1 - \delta)c > 1 - \delta + \delta w_{CC} - (1 - \delta)b$$

because $b > 1$, $c > \delta/(1 - \delta)$ and $w_{DD} - w_{CC} \geq -1$. Since the left-hand side is larger than the right-hand side term by term, (1) is satisfied.

Since DD is strictly risk-dominant in $PD(w)$, by KM (Lemma 5.5), CC is not 0-robust in $PD(w)$. This contradicts Theorem 1. \square

Together Propositions 4 and 5 show that the notion of robustness in repeated games developed in Section 3 can be quite tractable and yields sharp predictions. A corollary of Proposition 4 is that grim-trigger strategies are not the most robust way to sustain cooperation in the repeated PD. This is because grim-trigger strategies punish both players when a unilateral deviation occurs, whereas the stage-dominant strategies we constructed in the proof of Lemma 4 only punish the deviator, while rewarding the other player. This enhances their robustness to incomplete information, and as a result it is possible (if $PD \in \text{int } \mathcal{G}_{DC/CC}$ and $1/(1 - \delta) \leq b + c$) that the profile of grim-trigger strategies is not strongly robust, but cooperation can be sustained by some stage-dominant (and hence strongly robust) PPE.

5 A Folk Theorem in Robust Strategies

In this section, we prove a folk theorem in strongly robust PPEs which is an analogue of Fudenberg, Levine and Maskin (1994, henceforth FLM) but imposes stronger identifiability conditions on the monitoring structure. Under these conditions, we show that every interior point of the set of feasible and individually rational payoff profiles can be sustained by some strongly robust PPE for δ sufficiently close to 1. It implies that when public outcomes are sufficiently informative, then, as δ goes to 1, requiring strong robustness does not impose any

essential restriction on the set of equilibrium payoff profiles. We also provide an example in which the folk theorem in strongly robust PPEs does not hold under FLM's weaker identifiability conditions. This occurs because robustness constraints require us to control continuation payoffs upon joint deviations rather than just unilateral deviations.

The monitoring structure (Y, π) satisfies the *strong full rank condition* if $(\pi(\cdot | a))_{a \in A}$ is linearly independent. The strong full rank condition implies $|Y| \geq |A|$. Conversely, if $|Y| \geq |A|$, then the strong full rank condition is generically satisfied. As its name suggests, the strong full rank condition is more demanding than FLM's pairwise full rank condition.

For each $i \in N$, let

$$V^* = \left\{ v \in \text{co } g(A) \mid \forall i \in N, v_i \geq \min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} g_i(a_i, a_{-i}) \right\}$$

be the set of feasible and individually rational payoff profiles. Note that we use pure-action minimax values since strongly robust PPEs are strict and hence pure. Let $V^{\text{rob}}(\delta)$ denote the set of strongly robust PPE payoff profiles in Γ_G under discount factor δ to emphasize that it depends on δ . The following result holds.

Theorem 3 (folk theorem). *Assume that (Y, π) satisfies the strong full rank condition. For every compact $K \subset \text{int } V^*$, there exists $\underline{\delta} < 1$ such that, for every $\delta > \underline{\delta}$, $K \subseteq V^{\text{rob}}(\delta)$.*

We now describe an example showing that the folk theorem in strongly robust PPEs may fail if the strong full rank condition is not satisfied. Consider the two-by-two game G_0 with action sets $A_1 = A_2 = \{L, R\}$ and public outcomes $Y = \{y_L, y_R, y_M\}$. If both players choose the same action $a \in \{L, R\}$, then signal y_a is realized with certainty. If player 1 chooses L and player 2 chooses action R , then signal y_M is realized with certainty. If player 1 chooses R and player 2 chooses L , then all signals are realized with equal probability. Note that FLM's pairwise full rank condition is satisfied at every action profile, but the strong full

rank condition is not. Expected stage-game payoffs are given by

	L	R
L	3, 3	0, 1
R	1, 0	0, 0

so that minimax values are 0 for both players.¹⁷ The following result holds.

Proposition 6. *The repeated game Γ_{G_0} is such that for any $\delta \in (0, 1)$, if $v \in V^{\text{rob}}(\delta)$, then $v_1 - v_2 \leq 1/2$.*

This implies in particular that $V^{\text{rob}}(\delta)$ is bounded away from $(1, 0)$ so that the folk theorem does not hold in robust strategies for this game. The proof, given in Appendix A.10, is closely related to the argument developed by FLM in their counter-example to the folk theorem when pairwise full rank does not hold. A subtle difference is that FLM are able to construct a counter-example in which equilibrium values are bounded away from the individually rational set in the direction $(1, 1)$. Here, we show that strongly robust equilibrium values are bounded away from the individually rational set in the direction $(1, -1)$. The reason for this is that, upon unilateral deviation, values which enforce LL along the line orthogonal to $(1, 1)$ punish the deviator but reward the player who behaved appropriately. This enforces behavior in stage-dominant strategies. In contrast, upon unilateral deviation, values which enforce RL along the line orthogonal to $(1, -1)$ punish both the deviator and the player who behaved appropriately. This reduces the robustness of RL and enables us to construct a counter-example. If the strong full rank condition held and a fourth informative signal allowed us to identify joint deviations, we could enforce RL in stage-dominant strategies by making continuation payoffs upon such joint deviations particularly low.

¹⁷These expected payoffs can be associated with outcome-dependent realized payoffs $r_i(a_i, y) = 3$ if $y = y_L$, -3 if $(i, a_i, y) = (2, L, y_M)$, 1 if $(i, a_i, y) = (2, R, y_M)$ and 0 otherwise.

6 Alternative classes of perturbations

This paper considers the robustness of PPEs to payoff perturbations that are independent across periods. While we feel that this class of perturbations is a natural and informative benchmark, there are other reasonable classes of perturbations against which to test the robustness of PPEs.

In a number of applied settings, one may be interested in the robustness of PPEs against specific perturbations that have attractive interpretations. For instance the global games perturbations of Carlsson and van Damme (1993) corresponds to an information structure in which players make noisy private assessments of a common state of the world. As long as the perturbations considered are independent across time periods, our results can be thought of as sufficient conditions for equilibria to be robust. In some cases these sufficient conditions may in fact be quite tight. For instance, we know that for static two-by-two games with strict equilibria, risk-dominance is a necessary and sufficient condition for strong robustness, robustness in the sense of KM, as well as robustness against global games perturbations. The one-shot robustness principle suggests that this equivalence extends to repeated two-by-two games, and that robustness against global games perturbations coincides largely with strong robustness in repeated games. The results of Chassang (2007) imply that it is indeed the case for the class of trigger strategies.

Another important direction for research would be to allow for correlation between payoff shocks across time periods.¹⁸ ¹⁹ As the following example shows, including persistent payoff shocks would yield a notion of robustness radically different from the one we study in this paper. In particular, there is no hope of obtaining an extension of the one-shot deviation

¹⁸Note that applying KM's static approach to robustness on the normal form of repeated game Γ_G would allow for permanent payoff shocks. It would also rule out i.i.d. shocks since it requires players to know their entire payoff sequence with high probability.

¹⁹Note that elaborations with persistent payoff shocks are closely related to the type of perturbations studied in the reputation literature. The main difference is that we consider the effects of small perturbations keeping fixed the discount factor, while the reputation literature is interested in the effect of small but fixed perturbations as the discount factor goes to 1. See Mailath and Samuelson (2006) for a review of the reputation literature.

principle. Consider the following stage game G

	L	R
T	$2, 2$	$0, 0$
B	$0, 0$	$1, 1$

Game G is a coordination game such that TL is the unique strongly robust equilibrium. In addition, it can be shown that, for any discount factor δ , repeatedly playing TL is the unique strongly robust PPE of the repeated game Γ_G with perfect monitoring.²⁰ We now show that repeatedly playing TL is not robust if persistent perturbations are allowed. Consider for instance the perturbed game with private payoff shocks such that with probability ε payoffs are

	L	R
T	$2, 2$	$0, 3$
B	$0, 0$	$1, 3$

This corresponds to a permanent perturbation in payoffs. Following history (TR) , the row player will update that with probability 1, her opponent is of the crazy type and will always play R . From then on, the row player's best reply is to play B . This implies that repeatedly playing TL is not an equilibrium outcome following history (TR) . Hence the unique strongly robust PPE is not robust with respect to persistent payoff shocks.²¹

Characterizing equilibria that are robust even in the presence of persistent payoff shocks is delicate but clearly merits further exploration. A preliminary result is that our notion of robustness remains unchanged if we allow for vanishing levels of correlation between payoff perturbations across periods. 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 . Let P denote the players' common

²⁰See Appendix A.11 for a proof. Note that this example shows that the folk theorem in strongly robust PPEs may fail if V^* is not full dimensional even if $n = 2$ and monitoring is perfect. This contrasts with the standard folk theorem (Fudenberg and Maskin, 1986, Theorem 1).

²¹Note that one can construct similar examples if payoff shocks are not permanent (for instance if crazy type column players can become normal with positive probability every period).

prior over sequences $\{\omega_t\}_{t \in \mathbb{N}} \in \prod_{t \in \mathbb{N}} \Omega_t$. Distribution P_t is the marginal induced by P on Ω_t . We say that (\mathbf{U}, P) is a weakly correlated (ε, d) -elaboration of Γ_G whenever for every $t \in \mathbb{N}$, and P -almost every $(\omega_1, \dots, \omega_{t-1}) \in \prod_{\tau=1}^{t-1} \Omega_\tau$,

$$P(\{\omega_t \in \Omega_t \mid \forall i \in N, \forall \omega' \in Q_{it}(\omega_t), |u_{it}(\cdot, \omega') - g_i| \leq d\} \mid \omega_1, \dots, \omega_{t-1}) \geq 1 - \varepsilon.$$

Our robustness result extends to weakly correlated elaborations in the sense that, if a PPE is d -robust, then for any $d' < d$, it is also robust with respect to weakly correlated (ε, d') -elaborations for ε small enough.²²

7 Conclusion

This paper develops a notion of robustness to incomplete information that is adapted to the analysis of repeated games. Our approach to robustness is more restrictive than that of KM but closely related. Because perturbations in future periods may induce small changes in current expected payoffs with large probability, we allow perturbed payoffs to be slightly different (rather than identical) from those of the original game. Appendix C shows that this notion of strong robustness can be extended to multistage games with discounting.

Our main theoretical results show that the strong robustness of PPEs can be related to the strong robustness of the corresponding one-shot action profiles in appropriately augmented stage games. Specifically, we prove a one-shot robustness principle, analogous to the one-shot deviation principle, and show that it implies a factorization result for strongly robust PPEs that parallels the results of APS. We show the value of these characterizations by means of two examples. First, we compute explicitly the set of strongly robust PPE payoff profiles in the repeated Prisoners' Dilemma. We show that cooperation can be robustly

²²In the context of repeated games with imperfect public monitoring, one may be interested in considering i.i.d. perturbations in the monitoring structure. It can be shown that strong robustness is strictly more restrictive.

sustained if and only if both $(Cooperate, Cooperate)$ and $(Defect, Cooperate)$ are enforceable in PPE. Our analysis also emphasizes that grim-trigger strategies are 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. Second, we prove a folk theorem in robust PPEs for repeated games with imperfect public monitoring. We show that for discount factors close enough to 1, and under a strong full rank condition, any feasible and individually rational payoff profile can be achieved by a strongly robust PPE. The identifiability condition we use is stronger than that of FLM because robustness requires us to control all continuation payoffs upon joint deviations, rather than just upon unilateral deviations. With respect to applications, we believe that taking into account robustness issues can refine our understanding of dynamic games in important ways. When robustness is carefully considered, a player's payoff when others deviate will matter for the robustness of strategies. As Chassang and Padro i Miquel (2008) show in the context of military deterrence, this can lead to richer analysis and new comparative statics.

The particular class of perturbations we consider is important for our results to hold. One question is whether strongly robust equilibria are robust with respect to more general perturbations, for instance, allowing for significant intertemporal correlation, or doing away with the common prior assumption. Another question is how much selection could be achieved by considering more specific perturbations with appealing micro-foundations, such as global games perturbations. While we think that the set of perturbations we use is natural benchmark, especially for repeated games, we believe that studying robustness against other classes of elaborations is an important direction for future research.

A Proofs

A.1 Proof of Lemma 1

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 a $(0, d)$ -elaboration of G , G admits an equilibrium arbitrarily close to a^* . This implies that a^* is also a Nash equilibrium of G' , thus a $2d$ -strict equilibrium of G .

A.2 Proof of Proposition 1

The proof is by contradiction, and follows the structure of KM (Proposition 3.2). It uses Lemmas 5 and 6, which are of independent interest and given below.

Definition 7 (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 $\tilde{\omega} = a$, $\tilde{P}(\tilde{\omega}) = P^{\alpha^*}(a)$,
- (iii) $\tilde{Q}_i = \{\{a_i\} \times A_{-i} \mid a_i \in A_i\}$ and
- (iv) for $\tilde{\omega} \in \{a_i\} \times A_{-i}$,

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

if the denominator on the right-hand side is nonzero, and $\tilde{u}_i(\cdot, \tilde{\omega}) = g_i$ otherwise.²³

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

Lemma 5 (canonic normalization with respect to a Bayesian-Nash equilibrium). *Let \tilde{U} be the canonic normalization of U with respect to α^* . Then we have the following.*

- (i) *If U is an (ε, d) -elaboration of G with payoffs bounded by M , then \tilde{U} is an $(\tilde{\varepsilon}, \tilde{d})$ -elaboration of G , where $\tilde{\varepsilon} = n\varepsilon^{1/2}$ and $\tilde{d} = d + \varepsilon^{1/2}(|g| + M)$.*

²³The denominator is nonzero \tilde{P} -almost surely.

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

Proof. (ii) follows directly from the definition of the canonic normalization.

For (i), let

$$\Omega_d = \{\omega \in \Omega \mid \forall i \in N, \forall \omega' \in Q_i(\omega), |u_i(\cdot, \omega') - g_i| \leq d\}.$$

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

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

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

For $\tilde{\omega} = a \in A'$, $i \in N$ and $\tilde{\omega}' \in \tilde{Q}_i(\omega) = \{a_i\} \times A_{-i}$, we have

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

if $\sum_{\omega \in \Omega} \alpha_i^*(\omega)(a_i)P(\omega) > 0$, and $|\tilde{u}_i(\cdot, \tilde{\omega}') - g_i| = 0 \leq \tilde{d}$ otherwise.

In the case of $\varepsilon = 0$, we have $\tilde{P}(A') = 1$ since $A'_i = A_i$ for every $i \in N$. In the case of $\varepsilon > 0$, for each $a_i \in A_i \setminus A'_i$, we have

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

Summing up both sides for all $a_i \in A_i \setminus A'_i$, we have

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

thus $\tilde{P}((A_i \setminus A'_i) \times A_{-i}) \leq \varepsilon^{1/2}$. Thus, $\tilde{P}(A') \geq 1 - \sum_i \tilde{P}((A_i \setminus A'_i) \times A_{-i}) \geq 1 - n\varepsilon^{1/2} = 1 - \tilde{\varepsilon}$. \square

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 6 (locally unique equilibrium). *If a^* is the unique correlated equilibrium of G and is a strict equilibrium, then there exists $d > 0$ such that the unique Bayesian-Nash equilibrium of any $(0, d)$ -elaboration of G is to play a^* with probability 1.*

Proof. The proof is by contradiction. Assume that, for any $d > 0$, there exist a $(0, d)$ -elaboration $U_d = (N, \Omega_d, P_d, (A_i, u_{id}, Q_{id})_{i \in N})$ of G and a Bayesian-Nash equilibrium α_d of U_d such that $P_d^{\alpha_d}(a^*) < 1$. Since the canonic normalization of a $(0, d)$ -elaboration of G is also a $(0, d)$ -elaboration of G by Lemma 5, without loss of generality, we can assume that U_d takes a canonic form, and that α_d is truthtelling.

Since $P_d(a^*) < 1$, we define $\mu_d \in \Delta(A \setminus \{a^*\})$ by

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

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

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

whenever $\omega \in \{a_i\} \times A_{-i}$. As d goes to 0, payoff functions $u_d(\cdot, \omega)$ converge to g for every $\omega \in A$. Since $\mu_d \in \Delta(A \setminus \{a^*\})$, which is compact, as d goes to 0, we can extract a sequence of μ_d that converges to $\mu_0 \in \Delta(A \setminus \{a^*\})$. By continuity, we have that, for all $i \in N$, $a_i \in A_i \setminus \{a_i^*\}$ and $a'_i \in A_i$,

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

We now use distribution μ_0 to build a correlated equilibrium of G distinct from a^* . For $0 \leq q < 1$ define $\mu \in \Delta(A)$ by $\mu(a^*) = q$ and $\mu(a) = (1 - q)\mu_0(a)$ for every $a \in A \setminus \{a^*\}$. It follows from the family of inequalities (2) and the fact that a^* is a strict equilibrium of G that, for q close enough to 1, μ is a correlated equilibrium of G . This contradicts the premise that a^* is the unique correlated equilibrium of G . \square

We use ε -Bayesian-Nash equilibrium in the ex-ante sense. That is, α^* is an ε -Bayesian-

Nash equilibrium of U if

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

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

Proof of Proposition 1. By Lemma 6, we know that there exists $d > 0$ such that a^* is the unique Bayesian-Nash equilibrium of any $(0, d)$ -elaboration of G . Fix such d . 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 Bayesian-Nash equilibrium of U_ε induces probability less than $1 - \eta$ on a^* . Pick any such equilibrium α_ε . Without loss of generality, we can assume that there exists $M > 0$ such that $|u_\varepsilon| < M$ for all $\varepsilon > 0$. Let \tilde{U}_ε be the canonic normalization of U_ε with respect to α_ε . By Lemma 5, truthtelling is a Bayesian-Nash equilibrium of \tilde{U}_ε , $\tilde{P}_\varepsilon(a^*) < 1 - \eta$, and \tilde{U}_ε is an $(\tilde{\varepsilon}, \tilde{d})$ -elaboration of G , where $\tilde{\varepsilon} = n\varepsilon^{1/2}$ and $\tilde{d} = d + \varepsilon^{1/2}(|g| + M)$.

Consider the game \hat{U}_ε identical to \tilde{U}_ε except that $\hat{u}_{i\varepsilon}(\cdot, \omega) = g_i$ whenever $|\tilde{u}_{i\varepsilon}(\cdot, \omega) - g_i| > \tilde{d}$. By an argument identical to KM (Lemma 3.4), truthtelling is a $2M\tilde{\varepsilon}$ -Bayesian-Nash equilibrium of \hat{U}_ε . Note that game \hat{U}_ε is a $(0, \tilde{d})$ -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 convergent sequence of $(0, \tilde{d})$ -elaborations \hat{U}_ε such that $\hat{P}_\varepsilon(a^*) < 1 - \eta$. Denote by \hat{U}_0 the limit of the sequence.

By continuity, \hat{U}_0 is a $(0, d)$ -elaboration of G , truthtelling is a Bayesian-Nash equilibrium of \hat{U}_0 , and $\hat{P}_0(a^*) \leq 1 - \eta$. This contradicts the premise that a^* is the unique Bayesian-Nash equilibrium of all $(0, d)$ -elaborations. \square

A.3 Proof of Proposition 2

The proof of Proposition 2 is almost the same as that of Proposition 1. The only difference is to replace Lemma 6 by the following.

Lemma 7 (locally unique equilibrium for fixed d). *If a^* is the iteratively d -dominant equilibrium of G , then the unique Bayesian-Nash equilibrium of any $(0, d/2)$ -elaboration of G is to play a^* with probability 1.*

The proof of this lemma is straightforward, and hence omitted.

A.4 Proof of Proposition 3

We define the following notion.

Definition 8 ((\mathbf{p}, d) -dominance). For $d \geq 0$ and $\mathbf{p} = (p_1, \dots, p_n) \in (0, 1]^n$, an action profile a^* is a (\mathbf{p}, d) -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 $\lambda(a_{-i}^*) \geq p_i$,

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

If a^* is strictly \mathbf{p} -dominant with $\sum_i p_i < 1$, then it is (\mathbf{q}, d) -dominant for some \mathbf{q} with $\sum_i q_i < 1$ and some $d > 0$. Proposition 3 follows from the following lemma.

Lemma 8. *If a^* is (\mathbf{p}, d) -dominant with $\sum_i p_i < 1$, then it is $d/2$ -robust.*

Proof. Since a^* is (\mathbf{p}, d) -dominant, for all $i \in N$, $a_i \in A_i \setminus \{a_i^*\}$ and $\lambda \in \Delta(A_{-i})$ such that $\lambda(a_{-i}^*) \geq p_i$,

$$\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i, a_{-i}) \quad (3)$$

whenever $|g' - g| \leq d/2$.

For any $(\varepsilon, d/2)$ -elaboration $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ of G , let us define

$$\Omega_{d/2} = \{\omega \in \Omega \mid \forall i \in N, \forall \omega' \in Q_i(\omega), |u_i(\cdot, \omega') - g_i| \leq d/2\}.$$

By the definition of $(\varepsilon, d/2)$ -elaborations, we have that $P(\Omega_{d/2}) \geq 1 - \varepsilon$. As in KM, we are now interested in the set of states where event $\Omega_{d/2}$ is common \mathbf{p} -belief, which we denote by $C^{\mathbf{p}}(\Omega_{d/2})$. Proposition 4.2 (the critical path result) of KM implies that

$$P(C^{\mathbf{p}}(\Omega_{d/2})) \geq 1 - (1 - P(\Omega_{d/2})) \frac{1 - \min_i p_i}{1 - \sum_i p_i}.$$

Since $\sum_i p_i < 1$, for any $\eta > 0$, there exists $\varepsilon > 0$ small enough such that, for any $(\varepsilon, d/2)$ -elaboration U , $P(C^{\mathbf{p}}(\Omega_{d/2})) \geq 1 - \eta$. By (3) and KM (Lemma 5.2), U admits an equilibrium α^* such that $\alpha_i^*(\omega)(a_i^*) = 1$ for all $\omega \in C^{\mathbf{p}}(\Omega_{d/2})$. Equilibrium α^* satisfies $P^{\alpha^*}(a^*) \geq P(C^{\mathbf{p}}(\Omega_{d/2})) \geq 1 - \eta$, which concludes the proof. \square

A.5 Proof of Lemma 2

Fix any $t^0 \geq 1$ and $h^0 \in H_{t^0-1}$. Consider $\mathbf{U} = \{U_t\}$ such that $U_t = G$ for $t \neq t^0$ and $U_{t^0} = G' = (N, (A_i, g'_i)_{i \in N})$ 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 a $(0, d)$ -elaboration of G , $\Gamma_{\mathbf{U}}$ admits a PPE arbitrarily close to s^* . This implies that $s^*(h^0)$ is a Nash equilibrium of $G'(w_{s^*, h^0})$, thus a $2(1 - \delta)d$ -strict equilibrium of $G(w_{s^*, h^0})$.

A.6 Proof of Theorem 1

For an incomplete-information game $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ and $w: Y \rightarrow \mathbb{R}^n$, let $U(w)$ be the incomplete-information game with payoffs $(1 - \delta)u_i(a, \omega) + \delta \mathbb{E}[w_i(y)|a]$ for every $i \in N$, $a \in A$ and $\omega \in \Omega$. For a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete information games, a strategy profile σ of $\Gamma_{\mathbf{U}}$ and a history $h \in H$, let $w_{\sigma, h}$ be the contingent-payoff profile given by $w_{\sigma, h}(y) = (v_i(\sigma|h, y))_{i \in N}$ for each $y \in Y$. A strategy profile σ^* is a PPE of $\Gamma_{\mathbf{U}}$ if and only if $\sigma^*(h_{t-1}, \cdot)$ is a Bayesian-Nash equilibrium of $U_t(w_{\sigma^*, h_{t-1}})$ for all $h_{t-1} \in H$.

For the ‘‘only if’’ part, suppose that s^* is a d -robust PPE of Γ_G for some $d > 0$. By Lemma 2, $s^*(h)$ is a $2(1 - \delta)d$ -strict equilibrium of $G(w_{s^*, h})$ for every $h \in H$.

Pick any $t^0 \geq 1$ and $h^0 \in H_{t^0-1}$. We want to show that $s^*(h^0)$ is $(1 - \delta)d$ -robust in $G(w_{s^*, h^0})$. That is, for every $\eta > 0$, there exists $\varepsilon > 0$ such that every $(\varepsilon, (1 - \delta)d)$ -elaboration of $G(w_{s^*, h^0})$ has a Bayesian-Nash equilibrium that puts probability at least $1 - \eta$ on $s^*(h^0)$. Let $M = |g| + (1 - \delta)d$ and $M' = |g| + d$. Fix any $\eta > 0$. Without loss of generality, we can assume η to be small enough so that

- for every $t^1 > t^0$, $h^1 \in H_{t^1-1}$ and $\mathbf{U} = \{U_t\}$ with $|\mathbf{U}| < M'$ and $U_t = G$ for all $t \neq t^0$, if a strategy profile σ of $\Gamma_{\mathbf{U}}$ puts probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$, then $|w_{\sigma, h^1} - w_{s^*, h^1}| \leq (1 - \delta)d$, and
- if a^* is a $2(1 - \delta)^2d$ -strict equilibrium of some $G' = (N, (A_i, g'_i)_{i \in N})$, then G' has no other Nash equilibria in the η -neighborhood of a^* .

Since s^* is d -robust, there exists $\varepsilon > 0$ such that, for every sequence $\mathbf{U} = \{U_t\}$ of (ε, d) -elaborations of G with $|\mathbf{U}| < M'$, $\Gamma_{\mathbf{U}}$ has a PPE that puts probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$. Fix such ε . Pick any $(\varepsilon, (1 - \delta)d)$ -elaboration U of $G(w_{s^*, h^0})$. Without loss of generality, we can assume that payoffs in U are bounded by M . Then there exists an (ε, d) -elaboration U' of G such that $U'(w_{s^*, h^0}) = U$ and payoffs in U' are bounded by M' .

Consider the “one-shot” sequence $\mathbf{U} = \{U_t\}$ such that $U_t = G$ for all $t \neq t^0$ and $U_{t^0} = U'$. Let σ^* be a PPE of $\Gamma_{\mathbf{U}}$ that puts probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$. Note that $\sigma^*(h)$ is a Nash equilibrium of $G(w_{\sigma^*,h})$ for every $h \in H_{t-1}$ with $t \neq t^0$ and $\sigma^*(h^0, \cdot)$ is a Bayesian-Nash equilibrium of $U'(w_{\sigma^*,h^0})$.

We first show that $\sigma^*(h) = s^*(h)$ for every $t > t_0$ and $h \in H_{t-1}$. By the choice of η , we have $|w_{\sigma^*,h} - w_{s^*,h}| \leq (1 - \delta)d$. Then, since $s^*(h)$ is $2(1 - \delta)d$ -strict in $G(w_{s^*,h})$, $s^*(h)$ is $2(1 - \delta)^2d$ -strict in $G(w_{\sigma^*,h})$. Since $G(w_{\sigma^*,h})$ has no other Nash equilibria in the η -neighborhood of $s^*(h)$, $\sigma^*(h) = s^*(h)$.

Then we have $w_{\sigma^*,h^0} = w_{s^*,h^0}$ and hence $\sigma^*(h^0, \cdot)$ is a Bayesian-Nash equilibrium of $U'(w_{\sigma^*,h^0}) = U'(w_{s^*,h^0}) = U$ that puts probability at least $1 - \eta$ on $s^*(h^0)$.

For the “if” part, suppose that there exists $d > 0$ such that, for every $h \in H$, $s^*(h)$ is a d -robust PPE of $G(w_{s^*,h})$. Fix any d' with $0 < d' < d$. We will show that, for every $\eta > 0$ and $M > 0$, there exists $\varepsilon > 0$ such that, for every sequence $\mathbf{U} = \{U_t\}$ of (ε, d') -elaborations of G with $|\mathbf{U}| < M$, $\Gamma_{\mathbf{U}}$ has a PPE σ^* that puts probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$.

Fix any $M > 0$. Pick $\bar{\varepsilon} > 0$ and $\bar{\eta} > 0$ such that, for every $t \geq 1$, $h \in H_{t-1}$ and $\mathbf{U} = \{U_t\}$ of (ε, d') -elaborations of G with $|\mathbf{U}| < M$, if strategy profile σ of $\Gamma_{\mathbf{U}}$ puts probability at least $1 - \bar{\eta}$ on $s^*(h')$ for all $h' \in H_{t'-1}$ with $t' > t$, then $|w_{\sigma,h} - w_{s^*,h}| \leq d'$. 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 $a \in A$, since the set of contingent-payoff profiles $w_{s^*,h}$ for all $h \in H$ is a bounded subset of $\mathbb{R}^{n|A|}$, there exists a finite set of histories, $H(a)$, such that $s^*(h) = a$ for every $h \in H(a)$ and, whenever $s^*(h') = a$, then $|w_{s^*,h'} - w_{s^*,h}| \leq d''$ for some $h \in H(a)$.

For each $a \in A$ and $h \in H(a)$, since a is d -robust in $G(w_{s^*,h})$, there exists $\varepsilon_h > 0$ such that every (ε_h, d) -elaboration of $G(w_{s^*,h})$ has a Bayesian-Nash equilibrium that puts probability at least $1 - \eta$ on a . Let $\varepsilon = \min(\bar{\varepsilon}, \min_{a \in A} \min_{h \in H(a)} \varepsilon_h) > 0$. Then, for every $h \in H$, every (ε, d') -elaboration of $G(w_{s^*,h})$ has a Bayesian-Nash equilibrium that puts probability at least $1 - \eta$ 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 with $|\mathbf{U}| < M$. Now we 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 such that $U_t^T = U_t$ for $t \leq T$ and $U_t^T = G$ for all $t > T$. We backwardly construct a PPE σ^T of $\Gamma_{\mathbf{U}^T}$ as follows.

- For $h \in H_{t-1}$ with $t > T$, let $\sigma^T(h) = s^*(h)$.
- For $h \in H_{t-1}$ with $t \leq T$, let $\sigma^T(h, \cdot)$ be a Bayesian-Nash equilibrium of $U_t(w_{\sigma^T, h})$

that puts probability at least $1 - \eta$ on $s^*(h)$. Such a Bayesian-Nash equilibrium exists because $\sigma^T(h', \cdot)$ puts probability at least $1 - \eta$ on $s^*(h')$ for all $h' \in H_{t'-1}$ with $t' > t$ and thus $|w_{\sigma^T, h} - w_{s^*, h}| \leq d'$. Therefore, $U_t(w_{\sigma^T, h})$ is an (ε, d') -elaboration of $G(w_{s^*, h})$.

Since the set of all public-strategy profiles is a compact metric space in the product topology, let σ^* be the limit of $\{\sigma^T\}_{T \in \mathbb{N}}$ (take a subsequence if necessary). That is, $\sigma^T(h, \omega_t) \rightarrow \sigma^*(h, \omega_t)$ as $T \rightarrow \infty$ pointwise 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 $\Gamma_{\mathbf{U}}$. By the construction of σ^* , $\sigma^*(h, \cdot)$ puts probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$.

A.7 Stahl's Characterization

Here we summarize the results of Stahl (1991), which characterize V^{PPE} , the set of PPE payoff profiles of Γ_{PD} , as a function of its parameters b , c and δ . Given (b, c, δ) , we define the following parameters.

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

Let us denote

V_0 the set of feasible and individually rational values of G : $V_0 = \text{co}\{(0, 0), (1, 1), (0, p), (p, 0)\}$;

V_Q the set of values defined by $V_Q = \text{co}\{(0, 0), (1, 1), (0, q), (q, 0)\}$;

V_T the set of values defined by $V_T = \text{co}\{(0, 0), (0, b - c), (b - c, 0)\}$;

V_D the set of values defined by $V_D = \text{co}\{(0, 0), (1, 1)\}$.

The following result holds.

Lemma 9 (Stahl (1991)).

(i) If $\delta \geq \max\{(b-1)/b, c/(c+1)\}$, then $V^{\text{PPE}} = V_0$.

- (ii) If $b-1 \leq c \leq h$ and $\delta \in [(b-1)/b, c/(c+1))$, or $c > h$ and $\delta \in [\delta^*, c/(c+1))$, then $V^{\text{PPE}} = V_Q$.
- (iii) If $c < b-1$ and $\delta \in [c/b, (b-1)/b)$, then $V^{\text{PPE}} = V_T$.
- (iv) If $c > h$ and $\delta \in [(b-1)/b, \delta^*)$, then $V^{\text{PPE}} = V_D$.
- (v) If $\delta < \min\{c/b, (b-1)/b\}$, then $V^{\text{PPE}} = \{(0, 0)\}$.

A.8 Proof of Lemma 4

The *PPE Pareto frontier* is the set of $v \in V^{\text{PPE}}$ such that there is no $v' \in V^{\text{PPE}}$ that Pareto-dominates v . We say that a PPE is *Pareto-efficient* if it induces a payoff profile on the PPE Pareto frontier. We begin with the following lemma. We say that $V \subseteq \mathbb{R}^n$ is *self-generating with respect to* $\text{co } B$ if $V \subseteq \text{co } B(V)$. (Recall that $B(V)$ is the set of all payoff profiles that are (not necessarily robustly) generated by V .)

Lemma 10 (PPE Pareto frontier of games in $\mathcal{G}_{DC/CC}$). *Let $\text{PD} \in \mathcal{G}_{DC/CC}$. The following hold.*

- (i) *The PPE Pareto frontier is self-generating with respect to $\text{co } B$.*
- (ii) *No Pareto-efficient PPE prescribes outcome DD on the equilibrium path.*
- (iii) *The PPE Pareto frontier can be sustained by PPEs that prescribe outcome CC permanently along the equilibrium play once it is prescribed, and that never prescribe outcome DD on or off the equilibrium path.*

Proof. From Stahl's characterization, we know that the set of PPE values of Γ_{PD} takes the form $V^{\text{PPE}} = \text{co}\{(0, 0), (1, 1), (0, \phi), (\phi, 0)\}$, where $\phi \geq 1$. We begin with point (i). Pick a Pareto-efficient PPE s^* . Note that continuation payoff profiles of s^* on the equilibrium path are always on the PPE Pareto frontier (otherwise, replacing the continuation strategies by a Pareto-dominating PPE would improve on s^*). In what follows, we modify s^* so that continuation values are on the PPE Pareto frontier even off the equilibrium path. This is possible because points $(0, \phi)$ and $(\phi, 0)$ belong to the PPE Pareto frontier. Consider strategy profile \hat{s}^* that coincides with s^* on the equilibrium path, but such that, whenever player 1 deviates, continuation values are $(0, \phi)$, and whenever player 2 deviates alone, continuation values are $(\phi, 0)$. Since 0 is the minimax value for both players, the fact that s^* is a PPE

implies that \hat{s}^* is also a PPE. This shows that the PPE Pareto frontier is self-generating with respect to $\text{co } B$.

Let us turn to point (ii). Consider a Pareto-efficient PPE s^* . If there is an equilibrium history h at which DD is taken, then, the strategy profile \hat{s}^* obtained by skipping the history and instead playing as if the next period had already been reached is also a PPE and Pareto-dominates s^* . Hence, action DD is never used on the equilibrium path.²⁴

We now proceed with point (iii). From point (i), we know that the PPE Pareto frontier is self-generating with respect to $\text{co } B$. Since we have public randomization, this implies that the PPE Pareto frontier can be generated by PPEs whose continuation payoff profiles are always extreme points of the convex hull of the PPE Pareto frontier. This is the bang-bang property of APS. There are three such points, $(0, \phi)$, $(\phi, 0)$ and $(1, 1)$. Because $(1, 1)$ is not the weighted average of action profiles other than CC , this implies that, in any PPE that sustains values $(1, 1)$, outcome CC is played permanently on the equilibrium path. Inversely, when values $(0, \phi)$ are delivered, the current action profile is CD (otherwise, player 1 would get strictly positive value), and, when values $(\phi, 0)$ are delivered, the current action profile is DC . These imply that Pareto-efficient PPEs taking a bang-bang form are such that, once CC is prescribed, it is prescribed forever along the equilibrium play. Also, by point (ii), such PPEs never prescribe DD on or off the equilibrium path. \square

Proof of Lemma 4. Let us consider $\text{PD} \in \text{int } \mathcal{G}_{DC/CC}$. Since, for every PD' sufficiently close to PD , CC is enforced by a PPE of $\Gamma_{\text{PD}'}$ with continuation payoff profile $(1, 1)$ after CC , we have $1 > (1 - \delta)b$.

For any $d \in (0, 1)$, let us denote by PD_d the game

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

By subtracting d from all payoffs and dividing them by $1 - d$, we obtain PD'_d with payoffs

	C	D
C	$1, 1$	$\frac{-c-d}{1-d}, \frac{b-d}{1-d}$
D	$\frac{b-d}{1-d}, \frac{-c-d}{1-d}$	$0, 0$

²⁴If players only play DD following h , one can simply replace the entire continuation equilibrium by some PPE that gives the players strictly positive value.

which is strategically equivalent to PD_d . Since $PD \in \text{int } \mathcal{G}_{DC/CC}$, there exists $\bar{d} \in (0, 1)$ such that, for $d \in (0, \bar{d})$, we have that $PD'_d \in \mathcal{G}_{DC/CC}$. This means that the set of PPE values of $\Gamma_{PD'_d}$ is a quadrangle $\text{co}\{(0, 0), (1, 1), (0, \phi'), (\phi', 0)\}$, where $\phi' \geq 1$. Note that, since DC is enforceable in PPE of $\Gamma_{PD'_d}$, we have $(1 - \delta)\frac{c-d}{1-d} + \delta\phi' \geq 0$. By Lemma 10, we know that the PPE Pareto frontier of $\Gamma_{PD'_d}$ is sustained by a class of PPEs such that continuation payoffs are always on the PPE Pareto frontier, once action profile CC is prescribed, it is prescribed forever along the equilibrium play, and action profile DD is never prescribed on or off the equilibrium path. Let us denote by \mathcal{E} this class of strategy profiles.

Since game PD'_d is strategically equivalent to game PD_d , strategy profiles in \mathcal{E} are also PPEs of Γ_{PD_d} and generate its PPE Pareto frontier. The PPE Pareto frontier of Γ_{PD_d} is obtained by multiplying equilibrium values of $\Gamma_{PD'_d}$ by $1 - d$ and adding d , which we denote by ℓ_d . ℓ_d is the piecewise line that connects (d, ϕ) , $(1, 1)$ and (ϕ, d) , where $\phi = (1 - d)\phi' + d \geq [(1 - \delta)c + d]/\delta$. Note that, in Γ_{PD_d} , continuation payoffs of these PPEs are at least d at all histories.

Let us now show that strategy profiles in \mathcal{E} are also PPEs of Γ_{PD} . This occurs because PD differs from PD_d only in that the payoff profile from DD is $(0, 0)$ rather than (d, d) . Since strategy profiles in \mathcal{E} never use outcome DD and $d > 0$, whenever the one-shot incentive compatibility holds in Γ_{PD_d} , it also holds in Γ_{PD} . Hence strategy profiles in \mathcal{E} are PPEs of Γ_{PD} . Since payoff profiles upon CD , DC and CC are the same in PD and PD_d , \mathcal{E} generates ℓ_d in Γ_{PD} , and continuation payoff profiles of \mathcal{E} in Γ_{PD} are always in ℓ_d . (ℓ_d may not be the PPE Pareto frontier of Γ_{PD} .)

We now reach the final step of the proof. First, permanent defection is strongly robust, and thus $(0, 0) \in V^{\text{rob}}$. Pick any $s^* \in \mathcal{E}$ that attains $v \in \ell_d$. Let us show that there exists \hat{s}^* such that it attains v and $\hat{s}^*(h)$ is iteratively $(1 - \delta)d$ -dominant in $PD(w_{\hat{s}^*, h})$ for $d \in (0, \min\{\bar{d}, b - 1, c, 1 - (1 - \delta)b\})$. For each history h , we modify continuation strategies as follows.

- If $s^*(h) = CD$, then replace off-path continuation-payoff profiles by $w(CC) = w(DC) = w(DD) = (0, 0)$, where $(0, 0)$ is generated by permanent defection. Since $s^* \in \mathcal{E}$, we have that the value from playing CD at h is at least d . This yields that CD is iteratively $(1 - \delta)d$ -dominant in $PD(w_{\hat{s}^*, h})$. If $s^*(h) = DC$, a symmetric change makes DC iteratively $(1 - \delta)d$ -dominant in a game $PD(w_{\hat{s}^*, h})$, where off-path continuation-payoff profiles are set to $(0, 0)$ while on-path continuation-payoff profiles are not changed.
- If $s^*(h) = CC$, then replace off-path continuation-payoff profiles by $w(DD) = (0, 0)$, $w(DC) = (d, \phi)$ and $w(CD) = (\phi, d)$. Since $s^* \in \mathcal{E}$, the on-path continuation-payoff

profile is $(1, 1)$. Since $1 > (1 - \delta)b + d$ and $-(1 - \delta)c + \delta\phi \geq d$, CC is iteratively $(1 - \delta)d$ -dominant in $\text{PD}(w_{\hat{s}^*, h})$.

It results from this that every payoff profile in $\text{co}(\{(0, 0)\} \cup \ell_d) = \text{co}\{(0, 0), (1, 1), (d, \phi), (\phi, d)\}$ is sustained by some PPE that prescribes the iteratively $(1 - \delta)d$ -dominant equilibrium of the corresponding augmented game at every history. By taking d to 0, we obtain that, for every $v \in \{(0, 0), (1, 1)\} \cup \text{int } V^{\text{PPE}}$, there exist $d > 0$ and a PPE with payoff profile v that prescribes the iteratively $(1 - \delta)d$ -dominant equilibrium of the corresponding augmented game at every history. This concludes the proof when $\text{PD} \in \text{int } \mathcal{G}_{DC/CC}$. A similar proof holds when $\text{PD} \in \text{int } \mathcal{G}_{DC}$. \square

A.9 Proof of Theorem 3

Let $\Lambda = \{\lambda \in \mathbb{R}^n \mid |\lambda| = 1\}$ be the set of n -dimensional unit vectors. For each $\lambda \in \Lambda$ and $k \in \mathbb{R}$, let $H(\lambda, k) = \{v \in \mathbb{R}^n \mid \lambda \cdot v \leq k\}$. Following Fudenberg and Levine (1994), for each $\lambda \in \Lambda$ and $\delta < 1$, we define the *maximal score* $k(\lambda, \delta)$ by the supremum of $\lambda \cdot v$ such that v is d -robustly generated by $H(\lambda, \lambda \cdot v)$ under discount factor δ with some $d > 0$. (If there is no such v , let $k(\lambda, \delta) = -\infty$.) As in Lemma 3.1 (i) of Fudenberg and Levine (1994), $k(\lambda, \delta)$ is independent of δ , thus denoted $k(\lambda)$. Let $Q = \bigcap_{\lambda \in \Lambda} H(\lambda, k(\lambda))$. Q characterizes the limit of strongly robust PPE payoff profiles as $\delta \rightarrow 1$.

Lemma 11. (i) $V^{\text{rob}}(\delta) \subseteq Q$ for every $\delta < 1$.

(ii) If $\dim Q = n$, then, for any compact subset K of $\text{int } Q$, there exists $\underline{\delta} < 1$ such that $K \subseteq V^{\text{rob}}(\delta)$ for every $\delta > \underline{\delta}$.

We omit the proof, for it only replaces the one-shot deviation principle in the proof of Theorem 3.1 of Fudenberg and Levine (1994) by Theorem 1.

Let e_i be the n -dimensional coordinate vector whose i -th component is 1 and others are 0.

Lemma 12. Suppose that (Y, π) satisfies the strong full rank condition.

(i) $k(\lambda) = \max_{a \in A} \lambda \cdot g(a)$ for any $\lambda \in \Lambda \setminus \{-e_1, \dots, -e_n\}$.

(ii) $k(-e_i) = -\min_{a_{-i} \in A_{-i}} \max_{a_i \in A_i} g_i(a)$.

(iii) $Q = V^*$.

Proof. Fix δ . For (i), first consider the case that λ has at least two nonzero components. Pick any $a^0 \in A$. Let $Y = \{y^1, \dots, y^L\}$ with $L = |Y|$. Arrange $A = \{a^0, a^1, \dots, a^K\}$ in a “lexicographic” order that puts $a_i^0 > a_i$ for $a_i \neq a_i^0$, i.e., $1 = k_n < \dots < k_1 < k_0 = K+1$ such that $k_i = |A_{i+1} \times \dots \times A_n|$ and $i = \min\{j \in N \mid a_j^k \neq a_j^0\}$ for every k with $k_i \leq k < k_{i-1}$. Let $\Pi_i(a^0)$ be a $(k_{i-1} - k_i) \times L$ matrix whose (k, l) -component is $\pi(a^{k_i+k-1})(y^l) - \pi(a_i^0, a_{-i}^{k_i+k-1})(y^l)$. By the strong full rank condition, $\begin{pmatrix} \Pi_i(a^0) \\ \Pi_j(a^0) \end{pmatrix}$ has full row rank for every $i \neq j$.

First, we show that, for every $d > 0$, there exists w such that

$$(1 - \delta)g_i(a^k) + \delta \sum_{y \in Y} \pi(a^k)(y)w_i(y) = (1 - \delta)g_i(a_i^0, a_{-i}^k) + \delta \sum_{y \in Y} \pi(a_i^0, a_{-i}^k)(y)w_i(y) - d$$

for every $i \in N$ and k with $k_i \leq k < k_{i-1}$, and $\lambda \cdot w(y) = \lambda \cdot g(a^0)$ for each $y \in Y$. Note that these conditions are written as a system of linear equations in the following matrix form:

$$\begin{pmatrix} \delta \Pi_n(a^0) & \cdots & O \\ \vdots & \ddots & \vdots \\ O & \cdots & \delta \Pi_1(a^0) \\ \lambda_n I & \cdots & \lambda_1 I \end{pmatrix} \begin{pmatrix} w_n(y^1) \\ \vdots \\ w_n(y^L) \\ \vdots \\ w_1(y^1) \\ \vdots \\ w_1(y^L) \end{pmatrix} = \begin{pmatrix} (1 - \delta)(g_n(a_n^0, a_{-n}^1) - g_n(a^1)) - d \\ \vdots \\ (1 - \delta)(g_n(a_n^0, a_{-n}^{k_{n-1}-1}) - g_n(a^{k_{n-1}-1})) - d \\ \vdots \\ (1 - \delta)(g_1(a_1^0, a_{-1}^{k_1}) - g_1(a^{k_1})) - d \\ \vdots \\ (1 - \delta)(g_1(a_1^0, a_{-1}^{K_1}) - g_1(a^K)) - d \\ \lambda \cdot g(a^0) \\ \vdots \\ \lambda \cdot g(a^0) \end{pmatrix},$$

where I is the identity matrix of size L . Since λ has at least two nonzero components, and $\begin{pmatrix} \Pi_i(a^0) \\ \Pi_j(a^0) \end{pmatrix}$ has full row rank for every $i \neq j$, the matrix

$$\begin{pmatrix} \delta \Pi_n(a^0) & \cdots & O \\ \vdots & \ddots & \vdots \\ O & \cdots & \delta \Pi_1(a^0) \\ \lambda_n I & \cdots & \lambda_1 I \end{pmatrix}$$

has full row rank. Thus the system of equations has a solution w .

Now note that a_1^0 is strictly dominant for player 1 in $G(w)$. More generally, a_i^0 is strictly dominant for player i in $G(w)$ if players $1, \dots, i-1$ follow a_1^0, \dots, a_{i-1}^0 . Thus a^0 is iteratively d -dominant in $G(w)$. By Proposition 2, a^0 is strongly robust in $G(w)$, thus $k(\lambda) \geq \lambda \cdot g(a^0)$. Since this holds for any $a^0 \in A$, we have $k(\lambda) \geq \max_{a \in A} \lambda \cdot g(a)$. The other direction of the inequality is obvious.

Second, suppose that λ is a coordinate vector. Without loss of generality, we assume $\lambda = e_n$. Let $a^0 \in \arg \max_{a \in A} g_n(a)$. Arrange $A = \{a^0, \dots, a^K\}$ as in the first case. Since (Y, π) satisfies the strong full rank condition, $\Pi_i(a^0)$ has full row rank for every $i \in N$. Thus, for every $d > 0$, there exist $\kappa > 0$ and w such that

$$(1 - \delta)g_i(a^k) + \delta \sum_{y \in Y} \pi(a^k)(y)w_i(y) = (1 - \delta)g_i(a_i^0, a_{-i}^k) + \delta \sum_{y \in Y} \pi(a_i^0, a_{-i}^k)(y)w_i(y) - d$$

for every $i < n$ and k with $k_i \leq k < k_{i-1}$,

$$(1 - \delta)g_n(a_n^k, a_{-n}^0) + \delta \sum_{y \in Y} \pi(a_n^k, a_{-n}^0)(y)w_n(y) = (1 - \delta)g_n(a^0) + \delta \sum_{y \in Y} \pi(a^0)(y)w_i(y) - d$$

for every k with $1 \leq k < k_{n-1}$, and $g_n(a^0) - \kappa d \leq w_n(y) \leq g_n(a^0)$. As argued in the previous case, a^0 is iteratively d -dominant in $G(w)$. By Proposition 2, a^0 is $d/2$ -robust in $G(w)$. Also a^0 sustains $v = (1 - \delta)g(a^0) + \delta \mathbb{E}[w(y)|a^0]$ such that $v_n \geq g_n(a^0) - \kappa d$ and $w_n(y) \leq g_n(a^0)$ for every $y \in Y$. Let $v' = v - \kappa d \delta / (1 - \delta) e_n$ and $w'(y) = w(y) - \kappa d / (1 - \delta) e_n$ for every $y \in Y$. Then w' enforces (a^0, v') $d/2$ -robustly, $w'_n(y) \leq v'_n$ for every $y \in Y$, and $v'_n \geq g_n(a^0) - \kappa d / (1 - \delta)$. Since $d > 0$ is arbitrary, we have $k(e^n) \geq g_n(a^0)$. The other direction of the inequality is obvious.

The proof of (ii) is similar to the proof of the second case of (i). The only difference is to use a minimax action profile for each player.

(iii) follows from (i) and (ii). □

Proposition 3 follows from Lemmas 11 and 12.

A.10 Proof of Proposition 6

Fix δ and suppose that $\gamma := \sup\{v_1 - v_2 \mid v \in V^{\text{rob}}(\delta)\} > 1/2$. For any $\varepsilon \in (0, \gamma)$, there exists $v \in V^{\text{rob}}(\delta)$ such that $v_1 - v_2 > \gamma - \varepsilon$ and action profile RL is taken at the initial history.²⁵

²⁵If this is not the case, delete several initial periods. This always improves $v_1 - v_2$ since $g_1(a) \leq g_2(a)$ for all $a \neq RL$.

By Theorem 1, there exist $w(y_L), w(y_R), w(y_M) \in V^{\text{rob}}(\delta)$ that enforce (RL, v) robustly, i.e., such that RL is strongly robust in

$$G(w) = \begin{array}{c|cc} & L & R \\ \hline L & 3(1-\delta) + \delta w_1(y_L), 3(1-\delta) + \delta w_2(y_L) & \delta w_1(y_M), 1-\delta + \delta w_2(y_M) \\ R & v_1, v_2 & \delta w_1(y_R), \delta w_2(y_R) \end{array},$$

where

$$\begin{aligned} v_1 &= 1 - \delta + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)), \\ v_2 &= \frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)). \end{aligned}$$

Let $\gamma(y) = w_1(y) - w_2(y)$ for each $y \in Y$. By definition of γ , we have $\gamma(y) \leq \gamma$ for every $y \in Y$.

Since RL is a strict equilibrium of $G(w)$,

$$\frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)) > \delta w_2(y_R), \quad (4)$$

$$1 - \delta + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)) > 3(1 - \delta) + \delta w_1(y_L). \quad (5)$$

Also, since LR is not strictly $(1/2, 1/2)$ -dominant (KM, Lemma 5.5), either

$$3(1 - \delta) + \delta w_1(y_L) + \delta w_1(y_M) \leq 1 - \delta + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)) + \delta w_1(y_R), \quad (6)$$

or

$$1 - \delta + \delta w_2(y_M) + \delta w_2(y_R) \leq 3(1 - \delta) + \delta w_2(y_L) + \frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)). \quad (7)$$

Assume that (6) holds. Combining (6) and (4), we have $3(1-\delta)/\delta < -\gamma(y_L) + 2\gamma(y_R) - \gamma(y_M)$. This yields that

$$\begin{aligned} \gamma - \varepsilon &< v_1 - v_2 = 1 - \delta + \frac{\delta}{3}(\gamma(y_L) + \gamma(y_R) + \gamma(y_M)) \\ &\leq 1 - \delta + \frac{\delta}{3} \left(-3\frac{1-\delta}{\delta} + 3\gamma(y_R) \right) \leq \delta\gamma, \end{aligned}$$

thus $\gamma < \varepsilon/(1 - \delta)$. Since ε can be arbitrarily small, this contradicts with $\gamma > 1/2$.

Similarly, if (7) holds, combining (7) and (5) yields $3(1-\delta)/\delta < -2\gamma(y_L) + \gamma(y_R) + \gamma(y_M)$. Hence, we have

$$\begin{aligned} \gamma - \varepsilon < v_1 - v_2 &= 1 - \delta + \frac{\delta}{3}(\gamma(y_L) + \gamma(y_R) + \gamma(y_M)) \\ &\leq 1 - \delta + \frac{\delta}{3} \left(-\frac{3}{2} \frac{1-\delta}{\delta} + \frac{3}{2}\gamma(y_R) + \frac{3}{2}\gamma(y_R) \right) \leq \frac{1}{2}(1 - \delta) + \delta\gamma, \end{aligned}$$

thus $\gamma < 1/2 + \varepsilon/(1 - \delta)$. Since ε can be arbitrarily small, this contradicts $\gamma > 1/2$.

A.11 A Failure of the Folk Theorem under Perfect Monitoring

Consider the stage game G

	L	R
T	$2, 2$	$0, 0$
B	$0, 0$	$1, 1$

TL is the unique strongly robust equilibrium of G . We will show that, for each discount factor $\delta < 1$, the repetition of TL is the unique strongly robust equilibrium of the repeated game Γ_G with perfect monitoring. Minimax values are 1 for both players, and $V^* = \{v \in \mathbb{R}^2 \mid 1 \leq v_1 = v_2 \leq 2\}$ is not full dimensional. Let \underline{v} be the infimum of player 1's payoffs in strongly robust equilibria. For any strongly robust equilibrium payoff profile v , by Theorem 1, there exist $a \in A$ and $w: A \rightarrow \mathbb{R}^2$ with $w_1(a) = w_2(a) \geq \underline{v}$ such that a is strongly robust in $G(w)$. Since players receive the same payoff in $G(w)$, it can be shown that a strongly robust equilibrium of $G(w)$ must maximize player 1's payoff (this follows from the risk-dominance condition). This implies that $v_1 \geq (1 - \delta)g_1(TL) + \delta w_1(TL) \geq (1 - \delta)2 + \delta \underline{v}$. Since this holds for any strongly robust equilibrium payoff v_1 , we have $\underline{v} \geq (1 - \delta)2 + \delta \underline{v}$, thus $\underline{v} \geq 2$. Hence the only robust equilibrium value is $v = 2$ and repeatedly playing TL is the only strongly robust PPE of Γ_G . This holds for any discount factor δ .

B Public Randomization

Appendix B extends our framework to allow for public randomization. Given a complete-information game 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 distributed uniformly on $[0, 1)$ and independently of the past history. We write $\theta^t = (\theta_1, \dots, \theta_t) \in [0, 1)^t$, $\tilde{h}_{t-1} = (h_{t-1}, \theta^t) \in \tilde{H}_{t-1} = H_{t-1} \times [0, 1)^t$, and $\tilde{H} = \bigcup_{t \geq 1} \tilde{H}_{t-1}$. A

pure strategy of player i is a mapping $s_i: \tilde{H} \rightarrow A_i$ such that there exists a sequence $\{R_t\}$ of partitions consisting of finitely many subintervals of $[0, 1)$ such that $\tilde{s}_i(h_{t-1}, \cdot)$ is $R_1 \otimes \cdots \otimes R_t$ -measurable on $[0, 1)^t$ for every $h_{t-1} \in H$. Conditional on public history (h_{t-1}, θ^{t-1}) , a strategy profile \tilde{s} induces a probability distribution over sequences of future action profiles, which induces continuation payoffs

$$\forall i \in N, \forall h_{t-1} \in H, \forall \theta^{t-1} \in [0, 1)^{t-1}, \quad v_i(\tilde{s}|(h_{t-1}, \theta^{t-1})) = \mathbb{E} \left[(1 - \delta) \sum_{\tau=1}^{\infty} \delta^{\tau-1} g_i(a_{t+\tau-1}) \right].$$

Let $w_{\tilde{s}, \tilde{h}}$ be the contingent-payoff profile given by $w_{\tilde{s}, \tilde{h}}(y) = (v_i(\tilde{s}|(\tilde{h}, y)))_{i \in N}$ for each $y \in Y$. A strategy profile \tilde{s}^* is a PPE if $v_i(\tilde{s}^*|(h_{t-1}, \theta^{t-1})) \geq v_i(\tilde{s}_i, \tilde{s}_{-i}^*|(h_{t-1}, \theta^{t-1}))$ for every $h_{t-1} \in H$, $\theta^{t-1} \in [0, 1)^{t-1}$, $i \in N$ and \tilde{s}_i . We denote by \tilde{V}^{PPE} the set of PPE payoff profiles of $\tilde{\Gamma}_G$.

Given a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete-information games, we consider the corresponding dynamic game $\tilde{\Gamma}_{\mathbf{U}}$ with public randomization. A mapping

$$\tilde{\sigma}_i: \bigcup_{t \geq 1} (\tilde{H}_{t-1} \times \Omega_t) \rightarrow \Delta(A_i)$$

is a public strategy of player i if there exists a sequence $\{R_t\}$ of partitions consisting of finitely many subintervals of $[0, 1)$ such that $\tilde{\sigma}_i(\tilde{h}_{t-1}, \cdot)$ is Q_{it} -measurable on Ω_t for every $\tilde{h}_{t-1} \in \tilde{H}$, and $\tilde{\sigma}_i(h_{t-1}, \cdot, \omega_t)$ is $R_1 \otimes \cdots \otimes R_t$ -measurable on $[0, 1)^t$ for every $h_{t-1} \in H$ and $\omega_t \in \Omega_t$. A public-strategy profile $\tilde{\sigma}^*$ is a PPE if $v_i(\tilde{\sigma}^*|(h_{t-1}, \theta^{t-1})) \geq v_i(\tilde{\sigma}_i, \tilde{\sigma}_{-i}^*|(h_{t-1}, \theta^{t-1}))$ for every $h_{t-1} \in H$, $\theta^{t-1} \in [0, 1)^{t-1}$, $i \in N$ and public strategy $\tilde{\sigma}_i$ of player i .

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

Definition 9 (dynamic robustness with public randomization). For $d \geq 0$, a PPE \tilde{s}^* of $\tilde{\Gamma}_G$ is d -robust if, for every $\eta > 0$ and $M > 0$, there exists $\varepsilon > 0$ such that, for every sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of (ε, d) -elaborations of G with $|\mathbf{U}| < M$, game $\tilde{\Gamma}_{\mathbf{U}}$ has a PPE $\tilde{\sigma}^*$ such that $P_t^{\tilde{\sigma}^*(\tilde{h}_{t-1}, \cdot)}(\tilde{s}^*(\tilde{h}_{t-1})) \geq 1 - \eta$ for every $t \geq 1$ and $\tilde{h}_{t-1} \in \tilde{H}_{t-1}$.

A PPE s^* is *strongly robust* if it is d -robust for some $d > 0$.

Let \tilde{V}^{rob} denote the set of all payoff profiles of strongly robust PPEs in $\tilde{\Gamma}_G$.

The following is the one-shot robustness principle for repeated games with public randomization.

Proposition 7 (one-shot robustness principle with public randomization). *A strategy profile \tilde{s}^* is a strongly robust PPE of $\tilde{\Gamma}_G$ if and only if there exists $d > 0$ such that, for every $\tilde{h} \in \tilde{H}$, $\tilde{s}^*(\tilde{h})$ is a d -robust equilibrium of $G(w_{\tilde{s}^*, \tilde{h}})$.*

Proof. The proof of the “only if” part is essentially the same as that of Theorem 1, and thus omitted.

The proof of the “if” part is very similar to that of Theorem 1. One difference is in the last step, where we construct a sequence of PPEs $\tilde{\sigma}^T$ of “truncated” games $\tilde{\Gamma}_{\mathbf{U}^T}$, and then take the limit of these PPEs to obtain a PPE of the original game $\tilde{\Gamma}_{\mathbf{U}}$. Here, because \tilde{s}^* is adapted to some sequence $\{R_t\}$ of partitions consisting of finitely many subintervals of $[0, 1)$, we can construct a PPE $\tilde{\sigma}^T$ of $\tilde{\Gamma}_{\mathbf{U}^T}$ truncated at period T such that $\tilde{\sigma}^T(h_{t-1}, \cdot, \omega_t)$ is $R_1 \otimes \cdots \otimes R_t$ -measurable for every $h_{t-1} \in H$ 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 $\tilde{\sigma}^T(h_{t-1}, \theta^t, \omega_t) \rightarrow \tilde{\sigma}^*(h_{t-1}, \theta^t, \omega_t)$ pointwise as $T \rightarrow \infty$ for every $h_{t-1} \in H$, $\theta^t \in [0, 1)^t$ and $\omega_t \in \Omega_t$, and uniformly in θ^t on each cell of $R_1 \otimes \cdots \otimes R_t$ (take a subsequence if necessary). Then $\tilde{\sigma}^*$ is a PPE of $\tilde{\Gamma}_{\mathbf{U}}$. \square

Consider a mapping $\text{co } B^d$ from $V \subseteq \mathbb{R}^n$ to $\text{co } B^d(V)$. Similarly to Lemma 3, $\text{co } B^d$ is monotonic and admits the largest fixed point among all subsets of $\text{co } g(A)$. We denote it by \tilde{V}^d . Then we have the following characterization similar to Theorem 2.

Theorem 4 (characterization of \tilde{V}^{rob}). *If V is a compact convex set such that $\tilde{V}^{\text{rob}} \subseteq V \subseteq \text{co } g(A)$ and $B^d(V) \subseteq V$ for every $d > 0$, then*

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

Proof. For each $v \in \tilde{V}^{\text{rob}}$, let \tilde{s}^* be a strongly robust PPE of $\tilde{\Gamma}_G$ that attains v . Then, by Proposition 7, there exists $d > 0$ such that $\tilde{V}^* = \{v(\tilde{s}^*|(h_{t-1}, \theta^{t-1})) \in \mathbb{R}^n \mid h_{t-1} \in H, \theta^{t-1} \in [0, 1)^{t-1}\}$ is self-generating with respect to $\text{co } B^d$. So we have $v \in \tilde{V}^* \subseteq \tilde{V}^d$. Thus $\tilde{V}^{\text{rob}} \subseteq \bigcup_{d>0} \tilde{V}^d$.

For each $d > 0$, since \tilde{V}^d is self-generating with respect to $\text{co } B^d$, for each $v \in \tilde{V}^d$, there exist $\lambda(v, k) \geq 0$, $a(v, k) \in A$ and $w(v, k, \cdot): Y \rightarrow \tilde{V}^d$ for $k = 1, \dots, K(v)$ such that $\sum_{k=1}^{K(v)} \lambda(v, k) = 1$, $w(v, k, \cdot)$ enforces $(a(v, k), v(k))$ d -robustly for every k , and $v = \sum_{k=1}^{K(v)} \lambda(v, k) v(k)$. 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 $v \in \tilde{V}^d$ and construct \tilde{s}^* recursively as follows. For each θ_1 , let $\tilde{s}^*(\theta_1) = a(v, k(v, \theta_1))$. For each $y_1 \in Y$ and $(\theta_1, \theta_2) \in [0, 1)^2$, let $\tilde{s}^*(y_1, \theta_1, \theta_2) = a(w(v, k(v, \theta_1), y_1), k(w(v, k(v, \theta_1), \theta_2)))$ and so on. By construction, $\tilde{s}^*(\tilde{h})$ is d -robust in $G(w_{\tilde{s}^*, \tilde{h}})$ for every $\tilde{h} \in \tilde{H}$. Then, by Proposition 7, \tilde{s}^* is a strongly robust PPE of $\tilde{\Gamma}_G$ that attains v , and thus $v \in \tilde{V}^{\text{rob}}$. Thus $\tilde{V}^d \subseteq \tilde{V}^{\text{rob}}$.

The proof of the algorithm part is the same as the proof of Theorem 2. \square

C Extension to Dynamic Games

In Appendix C, we show that the one-shot robustness principle (Theorem 1) extends to general dynamic games with discounted stage payoffs. We also show by way of a counter example that the one-shot robustness principle does not hold anymore if we use robustness in the sense of KM.

C.1 Definition

We consider the following class of infinite-horizon dynamic games with public signals. 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 a public outcome $y_h \in Y_h$ with probability $\pi(a_h)(y_h)$, obtain stage-game payoffs $g_{ih}(a_h)$, and move to the next history $(h, y_h) \in H_t$. Thus $H_t = \{(h, y_h) \mid h \in H_{t-1}, y_h \in Y_h\}$. 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 in the entire game, but assume that the number of action profiles available at each stage game is finite and bounded uniformly in h . Equivalently, both the number of active players (those who have multiple available actions) at each history and the number of actions available to each active player are finite and bounded. Our class of dynamic games includes standard repeated games, repeated games with short-run players, and overlapping-generation games. We also assume that stage-game payoffs are bounded uniformly.

A strategy of player i is a mapping s_i defined on H such that $s_i(h) \in A_{ih}$ for each $h \in H$. Conditional on history $h \in H$, a strategy profile s induces a continuation payoff $v_i(s|h)$ for each player i . Let $w_{s,h}$ be the continuation-payoff profile given by $w_{s,h}(y) = (v_i(s|(h, y)))_{i \in N}$ for each $y \in Y_h$. A strategy profile s^* is a PPE if $v_i(s^*|h) \geq v_i(s_i, s_{-i}^*|h)$ for every $h \in H$, $i \in N$ and strategy s_i of player i .

We perturb Γ and consider a collection $\mathbf{U} = \{U_h\}_{h \in H}$ of 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 stage-game payoffs are bounded uniformly. At each history $h \in H$, $\omega_h \in \Omega_h$ is generated according to P_h independently of the past history. A public strategy of player i is a mapping σ_i such that, at every history $h \in H$, $\sigma_i(h, \cdot)$ is a Q_{ih} -measurable mapping from Ω_h to $\Delta(A_{ih})$. The total

payoff is given by the discounted sum of stage-game payoffs $u_{ih}(a_h, \omega_h)$ normalized by $1 - \delta$. We denote this game by $\Gamma_{\mathbf{U}}$.

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

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

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

Note that, in the above definition, we use history-dependent perturbations, a wider class of perturbations than what we use in Definition 4. Nevertheless, the one-shot robustness principle holds for dynamic games as we show below, which implies that the two definitions of strong robustness are equivalent for repeated games.

Note that total payoffs maybe represented using different stage-game payoffs and discount factors. Our notion of dynamic robustness depends which representation is chosen. For example, 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 $t \rightarrow \infty$, 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 8 (one-shot robustness principle for dynamic games). *A strategy profile s^* is a strongly robust PPE of Γ if and only if there exists $d > 0$ such that, for every $h \in H$, $s^*(h)$ is a d -robust equilibrium of $G_h(w_{s^*, h})$.*

The proof is essentially identical to that of Proposition 1.

C.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 illustrating why this strengthening

of robustness is necessary once we look at dynamic games. More specifically, we describe a game and a PPE of that game that is not dynamically robust in an acceptable sense, although its one-shot action profiles are robust (in the sense of KM) in all appropriately augmented stage games.

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

- Time is discrete, with $t \in \{1, \dots, T\}$.
- At each period t , there are two active players X_t and Y_t , which respectively take decisions $x_t \in \{0, 1\}$ and $y_t \in \{0, 1\}$.
- For every $t \in \{0, \dots, T - 1\}$, the payoffs to players X_t and Y_t are given by

$$\begin{array}{c|cc} & y_t = 1 & y_t = 0 \\ \hline x_t = 1 & a - by_{t+1}, 0 & 0, 1 \\ x_t = 0 & 0, 1 & 1, 0 \end{array},$$

where $1 < a < b < (1 + a)^2/4$. The term $-by_{t+1}$ corresponds to payoffs obtained by player X_t in period $t + 1$ conditional on $x_t = 1$ and $y_t = 1$.

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

$$\begin{array}{c|cc} & y_T = 1 & y_T = 0 \\ \hline x_T = 1 & a - b\lambda_T, 0 & 0, 1 \\ x_T = 0 & 0, 1 & 1, 0 \end{array},$$

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

Note that player Y_t 's payoffs depend only on the outcome of period t while player X_t 's payoffs depend on the outcome of periods t and $t + 1$. Such an overlapping-generation game is described by parameters a , b and λ_T . Let us define the function f such that

$$f(\lambda) = \begin{cases} \frac{1}{1+a-\lambda b} & \text{if } 0 \leq \lambda \leq a/b, \\ 1 & \text{if } a/b < \lambda \leq 1. \end{cases}$$

By the conditions on a and b , 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 by $G(\lambda)$ 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 13 (PPEs of Γ_T).

- (i) *In every PPE of Γ_T , player Y_t chooses $y_t = 1$ with probability $f^{T-t+1}(\lambda_T)$ and $y_t = 0$ with probability $1 - f^{T-t+1}(\lambda_T)$.*
- (ii) *If $b > a$ and $\lambda_T = \lambda_M$, then game Γ_T has a unique PPE such that, for all t , player X_t chooses $x_t = 1$ with probability $1/2$, and player Y_t chooses $y_t = 1$ with probability λ_M . This one-shot mixed-action profile is the unique correlated equilibrium of $G(\lambda_M)$.*

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 $f^{T-t}(\lambda_T)$. Then players X_t and Y_t are playing the stage game $G(f^{T-t}(\lambda_T))$. If $f^{T-t}(\lambda_T) \geq a/b$, then playing $x_t = 0$ is weakly dominant for player X_t , and hence, player Y_t chooses $y_t = 1$ with probability $f(f^{T-t}(\lambda_T)) = 1$. If $f^{T-t}(\lambda_T) < a/b$, then game $G(f^{T-t}(\lambda_T))$ has a unique mixed equilibrium in which player Y_t plays $y_t = 1$ with probability $f(f^{T-t}(\lambda_T)) = f^{T-t+1}(\lambda_T)$. The second result follows from the fact that λ_M is a fixed-point of f strictly below a/b . \square

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

Proposition 9 (sensitivity of mixed equilibria to perturbations). *Consider game Γ_T with $\lambda_T = \lambda_M$. The following results hold.*

- (i) *For any $\varepsilon \in (0, b\lambda_M/(1 + b\lambda_M)]$, game Γ_T^ε has a unique subgame-perfect equilibrium in which at any time $t < T$ player Y_t chooses $y_t = 1$ with probability $f^{T-t+1}(\lambda_M^\varepsilon)$, where $\lambda_M^\varepsilon = (1 - \varepsilon)\lambda_M - \varepsilon/b$.*
- (ii) *For any $\varepsilon \in (0, b\lambda_M/(1 + b\lambda_M)]$ and $\eta > 0$, there exists T large enough such that $|f^T(\lambda_M^\varepsilon) - \lambda_L| < \eta$.*

Proof. The proof of point (i) is identical to that of Lemma 13. Point (ii) is a consequence of the fact that λ_L is a stable fixed point of f with basin of attraction $[0, \lambda_M)$. \square

By point (ii) of Lemma 13, the unique PPE of Γ_T with $\lambda_T = \lambda_M$ is such that, at every period, players play the unique correlated equilibrium of the augmented stage game. However, by point (ii) of Proposition 9, the equilibrium can be destabilized by small payoff changes that happen sufficiently far in the future. This result can be leveraged to show that the one-shot robustness principle does not hold under KM's definition of robustness.

More precisely, while point (ii) of Proposition 9 applies to a family of games with increasing length T , a sequence $\{\Gamma_T\}_{T=1}^{\infty}$ of games can simply be regrouped in a single infinite horizon game Γ_{∞} in which players play game Γ_T in the interval of time $\{T(T-1)/2 + 1, \dots, T(T+1)/2\}$. Γ_{∞} has a unique PPE, such that, at each stage, equilibrium profiles are the unique correlated equilibrium of the appropriate augmented stage game. However, that equilibrium is not dynamically robust.

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.

References

- [1] Abreu, D., D. Pearce and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.
- [2] Aghion, P., D. Fudenberg and R. Holden (2007). "Subgame Perfect Implementation with Almost Perfect Information," Unpublished.
- [3] Angeletos, G-M., C. Hellwig and A. Pavan (2007). "Dynamic Global Games of Regime Change: Learning, Multiplicity and Timing of Attacks," *Econometrica*, 75, 711–756.
- [4] Argenziano, R (2008). "Differentiated Networks: Equilibrium and Efficiency," *RAND Journal of Economics*, 39, 747–769.
- [5] Bergemann, D. and S. Morris (2005). "Robust Mechanism Design," *Econometrica*, 6, 1771–1813.
- [6] Bhaskar, V., G. J. Mailath and S. Morris (2008). "Purification in the Infinitely-Repeated Prisoners' Dilemma," *Review of Economic Dynamics*, 11, 515–528.

- [7] Carlsson, H. and E. van Damme (1993). “Global Games and Equilibrium Selection,” *Econometrica* 61, 989–1018.
- [8] Chamley, C. (1999), “Coordinating Regime Switches,” *Quarterly Journal of Economics*, 114, 869–905.
- [9] Chassang, S. (2007). “Fear of Miscoordination and the Robustness of Cooperation in Dynamic Global Games with Exit,” Unpublished.
- [10] Chassang, S. and G. Padro i Miquel (2008). “Conflict and Deterrence under Strategic Risk,” Unpublished.
- [11] Dekel, E. and D. Fudenberg (1990). “Rational Behavior with Payoff Uncertainty,” *Journal of Economic Theory*, 52, 243–267.
- [12] Ely, J. C. and J. Välimäki (2002). “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, 102, 84–105.
- [13] Frankel, D., S. Morris and A. Pauzner (2003). “Equilibrium selection in global games with strategic complementarities,” *Journal of Economic Theory*, 108, 1–44.
- [14] Fudenberg, D., D. M. Kreps and D. K. Levine (1988). “On the Robustness of Equilibrium Refinements,” *Journal of Economic Theory*, 44, 354–380.
- [15] Fudenberg, D. and D. K. Levine (1994). “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, 62, 103–135.
- [16] Fudenberg, D., D. K. Levine and E. Maskin (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039.
- [17] Fudenberg, D. and E. Maskin (1986). “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533–554.
- [18] Fudenberg, D. and E. Maskin (1991). “On the Dispensability of Public Randomization in Discounted Repeated Games,” *Journal of Economic Theory*, 53, 428–438.
- [19] Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge: MIT Press.
- [20] Giannitsarou, C. and F. Toxvaerd (2007). “Recursive Global Games,” Unpublished.

- [21] Goldstein I. and A. Pauzner (2004). “Contagion of Self-Fulfilling Financial Crises due to diversification of investment portfolios,” *Journal of Economic Theory*, 119, 151–183.
- [22] Harsanyi, J. C. (1973). “Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *International Journal of Game Theory*, 2, 1–23.
- [23] Hörner, J. and W. Olszewski (2008). “How Robust is the Folk Theorem?,” forthcoming in the *Quarterly Journal of Economics*.
- [24] Izmalkov, S. and M. Yildiz (2008). “Investor Sentiments,” forthcoming in *American Economic Journal: Microeconomics*.
- [25] Kajii, A. and S. Morris (1997). “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283–1309.
- [26] Mailath, G. J. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford University Press.
- [27] Mailath, G. J. and S. Morris (2002). “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory* 102, 189–228.
- [28] Mailath, G. J. and S. Morris (2006). “Coordination failure in repeated games with almost-public monitoring,” *Theoretical Economics* 1, 311–340.
- [29] Mailath, G. and W. Olszewski (2008). “Folk Theorems with Bounded Recall and (Almost) Perfect Monitoring,” Unpublished.
- [30] Miller, D (2007). “The Dynamic Cost of Ex Post Incentive Compatibility in Repeated Games of Private Information,” Unpublished.
- [31] Morris, S. and H. S. Shin (1998). “Unique Equilibrium in a Model of Self-Fulfilling Attacks,” *American Economic Review* 88, 587–597.
- [32] Morris, S. and H. S. Shin (2003). “Global Games: Theory and Applications,” in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, edited by M. Dewatripont, L. Hansen and S. Turnovsky. Cambridge: Cambridge University Press, 56–114.
- [33] Oury, M. and O. Tercieux (2008). “Continuous Implementation,” Unpublished.

- [34] Oyama, D. and O. Tercieux (2007). “Robust Equilibria under Non-Common Priors,” Unpublished.
- [35] Stahl, D. O. (1991). “The Graph of Prisoners’ Dilemma Supergame Payoffs as a Function of the Discount Factor,” *Games and Economic Behavior* 3, 368–384.
- [36] Weinstein, J. and M. Yildiz (2007). “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.