

Robustness to Incomplete Information in Repeated Games*

Sylvain Chassang Satoru Takahashi
Princeton University Princeton University

November 19, 2009

Abstract

This paper extends the framework of Kajii and Morris (1997) to study the question of robustness to incomplete information in repeated games. We show that dynamically robust equilibria can be characterized using a one-shot robustness principle that extends the one-shot deviation principle. Using this result, we compute explicitly the set of dynamically robust equilibrium values in the repeated prisoners' dilemma. We show that robustness requirements have sharp intuitive implications regarding when cooperation can be sustained, what strategies are best suited to sustain cooperation, and how changes in payoffs affect the sustainability of cooperation. We also show that a folk theorem in dynamically robust equilibria holds, but requires stronger identifiability conditions than the pairwise full rank condition of Fudenberg, Levine and Maskin (1994).

*We are particularly grateful to Drew Fudenberg, Stephen Morris, Larry Samuelson and Olivier Tercieux for comments and advice. We thank Dilip Abreu, Eduardo Faingold, Roger Guesnerie, Johannes Hörner, Marcin Peški, Wolfgang Pesendorfer, Phil Reny, Yuliy Sannikov, Hugo Sonnenschein, Muhamet Yildiz as well as seminar participants at Duke, Johns Hopkins, Paris School of Economics, Penn State, Princeton, Rochester, Rutgers, Toronto, UT Austin, the University of Chicago, as well as the GRIPS International Workshop, the conference the 2008 Stony Brook Game Theory Festival, and the 2nd Singapore Economic Theory Annual Workshop for helpful conversations. Paul Scott and Takuo Sugaya provided excellent research assistance. Chassang: chassang@princeton.edu, Takahashi: satorut@princeton.edu.

1 Introduction

This paper formalizes and explores a notion of robustness to incomplete information in repeated games. We characterize dynamically robust equilibria 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). An important application of our work is that grim-trigger strategies are not the most robust way to sustain cooperation. Selective-punishment strategies – which punish only the most recent offender rather than all players – are more robust than grim-trigger because upon unilateral deviation, they allow to punish the player who defected while rewarding the player who cooperated. Concerns of robustness can also change comparative statics. In particular, diminishing payoffs obtained off of the equilibrium path can make cooperation harder to sustain.

Our notion of robustness to incomplete information extends the framework of Kajii and Morris (1997, henceforth KM) to repeated games. Given a complete-information game G , KM consider incomplete-information games U that are elaborations of G in the sense that with high probability every player knows that her payoffs in U are exactly those in G . A Nash equilibrium of G is robust if, for every elaboration U sufficiently close to G , it is close to some Bayesian-Nash equilibrium of U . Our approach to robustness in repeated games is as follows. Given a repeated game Γ_G with complete-information stage game G , we study dynamic games $\Gamma_{\mathbf{U}}$ given by sequences $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of independent incomplete-information stage games, all of which are elaborations of G .¹ A perfect public equilibrium of Γ_G is dynamically robust if for every sequence \mathbf{U} of elaborations sufficiently close to G it is close to some perfect public equilibrium of $\Gamma_{\mathbf{U}}$.

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 analogous to the one-shot deviation principle.

¹We show in Section 7 that our characterizations are unchanged when elaborations are correlated over time, provided that past private information becomes public sufficiently fast.

More precisely, an equilibrium of Γ_G is dynamically 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 Abreu, Pearce and Stacchetti (1990, henceforth APS). Specifically, equilibrium values sustained by dynamically robust equilibria of Γ_G essentially correspond to the largest fixed point of a robust value mapping that associates future continuation values with current values generated by robust equilibria of corresponding augmented stage games.

Our two main applications highlight the practical value of these characterizations. First, for any discount factor, we compute explicitly the set of dynamically robust equilibrium values in the repeated prisoners' dilemma. We show that, whenever outcome $(Defect, Cooperate)$ can be enforced under complete information, the set of dynamically 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 dynamically robust equilibria shrinks to permanent defection. In addition, we highlight that grim-trigger strategies are not very well suited to sustain robust cooperation. Indeed, selective-punishment strategies that punish only deviators upon unilateral deviation are robust over a larger set of parameter values. The reason for this is that selective punishment strategies manage to punish defectors while rewarding cooperators. We also highlight that comparative statics which hold under complete information may be overturned once robustness becomes a concern.

Second, we show that a folk theorem in dynamically 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. As a corollary, this folk theorem provides an existence result for dynamically robust equilibria

for discount factors close to one. This is useful given that the existence of robust equilibria is not guaranteed in general static games (see for instance Example 3.1 in KM).

Our approach to robustness is closely related to that of KM and has a similar interpretation. Since the pioneering work of Rubinstein (1989) and Carlsson and van Damme (1993), who show that strict equilibria of two-by-two games can be destabilized by arbitrarily small perturbations, the question of robustness to incomplete information has received much attention. Work on this topic is of two kinds. A variety of applied work uses robustness to incomplete information as a criterion for equilibrium selection.² A complementary literature explores robustness to incomplete information to ensure that specific equilibria of interest are robust to reasonable perturbations in the information structure.³ KM, as well as this paper, provide a benchmark for both types of studies by analyzing the robustness of equilibria to all small perturbations in the information structure.⁴ By considering a large class of possible perturbations, rather than focusing on specific ones, this approach provides general sufficient conditions that guarantee the robustness of equilibria, and establishes informative bounds on how much selection can be achieved using perturbations in the information structure.⁵

This paper contributes to the literature on repeated games by highlighting how robustness concerns affect the efficient provision of dynamic incentives. In this sense, our paper extends the work of Giannitsarou and Toxvaerd (2007) or Chassang (2009), who analyze dynamic global games in which the question of efficient punishment schemes does not arise. Giannitsarou and Toxvaerd (2007) show that, in a finite-horizon game with strategic com-

²See, 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.

³See for instance Bergemann and Morris (2005), Oury and Tercieux (2008), or Aghion, Fudenberg and Holden (2008).

⁴KM as well as Monderer and Samet (1989) or this paper consider perturbations that are small from an ex ante perspective. Weinstein and Yildiz (2007) consider perturbations that are close from an interim perspective in the product topology on the universal type space. See Dekel, Fudenberg and Morris (2006), Di Tillo and Faingold (2007), Chen and Xiong (2008) or Ely and Pęski (2008) for recent work exploring in details various topologies on informational types. Note also that KM, as well as this paper, maintain the common prior assumption. Oyama and Tercieux (2007) and Izmalkov and Yildiz (2008) consider incomplete information perturbations that do not satisfy the common prior assumption.

⁵These bounds are tight in the context of two-by-two coordination games since it can be shown that global-game perturbations are in fact most destabilizing.

plementarities, a global-game perturbation à la Carlsson and van Damme (1993) selects a unique equilibrium. Chassang (2009) considers an infinite-horizon exit game and shows that, even though the global-game perturbation does not yield uniqueness, 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 all sequences 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 simplifies the analysis and in particular allows us to do away with the strategic complementarity assumptions that are frequently used in the global games literature (see for instance Frankel, Morris and Pauzner, 2003).

Our framework is similar to the one Bhaskar, Mailath and Morris (2008) use to study dynamic robustness of a specific equilibrium in the repeated prisoners' dilemma. They focus on a 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. In contrast, we follow KM and add payoff shocks that may be correlated across players.

This paper is also 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 on past histories. We consider perturbations that depend neither on strategies nor on past histories.

Finally, much of the refinement literature is concerned with robustness in dynamic games in some form or another. 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 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.

The paper is structured as follows. Section 2 provides a motivating example. Section 3 defines robustness in static games. Section 4 formalizes our notion of dynamic robustness for repeated games and provides the main characterization results. Section 5 applies the results of Section 4 to study how concerns of robustness changes analysis in the repeated prisoners' dilemma. Section 6 proves a folk theorem in dynamically robust equilibria for repeated games with imperfect public monitoring. Section 7 extends our analysis to allow for incomplete information perturbations that do not satisfy the common-prior assumption, as well as persistent payoff shocks. Section 8 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 A Motivating Example

This section illustrates how considering incomplete-information perturbations can enrich the analysis of simple repeated games in realistic ways. We also emphasize the value of a systematic approach to robustness.

2.1 The Repeated Prisoners' Dilemma

Throughout this section, let PD denote the two-player prisoners' dilemma with actions $A_1 = A_2 = \{C, D\}$ and payoffs

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

where $b > 1$, $c > 0$ and $b - c < 2$. Let $A = A_1 \times A_2$. We denote by Γ_{PD} the infinitely repeated version PD with discount factor $\delta \in (0, 1)$. Let $H_t = A^t$ denote histories of length t . We allow players to condition their behavior on a public randomization device but omit it from histories for concision.

The analysis of the repeated prisoners' dilemma is greatly simplified by the penal code approach of Abreu (1988). Without loss of efficiency, to enforce cooperation it is sufficient to consider grim-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). Conditional on the other player cooperating, grim-trigger strategies provide players with the highest incentives to cooperate as well. Under complete information, grim-trigger strategies form a subgame-perfect equilibrium if and only if $\delta/(1 - \delta) \geq b - 1$. In words, cooperation is sustainable whenever the value of future cooperation is greater than the short term gains from deviation. Note that the cost c of cooperating while one's partner is defecting does not affect the sustainability of cooperation.

Throughout the paper we examine the robustness of these insights with respect to small misspecifications in the structure of the game of the kind considered by Rubinstein (1990), Carlsson and van Damme (1993) or Morris and Shin (1998). Does cost c start playing a more significant role in determining the sustainability of cooperation? Does the grim-trigger strategy remain an optimal way to sustain cooperation?

2.2 An Incomplete-Information Perturbation

Consider for instance the following perturbation of Γ_{PD} . In every period t , payoffs depend on an i.i.d. state ω_t uniformly distributed over $\{1, 2, \dots, L\}$ with integer $L \geq 1$. If $\omega_t \in \{1, 2, \dots, L - 1\}$, then players are in a normal state with payoffs given by PD. If $\omega_t = L$, then player 1 is “tempted” to play D with payoffs given by

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

where $B > b/(1 - \delta)$ so that D is clearly a dominant action for player 1 in the temptation state. We assume that player 1 is informed and observes a signal $x_{1,t} = \omega_t$ while player 2 observes only a noisy signal $x_{2,t} = \omega_t - \xi_t$, where ξ_t is an even coin flip over $\{0, 1\}$. We denote by Γ_L this perturbed repeated prisoners’ dilemma. A public strategy σ_i of player i is a mapping $\sigma_i : \bigcup_{t \geq 0} H_t \times \{2 - i, \dots, L\} \rightarrow \Delta(\{C, D\})$. A perfect public equilibrium (PPE) is a perfect Bayesian equilibrium in public strategies.

Fix B and consider $\{\Gamma_L \mid L \geq 1\}$. As L goes to infinity, the players will agree up to any arbitrary order of beliefs that they play the standard prisoners’ dilemma with high probability. The question we ask is as follows: when is it that an equilibrium of the complete information game Γ_{PD} approximately coincides with an equilibrium of the perturbed game Γ_L for L large enough? We formalize this question with the following notion of robustness.

Definition 1 (robustness with respect to Γ_L). A pure subgame-perfect equilibrium s^* of Γ_{PD} is *robust* to the class of perturbed games $\{\Gamma_L \mid L \geq 1\}$ if, for every $\eta > 0$, there exists \bar{L} such that, for every $L \geq \bar{L}$, Γ_L has a PPE σ^* such that $\text{Prob}(\sigma^*(h_{t-1}, \cdot) = s^*(h_{t-1})) \geq 1 - \eta$ for every $t \geq 1$ and $h_{t-1} \in H_{t-1}$.⁶

Proposition 1 (robustness of grim-trigger strategies). *If $\delta/(1 - \delta) > b - 1 + c$, then grim-trigger strategies are robust to $\{\Gamma_L \mid L \geq 1\}$. Conversely, if $\delta/(1 - \delta) < b - 1 + c$, then grim-trigger strategies are not robust to $\{\Gamma_L \mid L \geq 1\}$.*

⁶This definition extends without difficulty to equilibria in mixed strategies.

Note that condition

$$\frac{\delta}{1-\delta} > b - 1 + c \tag{1}$$

corresponds to outcome CC being strictly risk-dominant in the one-shot game augmented with continuation values

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

Section 4 provides a one-shot robustness principle extends this property to more general environments.

Condition (1) highlights that losses c matter as much as the deviation temptation b to determine the robustness of cooperation in trigger strategies. This contrasts with the condition for cooperation to be sustainable under complete information, $1/(1-\delta) \geq b$, where losses c play no role in determining the feasibility of cooperation. As the next section highlights, this difference can matter significantly for applications.

2.3 Implications

2.3.1 Comparative Statics

We now illustrate how considerations of robustness can change comparative statics by means of a simple example. We interpret the repeated prisoners' dilemma as a model of two firms in a joint venture. Each firm can either put all its efforts in the joint venture (cooperate) or redirect some of its efforts to a side project (defect). Imagine that payoffs are parameterized by the degree of interdependence $I \in [0, 1]$ of the two firms, which is exogenously specified by the nature of the joint venture project. Interdependence affects payoffs as follows:

$$\begin{aligned} b &= b_0 - b_1 I, \\ c &= c_0 + c_1 I, \end{aligned}$$

where b_0 , b_1 , c_0 and c_1 are strictly positive, $b_0 - b_1 > 1$ (so that players may be tempted to deviate even when $I = 1$) and $b_0 - c_0 < 2$ (so that cooperation is efficient even when $I = 0$). The greater the degree of interdependence I , the costlier it is for the two firms to function independently. The cost of functioning independently depends on whether the firm abandons the joint venture first or second. In particular, in many realistic environments, one may expect that $c_1 > b_1$, i.e. upon unilateral defection, increased interdependence hurts the defector less than the cooperator.⁷ The question is whether or not greater interdependency facilitates the sustainability of cooperation.⁸

Under complete information, cooperation is sustainable under grim strategies if and only if

$$\frac{\delta}{1 - \delta} \geq b - 1 = b_0 - 1 - b_1 I.$$

Greater interdependence reduces the value of unilateral deviation and hence facilitates the sustainability of cooperation. In contrast, grim-trigger strategies are robust to perturbations $\{\Gamma_L \mid L \geq 1\}$ if and only if

$$\frac{\delta}{1 - \delta} \geq b - 1 + c = b_0 - 1 + c_0 + (c_1 - b_1)I.$$

Hence, if $c_1 > b_1$, then greater interdependence reduces the sustainability of cooperation. Indeed, while greater interdependence diminishes the gains from unilateral deviation, it diminishes the payoffs of the player who still cooperates by an even greater amount. In the perturbed game Γ_L , players second guess each other's move and the losses from cooperating while one's partner is defecting loom large, to the extent that formerly unambiguous comparative statics can be overturned. This preemptive motive for defection does not exist in the complete-information environment, which highlights that taking robustness concerns seriously can significantly refine our intuitions.⁹

⁷This would reasonably be the case if the first mover can prepare better and has time to reduce her dependency on the other firm.

⁸Note that the analysis of Section 5 allows to tackle this question for general strategies and the results described here would be qualitatively similar.

⁹Chassang and Padro i Miquel (2009) make a similar point in the context of military deterrence using a related framework.

2.3.2 Grim Trigger, Selective Punishment and Robustness

A closer look at Condition (1) suggests that grim-trigger strategies may not be the most robust way to sustain cooperation. To see this, it is useful to distinguish predatory and preemptive incentives for defection. Cooperation under grim-trigger is robust to perturbation $\{\Gamma_L | L \geq 1\}$ if and only if

$$\frac{\delta}{1-\delta} > b - 1 + c.$$

Parameter $b - 1$ corresponds to a player's predatory incentives, i.e. her incentives to defect on an otherwise cooperative partner. Parameter c corresponds to a player's preemptive incentives, i.e. her incentives to defect on a partner whom she expects to defect. The role played by $b - 1$ and c in Proposition 1 highlights that making predatory incentives $b - 1$ small is good for robustness, but that making preemptive incentives c high is bad for robustness. While grim-trigger strategies minimize predatory incentives, they also increase preemptive incentives: a player who cooperates while her opponent defects suffers from long term punishment in addition to the short run cost $-c$. More sophisticated strategies that punish defectors while rewarding cooperators might support cooperation more robustly. To make this more specific, we now consider a different class of strategies, which we refer as selective-punishment strategies.

Selective-punishment strategies are described by the following automaton. There are 4 states: cooperation, C ; punishment of player 1, P_1 ; punishment of player 2, P_2 ; and defection, D . In state C prescribed play is CC ; in state P_1 prescribed play is CD ; in state P_2 prescribed play is DC ; in state D prescribed play is DD . If player i deviates unilaterally from prescribed play, then the state moves to P_i . If both players deviate, then the state moves to D . If both players play according to prescribed play, states C and D do not change whereas state P_i remains P_i with probability ρ and moves to C with probability $1 - \rho$. In selective-punishment strategies, players selectively punish a unilateral deviator while rewarding the player who is deviated upon.

Player i 's expected value in state P_i is denoted by v_P and characterized by equation $v_P = -c + \delta(\rho v_P + (1 - \rho)\frac{1}{1-\delta})$. Whenever $\frac{\delta}{1-\delta} \geq \max\{b - 1, c\}$ one can pick ρ such that

selective punishment strategies are in equilibrium. Furthermore, by picking ρ below but close to $1 - c(1 - \delta)/\delta$, one can take value v_P arbitrarily close to 0 in equilibrium.

Proposition 2 (robustness of selective-punishment strategies). *Whenever a pair of selective-punishment strategies forms a strict SPE of Γ_{PD} it is robust to $\{\Gamma_L \mid L \geq 1\}$.*¹⁰

By Propositions 1 and 2, if grim-trigger strategies are robust to $\{\Gamma_L \mid L \geq 1\}$, then so are selective-punishment strategies, but not vice-versa. The intuition for this is best explained by writing explicitly the one-shot game augmented with continuation values in state C :

	C	D
C	$1/(1 - \delta), 1/(1 - \delta)$	$-c + \delta v_R, b + \delta v_P$
D	$b + \delta v_P, -c + \delta v_R$	$0, 0$

where v_R is player j 's expected value in state P_i , and characterized by $v_R = b + \delta(\rho v_R + (1 - \rho)\frac{1}{1 - \delta})$. If the pair of selective punishment strategies forms a strict SPE, it must be that $\delta/(1 - \delta) > c$ and $1/(1 - \delta) > b + \delta v_P$. Since we have that $v_R > 1/(1 - \delta)$ and $v_P \geq 0$ it follows that playing C is a strictly dominant strategy of the augmented one-shot game. Dominant strategies are robust to small amounts of incomplete information.

The reason why selective-punishment strategies dominate grim-trigger strategies is that they allow to decrease v_P while increasing v_R . This reduces both predatory and preemptive incentives to defect. In contrast, grim-trigger strategies reduces predatory incentives but increase preemptive incentives.

2.4 The Need for a General Analysis

The example presented in this section shows that considering the impact of small perturbations in the information structure can suggest new and interesting insights on cooperation. The question remains: how much of this analysis is specific to the class of perturbations that we consider? Would selective-punishment strategies remain more robust than grim-trigger

¹⁰We say that an SPE is strict if at any history, a player's action are a strict best response.

strategies if we considered different classes of perturbation? Can anything be said about more general repeated games? Providing tractable answers to these questions has significant importance since much of the applied work on complete information repeated games focuses exclusively on predatory incentives and trigger strategies – see for instance Rotemberg and Saloner (1986), Bull (1987), Bagwell and Staiger (1990), or Baker, Gibbons and Murphy (1994,2002). Analyzing the implications of robustness concerns in these models may yield significant insights.

The remainder of the paper provides a framework that allows us to study the robustness to incomplete information of any given SPE without committing to a specific perturbation. Since we build on KM and consider robustness to an entire class of unspecified, small enough perturbations, the setup is necessarily quite abstract. However, we provide a characterization of dynamically robust equilibria that makes the analysis tractable and highlights how the intuitions developed in this section generalize. To illustrate the applicability of our results we characterize explicitly the set of dynamically robust equilibrium values in the repeated prisoners’ dilemma for any discount factor, and provide a folk theorem under imperfect public monitoring.

3 Robustness in Static Games

This section defines and characterizes robust equilibria in static games. Section 4 leverages these results by showing that the analysis of robustness in dynamic games can be reduced to the analysis of robustness in families of static games augmented with appropriate continuation values.

3.1 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 for 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 *pure 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 information partition over Ω . Let $|u| = \sup_{\omega \in \Omega} |u(\cdot, \omega)|$.¹¹ 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 $\alpha_i: \Omega \rightarrow \Delta(A_i)$.¹² 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 $P^\alpha(a) = \sum_{\omega \in \Omega} P(\omega) \prod_{i \in N} \alpha_i(\omega)(a_i)$ for each $a \in A$. 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 of G if, with probability at least $1 - \varepsilon$, state ω is such that every player in U knows that her payoff function is within distance d of 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 2 (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 $P^{\alpha^*}(a^*) \geq 1 - \eta$.

¹¹We assume the countability of Ω to guarantee the existence of Bayesian-Nash equilibria.

¹²With 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 .

A pure Nash equilibrium a^* of G is *strongly robust* if it is d -robust for some $d > 0$.¹³

In words, an equilibrium a^* of G is strongly robust if all sufficiently close elaborations of G admit a Bayesian-Nash equilibrium that puts high probability on action profile a^* . Note that 0-robustness corresponds to robustness in the sense of KM.¹⁴

3.2 Sufficient Conditions for Strong Robustness

Because the set of elaborations we consider allows for small shocks with a large probability, a strongly robust equilibrium a^* is necessarily strict. More precisely, the following holds.

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

We now provide sufficient conditions for an equilibrium a^* to be robust. These conditions essentially extend the results of KM to d -robustness with $d > 0$.¹⁵ We begin with the case where a^* is the unique correlated equilibrium of G .

Lemma 2 (strong robustness of unique correlated equilibria). *If a^* is the unique correlated equilibrium of G and a^* is strict, then a^* is strongly robust in G .*

A useful special case is the one where a^* is the only equilibrium surviving iterated elimination of strictly dominated actions. For $d \geq 0$, we say that an action profile a^* is an *iteratively d -dominant equilibrium* of 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

¹³To avoid unnecessary notations, 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 strongly robust.

¹⁴The notion of robustness for static games that we define here is a little more stringent than that of KM. Indeed, 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 large probability. In 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.

¹⁵For additional sufficient conditions ensuring the robustness of equilibria, see Ui (2001) or Morris and Ui (2005).

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 $a_{-i} \in \prod_{j \in N \setminus \{i\}} X_{j,t-1}$.

Lemma 3 (strong robustness of iteratively d -dominant equilibria). *If a^* is iteratively d -dominant in G , then it is $d/2$ -robust in G .*

KM provide an other sufficient condition for robustness which is particularly useful in applied settings. Following KM, for $\mathbf{p} = (p_1, \dots, p_n) \in (0, 1]^n$, we say that an action profile a^* is a **p-dominant equilibrium** of G if

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

for every $i \in N$, $a_i \in A_i$ and $\lambda \in \Delta(A_{-i})$ such that $\lambda(a_{-i}^*) \geq p_i$. An action profile a^* is a **strictly p-dominant equilibrium** of G if it is a strict Nash equilibrium and is **p-dominant** in G .

In words, an action profile a^* is strictly **p-dominant** if and only if every player has strict incentives to play a_i^* when she believes that the action profile of other players will be a_{-i}^* with probability greater than p_i . KM establish that, whenever a^* is a **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.

Lemma 4 (strong robustness of strictly **p-dominant equilibria**). *If a^* is a strictly **p-dominant equilibrium** of G with $\sum_i p_i < 1$, then it is strongly robust.*

In Appendix A.6, 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 **p-dominant equilibrium** with $\sum_i p_i < 1$ then no other action profile is 0-robust. Combined with Lemma 4, this implies that, if a game has a strictly **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 strictly risk-dominant equilibrium is the unique strongly robust equilibrium.

4 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 stochastically independent across periods. We show in Sections 4.2 and 4.3 that dynamically robust equilibria admit a convenient recursive representation. Section 7 extends our results to a larger class of correlated perturbations.

4.1 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) .¹⁶ For each $t \geq 1$, let $H_{t-1} = Y^{t-1}$ be the set of public histories of length $t - 1$, corresponding to possible histories at the beginning of period t . Let $H = \bigcup_{t \geq 1} H_t$ be the set of all finite public histories. A pure public strategy of player i is a mapping $s_i : H \rightarrow A_i$. Conditional on public 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

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

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). The restriction to public strategies corresponds to the assumption that, although player i observes her own actions a_i as well as past stage game payoffs $g_i(a)$ (or perhaps noisy signals of $g_i(a)$), she conditions her behavior only on public outcomes.

¹⁶We 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.

We define perturbations of Γ_G as follows. Consider a sequence $\mathbf{U} = \{U_t\}_{t \in \mathbb{N}}$ of incomplete-information elaborations $U_t = (N, \Omega_t, P_t, (A_i, u_{it}, Q_{it})_{i \in N})$ 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.¹⁷ 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 probability distribution over sequences of future action profiles and states, which allows us to define continuation payoffs $v_i(\sigma|h_{t-1})$ such that

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

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 3 (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 $d > 0$.

In words, we say that a PPE s^* of repeated game Γ_G is robust if every perturbed repeated game with independent and small perturbations admits an equilibrium σ^* that is close to s^* at every public history.

¹⁷Section 7 extends our results to the case of persistent perturbations, provided that past large payoff shocks are sufficiently public.

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 with respect to t . If we required only pointwise convergence of the sequence $\mathbf{U} = \{U_t\}$, i.e. that every U_t approach G , then the robustness criterion would become too restrictive. For example, consider a stage game G with a unique Nash equilibrium a^* , and perturbations $\mathbf{U}^T = \{U_t^T\}_{t \in \mathbb{N}}$ such that U_t^T is identical to G for $t \leq T$ and $u_{it}^T \equiv 0$ for every $i \in N$ and $t > T$. For each $t \geq 1$, U_t^T converges to G as $T \rightarrow \infty$. Since game $\Gamma_{\mathbf{U}^T}$ has a finite effective horizon, it follows from the standard backward induction that players play a^* for the first T periods in every PPE of $\Gamma_{\mathbf{U}^T}$. Thus the only dynamically robust equilibrium of Γ_G would be the repetition of a^* . This is why we focus on uniformly small perturbations.

4.2 A One-Shot Robustness Principle

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.

Given a stage game G and a one-period-ahead continuation-payoff profile $w: Y \rightarrow \mathbb{R}^n$ contingent on public outcomes, 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 $g'_i(a) = g_i(a) + \delta \mathbb{E}[w_i(y)|a]$ 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 appropriately augmented stage game.

Lemma 5 (strictness in augmented games). *If s^* is d -robust in Γ_G , then $s^*(h)$ is $2d$ -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 of the one-shot deviation principle for strongly robust PPEs.

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 $G(w_{s^*,h})$.*

This yields the following corollary.

Corollary 1. *If s^* is a finite-automaton PPE, then s^* is strongly robust if and only if, for every $h \in H$, $s^*(h)$ is strongly robust in $G(w_{s^*,h})$.*

If the stage game G is a two player, two action game and s^ is a finite-automaton PPE of Γ_G , then s^* is strongly robust if and only if, for every $h \in H$, $s^*(h)$ is strictly risk-dominant in $G(w_{s^*,h})$.*

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.

4.3 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 4 (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 $v = g(a) + \delta \mathbb{E}[w(y)|a]$.

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 $v = g(a) + \delta \mathbb{E}[w(y)|a]$ for $a \in A$ and $w: Y \rightarrow V$ such that a is a Nash equilibrium of $G(w)$. We say that V is *self-generating with respect to B^d* if $V \subseteq B^d(V)$. We denote the set of feasible values by $F = \frac{1}{1-\delta} \text{co } g(A)$.

Lemma 6 (monotonicity).

- (i) If $V \subseteq V' \subseteq F$, then $B^d(V) \subseteq B^d(V') \subseteq F$.
- (ii) B^d admits a largest fixed point V^d among all subsets of F .
- (iii) If $V \subseteq F$ and V is self-generating with respect to B^d , then $V \subseteq V^d$.

Note that by definition $B^d(V)$ and V^d are weakly decreasing in d with respect to set inclusion. We characterize V^{rob} using mapping B^d as follows.

Theorem 2 (characterization of V^{rob}). $V^{\text{rob}} = \bigcup_{d>0} V^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (B^d)^k(F)$.

V^{rob} is the limit of the largest fixed points V^d of B^d as d goes to 0. Theorem 2 corresponds to APS's self-generation, factorization and algorithm results (APS, Theorems 1, 2 and 5), which show that the set of all PPE payoff profiles is the largest bounded fixed point of the mapping B and can be computed by iteratively applying B to F . Since we require robust enforcement at every stage, mapping B is replaced by B^d .

5 Robustness in the Repeated Prisoners' Dilemma

In this section, we characterize strongly robust subgame-perfect equilibrium (SPE) payoff profiles in the repeated prisoners' dilemma with perfect monitoring.¹⁸ We show that, whenever outcome (*Defect*, *Cooperate*) can be enforced in a SPE under complete information, the set of strongly robust SPE payoff profiles is essentially equal to the set of SPE payoff profiles under complete information. Inversely, whenever (*Defect*, *Cooperate*) cannot be enforced in

¹⁸Note that, under perfect monitoring, PPEs simply correspond to SPEs.

a SPE under complete information, the set of strongly robust SPEs shrinks to permanent defection.

We also show that selective punishment strategies are more robust than grim trigger strategies. In fact, whenever selective punishment strategies form a strict SPE of the complete information games, then they are strongly robust. However, there exist more sophisticated strategies that can sustain cooperation in circumstances where selective punishment strategies cannot.

As in Section 2, let PD denote the two-player prisoners' dilemma with payoffs

	<i>C</i>	<i>D</i>
<i>C</i>	1, 1	− <i>c</i> , <i>b</i>
<i>D</i>	<i>b</i> , − <i>c</i>	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.¹⁹ We are interested in Γ_{PD} , the repeated prisoners' dilemma with public randomization devices and perfect monitoring.

5.1 Robust Cooperation in Grim-Trigger Strategies

As an illustration, we begin by studying the robustness of grim-trigger strategies as defined in Section 2. We showed that under complete information, grim-trigger strategies form an SPE if and only if $\delta/(1 - \delta) \geq b - 1$. We also showed that grim-trigger strategies are robust to the perturbations $\{\Gamma_{\text{PD}}^L | L \geq 1\}$ considered in Section 2 if and only if $\delta/(1 - \delta) > b - 1 + c$. We now show that this condition guarantees strong robustness (to any small enough perturbation) in the sense of Definition 3.

The proof follows from the one-shot robustness principle (Theorem 1): grim-trigger

¹⁹Formally, the framework of Section 4 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.

strategies form a strongly robust SPE if and only if CC is strongly robust in

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

and DD is strongly robust in

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

The latter property follows from Lemma 3 since D is strictly dominant in the augmented stage game at punishment histories. From Lemma 4 and KM (Lemma 5.5), it follows that grim-trigger 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 $\delta/(1 - \delta) > b - 1 + c$. Returning to the class of perturbations studied in Section 2, this means that whenever grim-trigger strategies are robust to perturbations $\{\Gamma_L | L \geq 1\}$ then they are robust to all small enough perturbations.

5.2 Characterizing Strongly Robust SPEs

Our characterization of strongly robust SPEs in the repeated prisoners' dilemma is in three steps. First, we provide a classification of prisoners' dilemma games under complete information. Then, we prove a fragility result which shows that if total surplus is low, so that a player would never accept to cooperate while the other defects, then the only strongly robust SPE is for players to defect at every history. In contrast, if there is enough surplus, so that one player may accept to cooperate while the other defects, then essentially every SPE value under complete information can be sustained by a strongly robust SPE.

5.2.1 A Classification of Prisoners' Dilemma Games

We classify prisoners' dilemma games according to the enforceability of action profiles. We say that action profile a is *enforceable under complete information* in Γ_{PD} if there exists an SPE of Γ_{PD} that prescribes a at some history.

Definition 5 (classification of prisoners' dilemma games). Fix δ . We define four classes of prisoners' dilemma games, $\mathcal{G}_{DC/CC}$, \mathcal{G}_{DC} , \mathcal{G}_{CC} and \mathcal{G}_\emptyset as follows:

- (i) $\mathcal{G}_{DC/CC}$ is the class of PD such that DC and CC are enforceable under complete information in Γ_{PD} .
- (ii) \mathcal{G}_{DC} is the class of PD such that DC is enforceable under complete information in Γ_{PD} , but CC is not.
- (iii) \mathcal{G}_{CC} is the class of PD such that CC is enforceable under complete information in Γ_{PD} , but DC is not.
- (iv) \mathcal{G}_\emptyset is the class of PD such that neither DC nor CC is enforceable under complete information in Γ_{PD} .

Note that DD is always enforceable under complete information. Stahl (1991) characterizes explicitly the set V^{SPE} of SPE payoff profiles under complete information as a function of parameters δ , b and c (Appendix A.11). See Figure 1 for a representation of classes of prisoners' dilemma games as a function of b and c , for δ fixed.

Stahl (1991) shows that, if $\text{PD} \in \mathcal{G}_{DC/CC}$, then $V^{\text{SPE}} = \text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta}), (0, \phi), (\phi, 0)\}$ with $\phi \geq \frac{1}{1-\delta}$. This means that, for $\text{PD} \in \mathcal{G}_{DC/CC}$, it is possible to punish one player while giving the other one her maximum continuation value. If $\text{PD} \in \mathcal{G}_{DC}$, then $V^{\text{SPE}} = \text{co}\{(0, 0), (0, \frac{b-c}{1-\delta}), (\frac{b-c}{1-\delta}, 0)\}$.²⁰ Finally, we have that if $\text{PD} \in \mathcal{G}_{CC}$, then $V^{\text{SPE}} = \text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta})\}$, and if $\text{PD} \in \mathcal{G}_\emptyset$, then $V^{\text{SPE}} = \{(0, 0)\}$.

²⁰Note that, if $\text{PD} \in \mathcal{G}_{DC}$, then $b > c$.

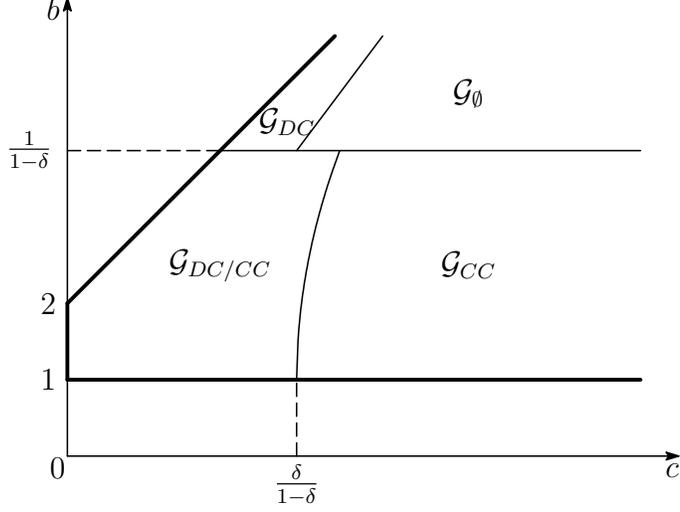


Figure 1: Classification of prisoners' dilemma games

5.2.2 A Fragility Result

The following proposition shows that, if DC is not enforceable under complete information, then the only strongly robust SPE is permanent defection.

Proposition 3 (fragile equilibria). *Fix δ . If $PD \in \mathcal{G}_{CC}$, then the only strongly robust SPE 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 SPE 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 , players have identical continuation values following history (h, a) . Furthermore, we have $c > \delta/(1 - \delta)$; otherwise, DC would be enforceable under complete information.

Given continuation values w , the augmented game $PD(w)$ at history h takes the form

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

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

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

$$\begin{aligned} & [\delta w_{DD} + c - \delta w_{CD}][\delta w_{DD} + c - \delta w_{DC}] \\ & > [1 + \delta w_{CC} - b - \delta w_{CD}][1 + \delta w_{CC} - b - \delta w_{DC}]. \end{aligned} \quad (2)$$

Note that each square bracket term of (2) is nonnegative because CC and DD are Nash equilibria of $PD(w)$. Also note that

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

because $b > 1$, $c > \delta/(1 - \delta)$ and $w_{DD} - w_{CC} \geq -1/(1 - \delta)$. Since the left-hand side is larger than the right-hand side term by term, (2) 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

5.2.3 A Robustness Result

We now show that if DC is enforceable under complete information then V^{rob} is essentially equal to V^{SPE} . Indeed, if action profile DC is enforceable under complete information, then, essentially every payoff profile $v \in V^{\text{SPE}}$ can be sustained by an SPE satisfying the following remarkable property, which we call *iterative stage dominance*.²¹

Lemma 7 (iterative stage dominance). *Fix δ . If either $PD \in \text{int } \mathcal{G}_{DC/CC}$ and $v \in \{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta})\} \cup \text{int } V^{\text{SPE}}$, or $PD \in \text{int } \mathcal{G}_{DC}$ and $v \in \{(0, 0)\} \cup \text{int } V^{\text{SPE}}$, then there exist $d > 0$ and an SPE s^* of Γ_{PD} with payoff profile v such that, for every public history h , $s^*(h)$ is iteratively d -dominant in the augmented game $PD(w_{s^*, h})$.*²²

²¹This 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.

²²We identify a prisoners' dilemma game by its parameters $(b, c) \in \mathbb{R}^2$, so the interior of a class of prisoners'

The detailed proof of Lemma 7, given in Appendix A.12, is lengthy, but the main idea of the argument is straightforward. We show that, for every SPE, its off-path behavior can be modified so that at each history the prescribed action profile is iteratively dominant in the appropriately augmented stage game. The proof exploits the fact that payoff profiles in V^{SPE} allow us to punish one player while rewarding the other.

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

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

Because DD is also an equilibrium of this game, CC is not iteratively dominant. This can be resolved by changing continuation strategies upon outcomes CD and DC . By Stahl's characterization, we know that V^{SPE} takes the form $\text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta}), (0, \phi), (\phi, 0)\}$, where $\phi \geq \frac{1}{1-\delta}$. Consider any public history of the form (CC, \dots, CC, CD) .²³ The grim-trigger strategy prescribes permanent defection. We replace this continuation strategy by an SPE 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 an SPE s_{DC} that attains $(0, \phi)$. Then the augmented game after (CC, \dots, CC) becomes

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

By assumption, CD and DC are enforceable under complete information, so $-c + \delta\phi \geq 0$. Thus C is weakly dominant for both players in this augmented game. Because $\text{PD} \in \text{int}\mathcal{G}_{DC/CC}$, C is in fact strictly dominant. The difficult part of the proof is to show that

dilemma games is derived from the standard topology on \mathbb{R}^2 .

²³We omit public randomizations to simplify notations.

strategy profiles s_{CD} and s_{DC} can be further modified so that their prescribed action profiles become iteratively dominant in corresponding augmented stage games as well.

The following follows directly from Lemma 3, Theorem 1 and Lemma 7.

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

$$\left\{ (0, 0), \left(\frac{1}{1-\delta}, \frac{1}{1-\delta} \right) \right\} \cup \text{int } V^{\text{SPE}} \subseteq V^{\text{rob}} \subseteq V^{\text{SPE}}.$$

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

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

Note that the selective punishment strategies described in Section 2 satisfy iterative stage dominance property of Lemma 7 whenever they form a strict SPE under complete information. Hence whenever they form a strict SPE, selective punishment strategies are a robust way to sustain cooperation. Selective punishment strategies also allow to sustain cooperation under a larger set of parameters (δ, b, c) than grim-trigger. However, it is possible to enforce CD under complete information although $1/(1-\delta) < c$ and selective punishment strategies are not in equilibrium (see Stahl, 1991). In such circumstances, Proposition 4 shows that it is still possible to sustain cooperation robustly, but the strategies used are more sophisticated.²⁴

6 A Folk Theorem in Strongly Robust PPEs

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 requires 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, if public outcomes

²⁴The proof of Lemma 7 provides a description of these strategies. Because $1/(1-\delta) < c$ it is not possible to enforce CD by promising the cooperating player permanent cooperation in the future. However, it may be possible to enforce CD by promising the cooperating player that play will be DC for some periods in the future.

are informative, then, as δ goes to 1, requiring robustness does not impose any essential restriction on the set of equilibrium payoff profiles. A useful corollary is that, for discount factor high enough, if the set of feasible and individually rational payoff profiles is full-dimensional, then there exist strongly robust PPEs. This is a valuable result since the existence of robust equilibria is not guaranteed in static games (see Example 3.1 in KM). 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) has *strong full rank* if $\{\pi(\cdot | a) \in \mathbb{R}^Y | a \in A\}$ 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 all pure action profiles.

Let us define

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

the set of feasible and individually rational values normalized to stage game units. Note that we use pure-action minimax values since strongly robust PPEs are pure. We denote by $NV^{\text{rob}}(\delta) \equiv (1 - \delta)V^{\text{rob}}$ the set of normalized strongly robust PPE payoff profiles in Γ_G given discount factor δ . The normalization by $(1 - \delta)$ ensures that equilibrium values are also expressed in fixed stage game units. The following result holds.

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

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 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 for every pure action profile, but the strong full rank condition is not. Expected stage-game payoffs for game G_0 are given by

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

so that minimax values are 0 for both players.²⁵ The following result holds.

Proposition 5 (failure of the folk theorem). *In the repeated game Γ_{G_0} , for every $\delta \in (0, 1)$, if $(v_1, v_2) \in NV^{\text{rob}}(\delta)$, then $v_1 - v_2 \leq 1/2$.*

This implies that $V^{\text{rob}}(\delta)$ is bounded away from $(1, 0)$ so that the folk theorem does not hold in strongly robust PPEs for this game. The proof, given in Appendix A.14, is closely related to the argument developed by FLM in their counter-example to the folk theorem when the pairwise full rank condition is not satisfied. A subtle difference is that FLM are able to construct a counter-example in which PPE payoff profiles are bounded away from a feasible and individually rational payoff profile in the direction of $(1, 1)$. Here, we show that strongly robust PPE payoff profiles are bounded away from a feasible and individually rational payoff profile in the direction of $(1, -1)$. The reason for this is that, upon unilateral deviation, continuation payoff profiles that enforce LL along the line orthogonal to $(1, 1)$ punish the deviator but reward the player who behaved appropriately. This enforces behavior in dominant actions. In contrast, upon unilateral deviation, continuation payoff profiles that 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

²⁵These 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.

a counter-example. If the strong full rank condition were satisfied and a fourth informative signal allowed us to identify joint deviations, then we could enforce RL in dominant actions by making continuation payoff profiles upon joint deviations particularly low.

7 Extensions

The notion of dynamic robustness we develop in Section 4 depends on the class of perturbations against which we test for robustness. In particular, we assume that players share a common prior and that perturbations are independent across periods. In this section, we discuss ways in which our framework can be extended to accommodate non-common priors and persistent shocks.

7.1 Non-Common Priors

This section considers two different classes of perturbations with non-common priors, depending on how much variation in priors is allowed across players. First, we show that our analysis of robustness to incomplete information is unchanged even if players have priors that are different but close to each other. We then discuss cases in which the players priors may differ significantly.

7.1.1 Approximately Common Priors

Consider an incomplete-information game $(U, (P_i)_{i \in N})$ with non-common priors, where $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ is an incomplete-information game with an “objective” prior P over Ω , and P_i is player i ’s prior over Ω . Let

$$m(P, P_i) = \sup_{\omega \in \Omega} \left| \frac{P_i(\omega)}{P(\omega)} - 1 \right|$$

with a convention that $q/0 = \infty$ for $q > 0$ and $0/0 = 1$. $m(P, P_i)$ measures the proximity between the “objective” prior and player i ’s prior.

Definition 6 (static robustness with almost common priors). For $d \geq 0$, a pure Nash equilibrium a^* of G is *d-robust to incomplete information with almost common priors* if, for every $\eta > 0$ and $M > 0$, there exists $\varepsilon > 0$ such that, for every (ε, d) -elaboration U of G with $|u| < M$ and profile of non-common priors $(P_i)_{i \in N}$ with $m(P, P_i) \leq \varepsilon$ for every $i \in N$, game $(U, (P_i)_{i \in N})$ has a Bayesian-Nash equilibrium α^* such that $P^{\alpha^*}(a^*) \geq 1 - \eta$.

A pure Nash equilibrium a^* of G is *strongly robust to incomplete information with almost common priors* if it is d -robust to incomplete information with almost common priors for some $d > 0$.

The following lemma shows that allowing for non-common priors with small $m(P, P_i)$ does not affect strong robustness in static games.

Lemma 8 (static equivalence of common and almost common priors). *If $d > d' > 0$ and a^* is d -robust to incomplete information with common priors in G , then a^* is d' -robust to incomplete information with almost common priors in G . Hence, strong robustness in the sense of Definition 6 is equivalent to that of Definition 2.*

Oyama and Tercieux (2009, Proposition 5.7) provide a similar result for p -dominant equilibria. We extend the definition of dynamic robustness given in Section 4 as follows.

Definition 7 (dynamic robustness with almost common priors). For $d \geq 0$, a PPE s^* of Γ_G is *d-robust to incomplete information with almost common priors* 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 $|\mathbf{U}| < M$ and every sequence $\{(P_{it})_{i \in N}\}_{t \in \mathbb{N}}$ of non-common priors with $m(P_t, P_{it}) \leq \varepsilon$ for every $i \in N$ and $t \geq 1$, the induced dynamic incomplete-information game with non-common priors 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 to incomplete information with almost common priors* if it is d -robust for some $d > 0$.

Similarly to Theorem 1, the one-shot robustness principle holds. Namely, a PPE is strongly robust to incomplete information with almost common priors in Γ_G if and only if

there exists $d > 0$ such that, for every $h \in H$, $s^*(h)$ is d -robust to incomplete information with almost common priors in $G(w_{s^*,h})$. Therefore, Theorem 1 and Lemma 8 imply the following.

Proposition 6 (dynamic equivalence of common and almost common priors). *If a PPE is strongly robust to incomplete information with common priors, then it is also strongly robust to incomplete information with almost common priors. Hence, strong robustness in the sense of Definition 7 is equivalent to that of Definition 3.*

7.1.2 General Non-Common Priors

In the case where players have significantly different priors $(P_i)_{i \in N}$, in the sense that $m(P, P_i)$ is large, robustness to such perturbations is a much more stringent requirement than robustness to common-prior perturbations. In a generic static game, Oyama and Tercieux (2009) show that a Nash equilibrium is robust to incomplete information with non-common priors if and only if it is iteratively dominant. One can extend their result to dynamic settings and show that a PPE is strongly robust to incomplete information with non-common priors if and only if it is iteratively stage-dominant. Some of our results still apply in this case. For instance, in the repeated prisoners' dilemma, our characterization of strongly robust SPE payoff profiles relies on iterative stage dominance (Lemma 7). As a consequence, the set of strongly robust SPE payoff profiles is essentially the same whether we assume common priors or not. Similarly, the folk theorem (Theorem 3) holds without the common prior assumption because our proof (especially Lemma 16 in Appendix A) relies only on iterative stage dominance.

7.2 Persistent Shocks

We now extend the class of perturbations against which we test robustness to allow for payoff shocks that are correlated across periods. We show that our notion of robustness is unchanged if asymmetric information is short lived as long as the players are in “normal”

states of the world (where “normal” will be made precise shortly). We note that this is no longer true if players have persistent informational asymmetry.²⁶

The class of correlated perturbations we consider is described as follows. In addition to payoff-relevant states ω_t and information sets $Q_{i,t}$, we allow players to observe public signals $z_t \in Z_t$, where Z_t is countable. We refer to z_t as regimes and denote by $Z_t^* \subseteq Z_t$ a set of “normal” regimes, which will be defined shortly. Let P be the probability distribution over $\prod_{t \geq 1} (\Omega_t \times Z_t)$. Distribution P may exhibit significant correlation between states $(\omega_t)_{t \in \mathbb{N}}$.

We say that P is ε -persistent along normal regimes if

$$\left| \frac{P(\omega_t | z_1, \omega_1, \dots, z_{t-1}, \omega_{t-1}, z_t)}{P(\omega_t | z_1, \dots, z_t)} - 1 \right| \leq \varepsilon$$

for every $t \geq 1$ and $(z_1, \omega_1, \dots, z_t, \omega_t) \in \prod_{\tau=1}^t (Z_\tau^* \times \Omega_\tau)$. In words, this means that while players have always been in a normal regime, then conditional on past regimes, private information over past states does not affect beliefs over the current state much. Note that once a abnormal regime is reached, past private information may become very relevant.

A sequence $\hat{\mathbf{U}} = (N, (\Omega_t, (A_i, u_{it}, Q_{it})_{i \in N}, Z_t)_{t \in \mathbb{N}}, P)$ of incomplete-information games that embed G with intertemporal correlation P is a *sequence of correlated (ε, d) -elaborations of G* if, P is ε -persistent along normal regimes, and conditional on each sequence $(z_1, \dots, z_t) \in \prod_{\tau=1}^t Z_\tau^*$ of past normal regimes, the stage game is close to G with high probability, i.e.,

$$P(\{\omega_t \in \Omega_t | \forall i \in N, \forall \omega'_t \in Q_{it}(\omega_t), |u_{it}(\cdot, \omega'_t) - g_i| \leq d\} | z_1, \dots, z_t) \geq 1 - \varepsilon,$$

and a regime in the next period is normal with high probability, i.e., $P(z_{t+1} \in Z_{t+1}^* | z_1, \dots, z_t) \geq 1 - \varepsilon$. Note that this need only hold conditional on past regimes being normal. In particular abnormal regimes can be arbitrarily persistent.

²⁶Perturbations with persistent asymmetric information are closely related to the type of perturbations studied in the reputation literature. See Mailath and Samuelson (2006) for a review of the reputation literature. See also Angeletos, Hellwig and Pavan (2007) for an analysis of the learning patterns that arise in a dynamic game of regime change where fundamentals are correlated across time.

An example. The class of correlated (ε, d) -elaborations includes the following perturbed prisoners' dilemma. In each period players have private information over whether or not the game will stop next period. More formally, in each period t a state $\omega_t \in \{1, \dots, L, L+1\}$ is drawn, players observe a public signal $z_t = \omega_{t-1}$ and a private signal $x_{i,t}$ where $x_{1,t} = \omega_t$ and $x_{2,t} = \omega_t - \xi_t$, with ξ_t an even coin flip over $\{0, 1\}$. Conditional on any $\omega_{t-1} \in \{1, \dots, L-1\}$, ω_t belongs to $\{1, \dots, L-1\}$ with high probability. If $\omega_{t-1} = L$, then $\omega_t = L+1$. Finally, state $L+1$ is absorbing. In states $\omega_t \in \{1, \dots, L\}$ payoffs are the payoffs of the original prisoners' dilemma. In state $L+1$ all payoffs are identically 0. State $L+1$ corresponds to the de facto end of the game. In state L , player 1 knows that the game will end next period, while player 2 may be uncertain.

Proposition 7 shows that robustness against such correlated (ε, d) -elaborations is equivalent to robustness against independent (ε, d) -elaborations. We say that a public history h_{t-1} is normal if and only if all past regimes are normal (i.e. for all $s \leq t-1$, $z_s \in Z_s^*$).

Definition 8 (dynamic robustness with persistent shocks). For $d \geq 0$, a PPE s^* of Γ_G is *d-robust to persistent incomplete information with public regimes* if, for every $\eta > 0$ and $M < \infty$, there exists $\varepsilon > 0$ such that, for every sequence $\hat{\mathbf{U}}$ of correlated (ε, d) -elaborations of G with $|\hat{\mathbf{U}}| < M$, the induced dynamic incomplete-information game has an equilibrium that puts probability at least $1 - \eta$ on $s^*(h_{t-1})$ at every normal public history $h_{t-1} \in H_{t-1}$.

A PPE s^* of Γ_G is *strongly robust to persistent incomplete information with public regimes* if it is *d-robust to persistent incomplete information with public regimes* for some $d > 0$.

Conditional on each public history, players may have different priors over current payoff shocks because they have observed different past signals. However, as long as past public regimes are normal, their beliefs over the current state will be close in the sense of Section 7.1. Therefore, Proposition 6 implies the following.

Proposition 7 (equivalence of perturbation classes). *If a PPE is strongly robust to independent incomplete information, then it is also strongly robust to persistent incomplete information with public regimes. Hence, strong robustness in the sense of Definition 8 is equivalent to that of Definition 3.*

This shows that correlations across shocks do not change our notion of robustness as long as asymmetric information is short-lived while players are in a “normal” regime. Note that this result does not hold anymore if asymmetric information is long lived: Appendix A.18 describes a game with a unique strongly robust equilibrium which is not robust anymore once persistent asymmetric information is allowed.

8 Conclusion

This paper provides a framework to study the robustness of repeated games strategies without committing to a specific incomplete information perturbation, and highlights the applied implications of robustness considerations.

Our main technical contribution is a one-shot robustness principle which reduces the analysis of robust equilibria in dynamic games to the analysis of robust equilibria in appropriate families of static games. This implies a factorization result for strongly robust PPE payoff profiles. We show the practical value of these characterizations by means of two examples.

First, we compute explicitly the set of strongly robust SPE payoff profiles in the repeated prisoners’ dilemma. We show that cooperation can be sustained by a strongly robust SPE if and only if both $(Cooperate, Cooperate)$ and $(Defect, Cooperate)$ are enforceable under complete information. In the spirit of Chassang and Padro i Miquel (2008) we also show that concerns of robustness can significantly affect comparative statics. Finally, our analysis implies that selective punishment strategies are more effective than grim-trigger strategies in sustaining cooperation in strongly robust SPEs. This occurs because grim trigger strategies minimize predatory incentives but increase preemptive incentives. In contrast, selective punishment strategies minimize both predatory and preemptive incentives.

Second, we prove a folk theorem in strongly robust PPEs for repeated games with imperfect public monitoring under the strong full rank condition. The identifiability conditions we use are stronger than those of FLM because robustness requires us to control all continuation payoff profiles upon joint deviations, rather than just upon unilateral deviations.

Our approach is necessarily dependent on the class of perturbations against which we test for robustness. While we think of the class of perturbations we consider as a natural and informative benchmark, one may reasonably worry whether studying other classes of perturbations would lead to very different results. In this respect it is informative to note that our results are unchanged if players have almost common priors or when payoff shocks are correlated across periods but private information is short lived.

A Proofs

A.1 Proof of Proposition 1

Since each Γ_L is an infinite-horizon incomplete-information game induced by a sequence of $(1/L, 0)$ -elaborations of PD (see Section 3.1 for definitions), the robustness of grim-trigger strategies is a straightforward application of Lemmas 3 and 4 and Theorem 1. See Section 5.1. To see this more explicitly, suppose that $1/(1-\delta) > b+c$. Consider a profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ of public strategies such that $\sigma_i^*(h_{t-1}, x_{i,t}) = C$ if $h_{t-1} = (CC, \dots, CC)$ and $x_{i,t} \leq L-1$, and $\sigma_i^*(h_{t-1}, x_{i,t}) = D$ otherwise. Note that, for every $x_{i,t} \leq L-1$, conditional on observing $x_{i,t}$, player i believes with probability at least $1/2$ that the opponent j observes $x_{j,t} \leq L-1$. Since CC is strictly risk-dominant in

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

σ^* is a PPE for sufficiently large L . Thus grim-trigger strategies are robust to $\{\Gamma_L \mid L \geq 1\}$.

Suppose that $1/(1-\delta) < b+c$. If grim-trigger strategies are robust to $\{\Gamma_L \mid L \geq 1\}$, then, for L large enough, Γ_L has a PPE σ^* that is close to grim-trigger strategies at every history. Let $v_i(\sigma^*|a)$ denote player i 's continuation payoff under σ^* after action profile a in the first period. Since σ^* is arbitrarily close to grim-trigger strategies, for $B > 0$ fixed and L large, $v_i(\sigma^*|CC)$ is arbitrarily close to 1, and $v_i(\sigma^*|CD)$, $v_i(\sigma^*|DC)$ and $v_i(\sigma^*|DD)$ arbitrarily close to 0. Since $1/(1-\delta) < b+c$, we can insure that, for L large enough,

$$1 + \delta v_i(\sigma^*|CC) - c + \delta v_i(\sigma^*|DC) < b + \delta v_i(\sigma^*|CD) + \delta v_i(\sigma^*|DD). \quad (3)$$

The rest of the proof shows by induction that both players play D in the first period under σ^* , which contradicts the robustness of grim-trigger strategies. If player 1 observes signal L , then, since B is sufficiently large, playing D is dominant for him. If player 2 observes signal L , she puts probability 1 on player 1 having observed signal L and playing D , and hence her best reply is to play D . Assume that, if player 1 observes signal k , then he plays D . If player 2 observes signal $k - 1$, then she puts probability at least $1/2$ on player 1 having observed k . By the induction hypothesis, this implies that she puts probability at least $1/2$ on player 1 playing D . Thus, by (3), her best reply is to play D . Similarly, if player 1 observes signal $k - 1$, he puts probability $1/2$ on player 2 having observed $k - 1$ and playing D . By (3), his best reply is to play D .

A.2 Proof of Proposition 2

Similarly to the first half of Proposition 1, Proposition 2 follows from Lemma 3 and Theorem 1. Indeed it is straight forward to check that at any history, the prescribed equilibrium action profile is iteratively dominant in the appropriate one-shot game augmented with continuation values.

A.3 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 a Nash equilibrium arbitrarily close to a^* . By the closedness of the set of Nash equilibria, a^* is also a Nash equilibrium of G' . Therefore, a^* is a $2d$ -strict equilibrium of G .

A.4 Proof of Lemma 2

The proof is by contradiction, and follows the structure of KM (Proposition 3.2). It uses Lemmas 9 and 10, which are of independent interest and given below.

Definition 9 (canonical 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 *canonical 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.²⁷

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 9 (canonical normalization with respect to a Bayesian-Nash equilibrium). *Let \tilde{U} be the canonical normalization of U with respect to α^* .*

(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 $\tilde{d} = d + \varepsilon^{1/2}(|g| + M)$.*

(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 canonical 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

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

²⁷The denominator is nonzero \tilde{P} -almost surely.

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 canonical normalization is that, given a set of players and an action space, they form a finite-dimensional class of games.

Lemma 10 (locally unique equilibrium). *If a^* is a strict equilibrium of G and G has no other correlated 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 canonical normalization of a $(0, d)$ -elaboration of G is also a $(0, d)$ -elaboration of G by Lemma 9, without loss of generality, we can assume that U_d takes a canonical 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 $a'_i \in A_i$,

$$\sum_{a_{-i} \in A_{-i}} g_i(a_i, a_{-i}) \mu_0(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} g_i(a'_i, a_{-i}) \mu_0(a_i, a_{-i}). \quad (4)$$

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 (4) 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 Lemma 2. By Lemma 10, 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 canonical normalization of U_ε with respect to α_ε . By Lemma 9, 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 , truth-telling 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.5 Proof of Lemma 3

The proof of Lemma 3 is almost the same as that of Lemma 2. The only difference is to replace Lemma 10 by the following.

Lemma 11 (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.6 Proof of Lemma 4

We define the following notion.

Definition 10 ((\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 $d > 0$. Lemma 4 follows from the following lemma.

Lemma 12 (strong robustness of (\mathbf{p}, d) -dominant equilibria). *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 $\lambda(a_{-i}^*) \geq p_i$,

$$\sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i^*, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} \lambda(a_{-i}) g_i'(a_i, a_{-i}) \quad (5)$$

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 (5) and KM (Lemma 5.2), U has 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.7 Proof of Lemma 5

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^* . By the closedness of the set of PPEs, s^* is also a PPE of $\Gamma_{\mathbf{U}}$, hence $s^*(h^0)$ is a Nash equilibrium of $G'(w_{s^*, h^0})$. Therefore, $s^*(h^0)$ is a $2d$ -strict equilibrium of $G(w_{s^*, h^0})$.

A.8 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 $u_i(a, \omega) + \delta \mathbb{E}[w_i(y)|a]$ 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 5, $s^*(h)$ is a $2d$ -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 d -robust in $G(w_{s^*, h^0})$. That is, for every $\eta > 0$, there exists $\varepsilon > 0$ such that every (ε, d) -elaboration of $G(w_{s^*, h^0})$ has a Bayesian-Nash equilibrium that puts probability at least $1 - \eta$ on $s^*(h^0)$.

Fix any $\eta > 0$. 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 $|\mathbf{U}| \leq 2|g|/(1 - \delta) + d$, $\Gamma_{\mathbf{U}}$ has a PPE that puts

probability at least $1 - \eta$ on $s^*(h)$ for every $h \in H$. Fix such ε . Pick any (ε, d) -elaboration $U = G(w_{s^*,h^0})$ of $G(w_{s^*,h^0})$. Without loss of generality, it is sufficient to consider elaborations such that $|U| \leq |g|/(1-\delta) + d$. Consider the “one-shot” sequence $\mathbf{U} = \{U_t\}$ such that $U_t = G$ for all $t \neq t^0$ and $U_{t^0} = U - \delta w_{s^*,h^0}$.²⁸ We have that $|\mathbf{U}| \leq 2|g|/(1-\delta) + d$. 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 $U(w_{\sigma^*,h^0})$.

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 $|w_{\sigma,h^1} - w_{s^*,h^1}| \leq d$, and
- if a^* is a $2(1 - \delta)d$ -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^* .

We now show that $\sigma^*(h) = s^*(h)$ for every $t > t_0$ and $h \in H_{t-1}$.²⁹ By the choice of η , we have $|w_{\sigma^*,h} - w_{s^*,h}| \leq d$. Then, since $s^*(h)$ is $2d$ -strict in $G(w_{s^*,h})$, $s^*(h)$ is $2(1 - \delta)d$ -strict in $G(w_{\sigma^*,h})$. Since $G(w_{\sigma^*,h})$ has no other Nash equilibria in the η -neighborhood of $s^*(h)$, $\sigma^*(h) = s^*(h)$.

Therefore, 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 $0 < d' < (1 - \delta)d$. 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 $|w_{\sigma,h} - w_{s^*,h}| \leq d'/(1 - \delta)$. Pick $d'' > 0$ such that $d'/(1 - \delta) + \delta d'' < d$. 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

²⁸ $U - \delta w_{s^*,h^0}$ denotes the incomplete information game with payoffs $u(\cdot, \omega) - \delta w_{s^*,h^0}$.

²⁹Since U_t is a complete-information game G for $t \neq t_0$, we suppress ω_t from the notation $\sigma^*(h, \omega_t)$.

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 $(\varepsilon, d'/(1 - \delta))$ -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 $|w_{\sigma^T, h} - w_{s^*, h}| \leq d'/(1 - \delta)$. Therefore, $U_t(w_{\sigma^T, h})$ is an $(\varepsilon, d'/(1 - \delta))$ -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.9 Proof of Lemma 6

(i) holds by the definition of B^d . (ii) and (iii) follow from Tarski’s fixed point theorem.

A.10 Proof of Theorem 2

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 6, $v \in V^* \subseteq V^d$. 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 $\bigcup_{d>0} V^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (B^d)^k(F)$, which corresponds to APS's algorithm result. To this end, we define $\bar{B}^d(F)$ by the closure of $B^d(F)$. Denote $f^\infty(F) = \bigcap_{k=0}^{\infty} f^k(F)$ for $f = B^d$ or \bar{B}^d . By the monotonicity of B^d and \bar{B}^d (Lemma 6), we have $V^d \subseteq (B^d)^\infty(F) \subseteq (\bar{B}^d)^\infty(F)$ for every $d > 0$.

To prove the opposite direction of set inclusion, we show that, for each $d > 0$, $(\bar{B}^d)^\infty(F)$ is self-generating with respect to $B^{d/2}$, which implies that $(\bar{B}^d)^\infty(F) \subseteq V^{d/2}$ by Lemma 6. Pick any $v \in (\bar{B}^d)^\infty(F)$. For each $k \geq 1$, since we have $v \in (\bar{B}^d)^\infty(F) \subseteq (\bar{B}^d)^k(F)$, there exist $a^k \in A$ and $w^k: Y \rightarrow (\bar{B}^d)^{k-1}(F)$ such that w^k enforces (a^k, v) d -robustly. Since A and Y are finite and $(\bar{B}^d)^k(F)$ 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 $w^l(y) \in (\bar{B}^d)^{l-1}(F) \subseteq (\bar{B}^d)^{k-1}(F)$ for every $l \geq k$ and $(\bar{B}^d)^{k-1}(F)$ is compact, by taking $l \rightarrow \infty$, we have $w^*(y) \in (\bar{B}^d)^{k-1}(F)$. Since this holds for every $k \geq 1$, $w^*(y) \in (\bar{B}^d)^\infty(F)$. Thus $v \in B^{d/2}((\bar{B}^d)^\infty(F))$, and $(\bar{B}^d)^\infty(F)$ is self-generating with respect to $B^{d/2}$.

A.11 Stahl's Characterization

Here we summarize the results of Stahl (1991), which characterize V^{SPE} , the set of SPE 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 : $V_0 = \frac{1}{1-\delta} \text{co}\{(0, 0), (1, 1), (0, p), (p, 0)\}$;

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

V_T the set of values defined by $V_T = \frac{1}{1-\delta} \text{co}\{(0, 0), (0, b-c), (b-c, 0)\}$;

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

Lemma 13 (Stahl (1991)). V^{SPE} is characterized as follows.

- (i) If $\delta \geq \max\{(b-1)/b, c/(c+1)\}$, then $V^{\text{SPE}} = V_0$.
- (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 $V^{\text{SPE}} = V_Q$.
- (iii) If $c < b-1$ and $\delta \in [c/b, (b-1)/b)$, then $V^{\text{SPE}} = V_T$.
- (iv) If $c > h$ and $\delta \in [(b-1)/b, \delta^*)$, then $V^{\text{SPE}} = V_D$.
- (v) If $\delta < \min\{c/b, (b-1)/b\}$, then $V^{\text{SPE}} = \{(0, 0)\}$.

A.12 Proof of Lemma 7

The *SPE Pareto frontier* is the set of $v \in V^{\text{SPE}}$ such that there is no $v' \in V^{\text{SPE}}$ that Pareto-dominates v . We say that an SPE is *Pareto-efficient* if it induces a payoff profile on the SPE 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 14 (SPE Pareto frontier of games in $\mathcal{G}_{DC/CC}$). Let $\text{PD} \in \mathcal{G}_{DC/CC}$.

- (i) The SPE Pareto frontier is self-generating with respect to $\text{co} B$.
- (ii) No Pareto-efficient SPE prescribes outcome DD on the equilibrium path.
- (iii) The SPE Pareto frontier can be sustained by SPEs 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 SPE payoff profiles of Γ_{PD} takes the form $V^{\text{SPE}} = \text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta}), (0, \phi), (\phi, 0)\}$, where $\phi \geq \frac{1}{1-\delta}$. We begin with point (i). Pick a Pareto-efficient SPE s^* . Note that continuation payoff profiles of s^* on the equilibrium path are always on the SPE Pareto frontier (otherwise, replacing the continuation strategies by a Pareto-dominating SPE would improve on s^*). In what follows, we modify s^* so that continuation values are on the SPE Pareto frontier even off the equilibrium path. This is possible because points $(0, \phi)$ and $(\phi, 0)$ belong to the SPE 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 an SPE implies that \hat{s}^* is also an SPE. This shows that the SPE Pareto frontier is self-generating with respect to $\text{co } B$.

Let us turn to point (ii). Consider a Pareto-efficient SPE 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 an SPE and Pareto-dominates s^* . Hence, action DD is never used on the equilibrium path.³⁰

We now proceed with point (iii). From point (i), we know that the SPE Pareto frontier is self-generating with respect to $\text{co } B$. Since we have public randomization, this implies that the SPE Pareto frontier can be generated by SPEs whose continuation payoff profiles are always extreme points of the convex hull of the SPE Pareto frontier. This is the bang-bang property of APS. There are three such points, $(0, \phi)$, $(\phi, 0)$ and $(\frac{1}{1-\delta}, \frac{1}{1-\delta})$. Because $(\frac{1}{1-\delta}, \frac{1}{1-\delta})$ is not the discounted sum of payoffs upon action profiles other than CC , this implies that, in any SPE that sustains values $(\frac{1}{1-\delta}, \frac{1}{1-\delta})$, 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 SPEs 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 SPEs never prescribe DD on or off the equilibrium path. \square

Proof of Lemma 7. Let us consider $PD \in \text{int } \mathcal{G}_{DC/CC}$. Since, for every PD' sufficiently close to PD , CC is enforced by an SPE of $\Gamma_{PD'}$ with continuation payoff profile $(1, 1)$ after CC , we have $1 > (1 - \delta)b$.

³⁰If players only play DD following h , one can simply replace the entire continuation equilibrium by some SPE that gives the players strictly positive value.

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

which is strategically equivalent to PD_d . Since $\text{PD} \in \text{int } \mathcal{G}_{DC/CC}$, there exists $\bar{d} \in (0, 1)$ such that, for $d \in (0, \bar{d})$, we have that $\text{PD}'_d \in \mathcal{G}_{DC/CC}$. This means that the set of SPE payoff profiles of $\Gamma_{\text{PD}'_d}$ is a quadrangle $\text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta}), (0, \phi'), (\phi', 0)\}$, where $\phi' \geq \frac{1}{1-\delta}$. Note that, since DC is enforceable under complete information in $\Gamma_{\text{PD}'_d}$, we have $\frac{-c-d}{1-d} + \delta\phi' \geq 0$. By Lemma 14, we know that the SPE Pareto frontier of $\Gamma_{\text{PD}'_d}$ is sustained by a class of SPEs such that continuation payoffs are always on the SPE 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 SPEs of Γ_{PD_d} and generate its SPE Pareto frontier. The SPE Pareto frontier of Γ_{PD_d} is obtained by multiplying equilibrium values of $\Gamma_{\text{PD}'_d}$ by $1 - d$ and adding $d/(1 - \delta)$. We denote by ℓ_d this frontier: ℓ_d is the piecewise line that connects $(\frac{d}{1-\delta}, \phi)$, $(\frac{1}{1-\delta}, \frac{1}{1-\delta})$ and $(\phi, \frac{d}{1-\delta})$, where $\phi = (1 - d)\phi' + d/(1 - \delta) \geq c/\delta + d/[\delta(1 - \delta)]$. Note that, in Γ_{PD_d} , continuation payoffs of these SPEs are at least $d/(1 - \delta)$ at all histories.

Let us now show that strategy profiles in \mathcal{E} are also SPEs 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 SPEs 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 SPE 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 d -dominant in $\text{PD}(w_{\hat{s}^*, h})$ 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 d -dominant in $\text{PD}(w_{\hat{s}^*, h})$. If $s^*(h) = DC$, a symmetric change makes DC iteratively d -dominant in a game $\text{PD}(w_{\hat{s}^*, h})$, 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 $w(DD) = (0, 0)$, $w(DC) = (\frac{d}{1-\delta}, \phi)$ and $w(CD) = (\phi, \frac{d}{1-\delta})$. Since $s^* \in \mathcal{E}$, the on-path continuation-payoff profile is $(\frac{1}{1-\delta}, \frac{1}{1-\delta})$. Since $\frac{1}{1-\delta} > b + \frac{\delta}{1-\delta}d + d$ and $-c + \delta\phi \geq d$, CC is iteratively d -dominant in $\text{PD}(w_{\hat{s}^*, h})$.

It results from this that every payoff profile in $\text{co}(\{(0, 0)\} \cup \ell_d) = \text{co}\{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta}), (\frac{d}{1-\delta}, \phi), (\phi, \frac{d}{1-\delta})\}$ is sustained by some SPE that prescribes the iteratively d -dominant equilibrium of the corresponding augmented game at every history. By taking d to 0, we obtain that, for every $v \in \{(0, 0), (\frac{1}{1-\delta}, \frac{1}{1-\delta})\} \cup \text{int } V^{\text{SPE}}$, there exist $d > 0$ and an SPE with payoff profile v that prescribes the iteratively 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.13 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 15 (limit of strongly robust PPE payoff profiles). *We have the following.*

- (i) $NV^{\text{rob}}(\delta) \subseteq Q$ 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 $K \subseteq NV^{\text{rob}}(\delta)$ 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 16 (characterization of $k(\lambda)$). *Suppose that (Y, π) has strong full rank.*

(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) $Q = NV^*$.

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:

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

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 Lemma 3, 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) has strong full rank, $\Pi_i(a^0)$ 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 Lemma 3, 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). □

Theorem 3 follows from Lemmas 15 and 16.

A.14 Proof of Proposition 5

Suppose that $\gamma := \sup\{v_1 - v_2 \mid (v_1, v_2) \in NV^{\text{rob}}(\delta)\} > 1/2$ for some $\delta < 1$. For any $\varepsilon \in (0, \gamma)$, there exists $(v_1, v_2) \in V^{\text{rob}}(\delta)$ such that $(1 - \delta)(v_1 - v_2) > \gamma - \varepsilon$ and action profile RL is taken at the initial history.³¹ By Theorem 1, there exist $w(y_L), w(y_R), w(y_M) \in V^{\text{rob}}(\delta)$ that enforce $(RL, (v_1, v_2))$ robustly, i.e., such that RL is strongly robust in

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

where

$$v_1 = 1 + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)),$$

$$v_2 = \frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)).$$

³¹If this is not the case, delete several initial periods. This always increases $v_1 - v_2$ since $g_1(a) \leq g_2(a)$ for all $a \neq RL$.

Let $\gamma(y) = (1 - \delta)(w_1(y) - w_2(y))$ for each $y \in Y$. By the definition of γ , we have $\gamma(y) \leq \gamma$ for every $y \in Y$.

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

$$\frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)) > \delta w_2(y_R), \quad (6)$$

$$1 + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)) > 3 + \delta w_1(y_L). \quad (7)$$

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

$$3 + \delta w_1(y_L) + \delta w_1(y_M) \leq 1 + \frac{\delta}{3}(w_1(y_L) + w_1(y_R) + w_1(y_M)) + \delta w_1(y_R), \quad (8)$$

or

$$1 + \delta w_2(y_M) + \delta w_2(y_R) \leq 3 + \delta w_2(y_L) + \frac{\delta}{3}(w_2(y_L) + w_2(y_R) + w_2(y_M)). \quad (9)$$

If (8) holds, then (6) and (8) yield $3(1 - \delta)/\delta < -\gamma(y_L) + 2\gamma(y_R) - \gamma(y_M)$. Hence,

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

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

Similarly, if (9) holds, then (7) and (9) yield $3(1 - \delta)/\delta < -2\gamma(y_L) + \gamma(y_R) + \gamma(y_M)$.

Hence,

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

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

A.15 Proof of Lemma 8

Fix $\eta > 0$ and $M > 0$. Since a^* is d -robust, there exists $\varepsilon_0 > 0$ such that every (ε, d) -elaboration of G has a Bayesian-Nash equilibrium that puts probability at least $1 - \eta$ on a^* . Let $\varepsilon = \min(\varepsilon_0, (d - d')/M) > 0$. Fix an (ε, d') -elaboration $U = (N, \Omega, P, (A_i, u_i, Q_i)_{i \in N})$ of G with $|u| < M$ and priors $(P_i)_{i \in N}$ with $m(P, P_i) \leq \varepsilon$. Consider an incomplete-information

game $U' = (N, \Omega, P, (A_i, u'_i, Q_i)_{i \in N})$ with common prior P , where each u'_i is defined by $u'_i(\cdot, \omega) = (P_i(\omega)/P(\omega))u_i(\cdot, \omega)$ if $P(\omega) > 0$, and $u'_i(\cdot, \omega) = 0$ otherwise. Since $|u'_i(\cdot, \omega) - u_i(\cdot, \omega)| \leq m(P, P_i)|u| < d - d'$ P -almost surely, U' is an (ε, d) -elaboration of G , hence has a Bayesian-Nash equilibrium α^* that puts probability at least $1 - \eta$ on a^* . By the construction of U' , α^* is also a Bayesian-Nash equilibrium of the non-common prior game $(U, (P_i)_{i \in N})$.

A.16 Proof of Proposition 6

The proof is essentially identical to that of Theorem 1.

A.17 Proof of Proposition 7

We show that strong robustness in the sense of Definition 8 is also characterized by the one-shot robustness principle, i.e., if there exists $d > 0$ such that $s^*(h)$ is d -robust in $G(w_{s^*, h})$, then s^* is strongly robust in Definition 8. The proof is similar to that of Theorem 1. For each sequence \hat{U} of correlated (ε, d) -elaborations of G , we construct a PPE $\hat{\sigma}^T$ of truncated game $\Gamma_{\hat{U}^T}$ close to s^* along normal regimes and take $T \rightarrow \infty$. For each sequence h_{t-1} of public signals, players' private information is summarized by the current public regime z_t . Thus, if $z_t \in Z_t^*$, then the continuation game is close to $G(w_{s^*, h_{t-1}})$, thus has a Bayesian-Nash equilibrium $\hat{\sigma}^T(h_{t-1}, z_t, \cdot)$ that puts high probability on $s^*(h_{t-1})$. If $z_t \notin Z_t^*$, then players' actions outside normal regimes are determined arbitrarily by Kakutani's fixed point theorem.

A.18 Privately Observed Regimes

While our framework extends easily to the case of publicly observed regime changes, this is not true anymore when there is durable asymmetric information over past payoff shocks. This is the kind of environment that the literature on reputation focuses on, and it is easy to come up with examples showing that there is no hope to obtain an analogue to the one-shot deviation principle.

To illustrate the type of issues that can occur, consider the following stage game G

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

G is a coordination game with unique strongly robust equilibrium TL . In addition, we

show below that, for any discount factor δ , repeatedly playing TL is the unique strongly robust SPE of the repeated game Γ_G with perfect monitoring. 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 $\varepsilon > 0$, 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 SPE is not robust with respect to persistent payoff shocks.³²

Let us now 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

Here we extend 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

³²Note that one can construct similar examples if payoff shocks are not permanent (for instance, if the column player of the crazy type can become normal with positive probability every period).

$\theta^t = (\theta_1, \dots, \theta_t) \in [0, 1]^t$, $\tilde{h}_{t-1} = (h_{t-1}, \theta^t) \in \tilde{H}_{t-1} = H_{t-1} \times [0, 1]^t$, 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 $R_1 \otimes \dots \otimes R_t$ -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 .

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 $R_1 \otimes \dots \otimes R_t$ -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 11 (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 8 (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 6, $\text{co } B^d$ is monotonic and admits the largest fixed point among all subsets of F . We denote it by \tilde{V}^d . Then we have the following characterization similar to Theorem 2.

Theorem 4 (characterization of \tilde{V}^{rob}). $\tilde{V}^{\text{rob}} = \bigcup_{d>0} \tilde{V}^d = \bigcup_{d>0} \bigcap_{k=0}^{\infty} (\text{co } B^d)^k(F)$.

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 8, 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 $v \in V^* \subseteq \tilde{V}^d$. 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 8, \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

Here, 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 . 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 $u_{ih}(a_h, \omega_h)$. 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 12 (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 3. 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 9 (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 Theorem 1.

C.2 Failure of the One-Shot Robustness Principle under Weak Robustness

The notion of strong robustness we develop in Section 3 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 an SPE 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 17 (SPEs of Γ_T).

- (i) *In every SPE 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 SPE 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 10 (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 17. 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 17, the unique SPE 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 10, 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 10 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 SPE, 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. (1988). "On the Theory of Infinitely Repeated Games with Discounting," *Econometrica*, 56, 383–396.
- [2] Abreu, D., D. Pearce and E. Stacchetti (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58, 1041–1063.
- [3] Aghion, P., D. Fudenberg and R. Holden (2008). "Subgame Perfect Implementation with Almost Perfect Information," Unpublished.
- [4] 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.
- [5] Argenziano, R. (2008). "Differentiated Networks: Equilibrium and Efficiency," *RAND Journal of Economics*, 39, 747–769.
- [6] Bagwell, K. and G. Staiger. (1990). "A Theory of Managed Trade A Theory of Managed Trade," *American Economic Review*, 80, 779–795.

- [7] Baker, G., R. Gibbons and K. Murphy. (1994). “Subjective Performance Measures in Optimal Incentive Contracts,” *Quarterly Journal of Economics*, 109, 1125–1156
- [8] Baker, G., R. Gibbons and K. Murphy. (2002). “Relational Contracts and the Theory of the Firm,” *Quarterly Journal of Economics*, 117, 39–84.
- [9] Bergemann, D. and S. Morris (2005). “Robust Mechanism Design,” *Econometrica*, 73, 1771–1813.
- [10] Bhaskar, V., G. J. Mailath and S. Morris (2008). “Purification in the Infinitely-Repeated Prisoners’ Dilemma,” *Review of Economic Dynamics*, 11, 515–528.
- [11] Bull, C. (1987). “The Existence of Self-Enforcing Implicit Contracts The Existence of Self-Enforcing Implicit Contracts,” *Quarterly Journal of Economics*, 102, 147–159.
- [12] Carlsson, H. and E. van Damme (1993). “Global Games and Equilibrium Selection,” *Econometrica*, 61, 989–1018.
- [13] Chamley, C. (1999), “Coordinating Regime Switches,” *Quarterly Journal of Economics*, 114, 869–905.
- [14] Chassang, S. (2009). “Fear of Miscoordination and the Robustness of Cooperation in Dynamic Global Games with Exit,” forthcoming, *Econometrica*.
- [15] Chassang, S. and G. Padro i Miquel (2008). “Conflict and Deterrence under Strategic Risk,” Unpublished.
- [16] Chen, Y.-C. and S. Xiong (2008). “Strategic Approximation in Incomplete-Information Games,” Unpublished.
- [17] Dekel, E. and D. Fudenberg (1990). “Rational Behavior with Payoff Uncertainty,” *Journal of Economic Theory*, 52, 243–267.
- [18] Dekel, E., D. Fudenberg and S. Morris (2006). “Topologies on Types,” *Theoretical Economics*, 1, 275–309.
- [19] Di Tillio, A. and E. Faingold (2007). “Uniform Topology on Types and Strategic Convergence,” Unpublished.
- [20] Ely, J. C. and J. Välimäki (2002). “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, 102, 84–105.

- [21] Frankel, D., S. Morris and A. Pauzner (2003). “Equilibrium selection in global games with strategic complementarities,” *Journal of Economic Theory*, 108, 1–44.
- [22] Fudenberg, D., D. M. Kreps and D. K. Levine (1988). “On the Robustness of Equilibrium Refinements,” *Journal of Economic Theory*, 44, 354–380.
- [23] Fudenberg, D. and D. K. Levine (1994). “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, 62, 103–135.
- [24] Fudenberg, D., D. K. Levine and E. Maskin (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 62, 997–1039.
- [25] Fudenberg, D. and E. Maskin (1986). “The Folk Theorem in Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, 54, 533–554.
- [26] Fudenberg, D. and J. Tirole (1991). *Game Theory*. Cambridge: MIT Press.
- [27] Giannitsarou, C. and F. Toxvaerd (2007). “Recursive Global Games,” Unpublished.
- [28] 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.
- [29] 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.
- [30] Hörner, J. and W. Olszewski (2008). “How Robust is the Folk Theorem?,” forthcoming in the *Quarterly Journal of Economics*.
- [31] Izmalkov, S. and M. Yildiz (2008). “Investor Sentiments,” forthcoming in *American Economic Journal: Microeconomics*.
- [32] Kajii, A. and S. Morris (1997). “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283–1309.
- [33] Mailath, G. J. and S. Morris (2002). “Repeated Games with Almost-Public Monitoring,” *Journal of Economic Theory*, 102, 189–228.
- [34] Mailath, G. J. and S. Morris (2006). “Coordination failure in repeated games with almost-public monitoring,” *Theoretical Economics*, 1, 311–340.

- [35] Mailath, G. and W. Olszewski (2008). “Folk Theorems with Bounded Recall and (Almost) Perfect Monitoring,” Unpublished.
- [36] Mailath, G. J. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford: Oxford University Press.
- [37] Miller, D (2007). “The Dynamic Cost of Ex Post Incentive Compatibility in Repeated Games of Private Information,” Unpublished.
- [38] Monderer, D. and D. Samet (1989). “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior*, 1, 170–190.
- [39] Morris, S. and H. S. Shin (1998). “Unique Equilibrium in a Model of Self-Fulfilling Attacks,” *American Economic Review*, 88, 587–597.
- [40] 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.
- [41] Morris, S. and T. Ui (2005). “Generalized Potentials and Robust Sets of Equilibria,” *Journal of Economic Theory*, 124, 45–78.
- [42] Oury, M. and O. Tercieux (2008). “Continuous Implementation,” Unpublished.
- [43] Oyama, D. and O. Tercieux (2009). “Robust Equilibria under Non-Common Priors,” Unpublished.
- [44] Rotemberg, J. and G. Saloner (1986). “A Supergame-Theoretic Model of Price Wars during Booms,” *American Economic Review*, 76, 390–407.
- [45] Rubinstein, A. (1989). “The Electronic Mail Game: A Game with Almost Common Knowledge,” *American Economic Review*, 79, 389–391.
- [46] 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.
- [47] Ui, T. (2001). “Robust Equilibria of Potential Games,” *Econometrica*, 69, 1373–1380.
- [48] Weinstein, J. and M. Yildiz (2007). “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.