Instability of Equilibria with Private Monitoring

Heller, Yuval

University of Oxford

20 December 2015

Instability of Belief-free Equilibria

Yuval Heller*

Department of Economics and Queen's College, University of Oxford

December 20, 2015

Abstract

Various papers have presented folk theorem results for repeated games with private monitoring that rely on belief-free equilibria. I show that these equilibria are not robust against small perturbations in the behavior of potential opponents. Specifically, I show that (1) essentially none of the belief-free equilibria is evolutionarily stable, (2) most of them are not neutrally stable, and (3) in games with two actions that admit a strict equilibrium, such as the Prisoner's Dilemma, the belief-free equilibria fail to satisfy even a very mild stability refinement.

JEL Classification: C73, D82. **Keywords**: Belief-free equilibrium, evolutionary stability, private monitoring, repeated Prisoner's Dilemma, communication.

1 Introduction

The theory of repeated games provides a formal framework to explore the possibility of cooperation in long-term relationships, such as collusion between firms. The various folk theorem results (e.g., Fudenberg and Maskin, 1986; Abreu, Pearce, and Stacchetti, 1990) have established that efficiency can be achieved under fairly general conditions when players observe commonly shared information about past action profiles.

In many real-life situations players privately observe imperfect signals about past actions. For example, each firm in a cartel privately observes its own sales, =which contain imperfect information about secret price cuts that its competitors offer to some of their customers. Formal analysis of private monitoring began with the pioneering work of Sekiguchi (1997). Since then, several papers have presented various folk theorem results that have shown that efficiency can be achieved also with private monitoring (see Kandori, 2002; Mailath and Samuelson, 2006, for surveys of this literature).

The most commonly used equilibrium in the literature on private monitoring is the *belief-free equilibrium* in which the continuation strategy of each player is a best reply to his opponent's strategy at every private history. These equilibria are called "belief-free" because a player's belief about his opponent's history is not needed to compute a best reply. Piccione (2002) and Ely and Välimäki (2002) present folk theorem results for the repeated Prisoner's Dilemma using belief-free equilibria under the assumptions that the monitoring technology is almost perfect and the players are sufficiently patient. Ely, Hörner, and Olszewski (2005),

*Email: yuval.heller@economics.ox.ac.uk. A previous version of this manuscript was called "Instability of Equilibria with Private Monitoring". I would like to express our deep gratitude to Mehmet Ekmekci, Peter Eso, Michihiro Kandori, Erik Mohlin, Satoru Takahashi, Jorgen Weibull, Yuichi Yamamoto, the associate editor and the referees, for many useful comments.

Miyagawa, Miyahara, and Sekiguchi (2008), and Yamamoto (2009, 2014) extend the folk theorem results that rely on belief-free equilibria to general repeated games and to costly observability. Kandori and Obara (2006) study a setup of imperfect *public* monitoring and show that belief-free private strategies can improve the efficiency relative to the maximal efficiency obtained by public strategies. Takahashi (2010) applies the belief-free equilibria to obtain folk theorem results for repeated games in which the players are randomly matched with a new opponent at each round.

The results of the present paper show that belief-free equilibria are not robust against small perturbations in the behavior of potential opponents, and that this instability is extreme in games with two actions that admit a strict equilibrium, such as the Prisoner's Dilemma.

Instability of Belief-free Equilibria in General Games One of the leading justifications for using Nash equilibrium to predict behavior is its interpretation as a stable convention in a population of potential players. Suppose that individuals in a large population are repeatedly drawn to play a game, and that initially all individuals play the strategy s^* but occasionally a small group of agents may experiment with a different strategy s'. If this induces the experimenting agents to gain more than the incumbents, then the population will move away from s^* toward s'. Thus, strategy s^* is *evolutionarily (neutrally) stable* (Maynard-Smith and Price, 1973) if (1) it is a best reply to itself (i.e., it is a symmetric Nash equilibrium),[1] and (2) it achieves a strictly (weakly) higher payoff against any other best-reply strategy s': $u(s^*, s') > u(s', s')$. One example of an evolutionarily stable strategy in the repeated Prisoner's Dilemma is the strategy of always defecting regardless of the history.

A belief-free equilibrium is *trivial* if it induces the play of a Nash equilibrium at all periods. My first result shows that only trivial belief-free equilibria may satisfy evolutionary stability. My second result characterizes a mild sufficient condition for a belief-free equilibrium to be neutrally unstable. In particular, if the underlying game admits an action profile that is a strict equilibrium, and the belief-free equilibrium plays this action profile with probability strictly between zero and one at some period, then the belief-free equilibrium is not neutrally stable. The intuition of these results is as follows. As observed by Ely, Hörner, and Olszewski (2005, Section 2.1) at each period t, the set of optimal actions in a belief-free equilibrium is independent of the private history. This implies that mutants who play a symmetric Nash equilibrium in an auxiliary game in which players are allowed to choose only from the set of optimal actions weakly outperform the incumbents (and strictly outperform the incumbents if this set includes a strict equilibrium).

Strong Instability in 2×2 Games with Strict Equilibria The existing notions of stability, namely, evolutionary and neutral stability, are arguably too-strong refinements, as demonstrated in the rock-paper-scissors game (see Section 2.3) that admits a unique Nash equilibrium that is not neutrally stable, but that is a plausible prediction for the long-run average behavior in the population (see, e.g., Benaïm, Hofbauer, and Hopkins, 2009). Motivated by this, I present a novel, and much weaker, notion of stability. I say that a strategy s is vulnerable to strategy s', if: (1) strategy s' is a best reply to s, and (2) strategy s' achieves a strictly higher payoff against itself: $u(s', s') > u(s, s')$. I say that a symmetric Nash equilibrium s^* is *weakly stable* if there does not exist a finite sequence of strategies $(s_1, ..., s_K)$, such that: (1) strategy s^* is

[1]To simplify the exposition I focus in most of the paper on symmetric equilibria in symmetric games, and I extend the analysis to asymmetric equilibria and asymmetric games in Section 4.

vulnerable to s_1, (2) each strategy s_k is vulnerable to s_{k+1}, and (3) strategy s_K is evolutionarily stable.[2]

The definition implies that any symmetric game admits a weakly stable strategy, and that if s^* is not weakly stable, then it is not a plausible prediction of long-run behavior. This is because as soon as a small group of agents experiments with playing s_1, the population diverges to s_1. If this is followed by an invasion of small groups of agents who play s_2, then the population diverges to s_2, and after a finite number of such sequential invasions, the population diverges to s_K, and it will remain in s_K in the long run (due to s_K being evolutionarily stable).[3] A simple example of a non-weakly stable equilibrium is the mixed equilibrium in a coordination game, for which every small perturbation takes the population to one of the pure equilibria.

My next result focuses on games with two actions that admit a strict equilibrium (such as the Prisoner's Dilemma), and shows that only trivial belief-free equilibria satisfy the mild refinement of weak stability. The intuition for the Prisoner's Dilemma is that any belief-free equilibrium is vulnerable to a deterministic strategy s' in which the players defect at each period in which defection is an optimal action with respect to the belief-free equilibrium, and this strategy s' is vulnerable to the evolutionarily stable strategy of always defecting.

Next, I focus on the repeated Prisoner's Dilemma and show that also among all the non-belief-free equilibria in the existing literature on private monitoring, only defection satisfies the mild refinement of weak stability.[4] In particular, I show that (1) always defecting is the unique belief-free review-strategy equilibria (Matsushima, 2004; Yamamoto, 2007; Deb, 2012; Yamamoto, 2012) that satisfies weak stability, and (2) none of the "belief-based" equilibria of Bhaskar and Obara (2002) is weakly stable.

1.1 Related Literature and Contribution

Conditionally Correlated Signals A few papers in the literature yield stable cooperation if the private signals are sufficiently correlated conditional on the action profile. Mailath and Morris (2002, 2006), Hörner and Olszewski (2009), and Mailath and Olszewski (2011) show that when the private signals are almost perfectly correlated conditional on the action profile (i.e., when there is *almost public monitoring*), then any sequential equilibrium of the nearby public monitoring game with bounded memory remains an equilibrium also with almost public monitoring. Some of these equilibria are evolutionarily stable, and, in particular, cooperation can be the outcome of an evolutionarily stable strategy.

Kandori (2011) presents the notion of weakly belief-free equilibria, in which the strategy of each player is a best reply to any private history of the opponent up to the actions of the previous round. Unlike standard belief-free equilibria, players need to form the correct beliefs about the signal obtained by the opponent in the previous round. Kandori (2011) demonstrates that if there is sufficient correlation between private signals (conditional on the action profile), then the game admits a strict weakly belief-based equilibrium that yields substantial cooperation. The strictness of the equilibrium implies that it satisfies the refinement of evolutionary stability. In the discussion paper version of his paper Kandori (2009) points out that the

[2] Remark 2 discusses the relation between weak stability and the structurally similar notion of "robustness against indirect invasions" of Van Veelen (2012).

[3] I assume that these experimentations are infrequent enough such that strategies that are outperformed following the entry of a group of experimenting agents become sufficiently rare before a new group of agents starts experimenting with a different behavior.

[4] Except equilibria that rely on sufficient correlation between private signals of different players (Mailath and Morris, 2002, 2006; Kandori, 2011), and equilibria that rely on communication between players at each round of the game (Compte, 1998; Kandori and Matsushima, 1998; Obara, 2009), as discussed in Section 1.1.

specific non-trivial belief-free equilibria in the existing literature are not evolutionarily stable in the repeated Prisoner's Dilemma. The present paper substantially strengthens Kandori's observation in at least two important ways: (1) I show that any non-trivial belief-free equilibrium of any underlying game is not evolutionarily stable (and under mild assumptions, it is also not neutrally stable), and (2) I show that when the underlying game is the Prisoner's Dilemma (or, more broadly, any underlying game with two actions and a strict equilibrium), then the belief-free equilibria fail to satisfy a very mild refinement of stability (i.e., the belief-free equilibria violate weak stability, rather than only violating evolutionary stability).

Communication and Conditionally Independent Signals Compte (1998), Kandori and Matsushima (1998), and Obara (2009) present folk theorem results that rely on (noiseless) communication between the players at each stage of the repeated game. The players use this communication to publicly report (possibly with some delay) the private signals they obtain. These equilibria are constructed such that the players have strict incentives while playing, and such that they are always indifferent between reporting the truth and lying regardless of the reporting strategy of the opponent. One can show that this property implies that these equilibria are neutrally stable, and hence also weakly stable.[5]

The present paper shows that all the mechanisms in the existing literature can yield only defection as the outcome of a weakly stable equilibrium in the repeated Prisoner's Dilemma with conditionally independent imperfect monitoring. I leave for future research the open question whether any new mechanism may yield cooperation as a stable outcome with conditionally independent private monitoring. This open question has interesting implications for antitrust laws. If the answer to this question is negative, then it would suggest that communication between players is critical to obtain collusive behavior whenever the private imperfect monitoring between the firms is such that the conditional correlation between the private signals is sufficiently low.[6,7]

One promising direction toward the solution of this open question might rely on the methods developed in Heller and Mohlin (2015) for the related setup of random matching and partial observation of the partner's past behavior. In that setup, Heller and Mohlin (2015) characterize conditions under which only defection is stable, and construct novel mechanisms to sustain stable cooperative equilibria whenever these conditions are not satisfied.

Robustness Sugaya and Takahashi (2013) show that "generically" only belief-free equilibria are robust to small perturbations in the monitoring structure. Our main result shows that belief-free equilibria (except for defection) are not robust to small perturbations in the behavior of the potential opponents. Taken together,

[5] The argument for neutral stability is sketched as follows. Having strict incentives while playing implies that any best-reply strategy must induce the same play on the equilibrium path, and differ from the incumbent strategy only by sending false reports. The fact that players are always indifferent between reporting the truth and lying implies that any such best-reply strategy yields the same payoff as the incumbent strategy (both when the opponent is an incumbent as well as when he is a mutant that follows a best-reply strategy).

[6] This empirical prediction can be tested experimentally by comparing how subjects play the repeated Prisoner's Dilemma with private monitoring and conditionally independent signals with and without the ability to communicate by exchanging "cheap talk" messages. Matsushima, Tanaka, and Toyama (2013) experimentally study this setup without communication, and their findings suggest that the subjects' behavior is substantially different than the predictions of the belief-free equilibria (in particular, subjects retaliate more severely when monitoring is more accurate). I am not aware of any experiment that studies this setup with communication.

[7] See also the recent related result of Awaya and Krishna (2015), which deals with sequential equilibria of oligopolies under some plausible private monitoring structures, and shows that cheap talk communication allows one to achieve a higher level of collusion relative to the maximal level that one can achieve without communication.

the two results suggest that defection is the unique equilibrium outcome of the repeated Prisoner's Dilemma that is robust to both kinds of perturbations.[8]

Structure The model is described in Section 2. Section 3 presents the results for symmetric games. Sections 4 extends the analysis to asymmetric games.

2 Model

2.1 Games with Private Monitoring

I analyze a two-player δ-discounted repeated game with private monitoring. I use the index $i \in \{1,2\}$ to refer to one of the players, and $-i$ to refer to the opponent. Each player i has a finite action set A_i and a finite set of signals Σ_i. An action profile is an element of $A_1 \times A_2$. I use ΔW to represent the set of probability distributions over a finite set W. Let ΔA_i and $\Delta A_1 \times \Delta A_2$ represent respectively the set of mixed actions for player i and mixed action profiles. For each player i let $\pi_i : A_1 \times A_2 \to R^n$ denote the payoff function, which is extended to mixed actions in the standard (linear) way.

For each possible action profile $(a_1, a_2) \in A_1 \times A_2$, the monitoring distribution $m(\cdot|a_1, a_2)$ specifies a joint probability distribution over the set of signal profiles $\Sigma_1 \times \Sigma_2$. When action profile a is played and signal profile (σ_1, σ_2) is realized, each player i privately observes his corresponding signal σ_i. Letting $\tilde{u_i}(a_i, \sigma_i)$ denote the payoff to player i from action a_i and signal σ_i, I can represent stage payoffs as a function of mixed action profiles alone:

$$u_i(\alpha_1, \alpha_2) = \sum_{(a_1,a_2)\in A_1\times A_2} \sum_{\sigma_i \in \Sigma_i} \alpha_1(a_1) \cdot \alpha_2(a_2) \cdot m_i(\sigma_i|a_1, a_2) \cdot \tilde{u}(a_i, \sigma_i).$$

To simplify the presentation of the results, I assume that the monitoring structure has full support, i.e., that each signal profile is observed with positive probability after each action profile. Formally:[9]

Assumption 1. *The monitoring structure has* full support: $m(\sigma_1, \sigma_2|a_1, a_2) > 0$ *for each action profile* $(a_1, a_2) \in A_1 \times A_2$ *and each signal profile* $(\sigma_1, \sigma_2) \in \Sigma_1 \times \Sigma_2$.

One example of a monitoring structure with full support is the *conditionally independent ϵ-perfect monitoring* in which each player privately observes his opponent's last action with probability $1-\epsilon$ and observes the opposite action with the remaining probability ϵ.

A t-length (private) history for player i is a sequence that includes the action played by the player and the observed signal in each of the previous t rounds of the game. Let $H_i^t := (A_i \times \Sigma_i)^t$ denote the set of all such histories for player i. Each player's initial history is the null history, denoted by ϕ. A pair of t-length

[8] Matsushima (1991) presents a related result by showing that defection is the unique *pure* equilibrium in the repeated Prisoner's Dilemma in which signals are conditionally independent and Nash equilibria are restricted to being independent of payoff-irrelevant private histories. As demonstrated by the "belief-based" equilibria of Bhaskar and Obara (2002), the uniqueness result does not hold for *mixed* equilibria (the mixed "belief-based" equilibria achieve cooperation even though the behavior of the players is independent of payoff-irrelevant private histories, and signals may be conditionally independent).

[9] The results can be adapted to a setup in which the monitoring structure does not have full support. The adaptation requires changing two definitions (and related minor adaptations to the proofs): (1) extending the set of of trivial belief-free equilibria in Definition 2, such that it relates only to histories that occur with positive probability, and (2) refining Definition 5 of weak stability by allowing the strategy s' to be neutrally stable, rather than evolutionarily stable (because if the monitoring structure does not have full support, then no strategy is evolutionarily stable).

histories (or history for short) is denoted by h^t. Let H^t denote the set of all t-length histories, $H = \cup_t H^t$ the set of all histories, and $H_i = \cup_t H_i^t$ the set of all private histories for i.

2.2 Belief-free Equilibria

A repeated-game (behavior) strategy for player i is a mapping $s_i : H_i \to \Delta(A_i)$. For history h_i^t , let $s_i|_{h_i^t}$ denote the continuation strategy derived from s_i following history h_i^t. Specifically, if $h_i\hat{h}_i$ denotes the concatenation of the two histories h_i and $\hat{h}_i$, then $s_i|_{h_i^t}$ is the strategy defined by $s_i|_{h_i^t}\left(\hat{h}_i\right) = s_i\left(h_i\hat{h}_i\right)$. Given a strategy profile $\overrightarrow{s} = (s_1, s_2)$, let $B_i\left(\overrightarrow{s}|h_{-i}^t\right)$ denote the set of continuation strategies for i that are best replies to $s_{-i}|_{h_{-i}^t}$.

Definition 1 (Ely, Hörner, and Olszewski 2005)**.** A strategy profile $\overrightarrow{s}^* = (s_1^*, s_2^*)$ is *belief-free* if for every h^t, $s_i^*|_{h_i^t} \in B_i\left(\overrightarrow{s}^*|h_{-i}^t\right)$ for $i \in \{1, 2\}$.

The condition characterizing a belief-free strategy profile is stronger than that characterizing a sequential equilibrium. In a sequential equilibrium, a player's continuation strategy is the player's best reply given his belief about his opponent's continuation strategy, that is, given a unique probability distribution over the opponent's private histories. In a belief-free strategy profile, a player's continuation strategy is his best reply to his opponent's continuation strategy at every private history. In other words, a sequential equilibrium is a belief-free strategy profile if it has the property that a player's continuation strategy is still the player's best reply when he secretly learns about his opponent's private history.

A simple kind of of a belief-free equilibrium, is a strategy profile in which the players play a Nash equilibrium of the underlying game at all periods, and this equilibrium is independent of the history of play. I call such belief-free equilibria trivial. Formally, let $NE\left(\left(A_1, A_2\right), \left(u_1, u_2\right)\right)$ denote the set of Nash equilibria of the underlying game.

Definition 2. A belief-free equilibrium (s_1, s_2) is *trivial* if for every two histories h^t, $\tilde{h}^t$ of length t:

$$\left(s_i\left(h_i^t\right), s_{-i}\left(h_{-i}^t\right)\right) = \left(s_i\left(\tilde{h}_i^t\right), s_{-i}\left(\tilde{h}_{-i}^t\right)\right) \in NE\left(\left(A_1, A_2\right), \left(u_1, u_2\right)\right).$$

2.3 Evolutionary Stability in Symmetric Games

In what follows, I study evolutionary stability in symmetric games (which are played within a single population). Section 4 extends the analysis to asymmetric games. In the setup of symmetric games I will omit the index i: $A := A_1 = A_2$, $\tilde{u} := \tilde{u}_1 = \tilde{u}_2$, $u := u_1 = u_2$, and $m := m_1 = m_2$. I will say that a strategy s is a symmetric Nash (belief-free) equilibrium if the symmetric strategy profile (s, s) is a Nash (belief-free) equilibrium.

I present a refinement of a symmetric Nash equilibrium that requires robustness to a small group of agents who experiment with a different behavior (see, Weibull, 1995, for an introductory textbook). Suppose that individuals in a large population are repeatedly drawn to play a two-person symmetric game, and that there is an underlying dynamic process of social learning in which more successful strategies (which induce higher average payoffs) become more frequent. Suppose that initially all individuals play the equilibrium strategy s^*. Now consider a small group of agents (called, *mutants*) who play a different strategy s'. If s' is not a best reply to s^*, then if the mutants are sufficiently rare they will be outperformed. If s' is a best reply to s^*, then the relative success of the incumbents and the mutants depends only on the average payoff they achieve when matched against a mutant opponent. If the incumbents achieve a higher payoff when matched against the mutants, then the mutants are outperformed. Otherwise, the mutants outperform the incumbents, and their strategy gradually takes over the population.

The formal definition is as follows. Let $U(s,s')$ denote the expected discounted payoff to a player following strategy s and facing an opponent who plays strategy s'.

Definition 3 (Maynard-Smith and Price, 1973; Maynard-Smith, 1982)**.** A symmetric Nash equilibrium s^* is neutrally (evolutionarily) stable if $U(s^*,s') \geq U(s',s')$ $(U(s^*,s') > U(s',s'))$ for each strategy $s^* \neq s' \in B(s^*)$.

Note that any strict symmetric equilibrium is evolutionarily stable. The key difference between evolutionary and neutral stability is whether the mutants are allowed to obtain the same payoff as incumbents. As a result neutrally stable strategies (which are not evolutionary stable) may be vulnerable to a random drift of the population away from the initial state. The existing literature typically uses evolutionary stability as the strongest notion of stability, and neutral stability as the weakest notion (and it studies various notions in between, such as evolutionarily stable sets; see Thomas, 1985).

One may argue that neutral stability is still "too strong" a refinement because: (1) some games do not admit any neutrally stable strategies, and (2) some equilibria that are not neutrally stable are plausible predictions of the time-average behavior in the game. This is demonstrated in the rock-paper-scissors game in Table 1 (left side). The unique symmetric equilibrium is $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$, which is not neutrally stable (because $R \in B\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$ and $U\left(\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right),R\right) = -\frac{1}{3} < U(R,R) = 0$. One can show that although $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$ is not neutrally stable, still, under mild assumptions on the dynamics, the time average of the aggregate play converges to $\left(\frac{1}{3},\frac{1}{3},\frac{1}{3}\right)$ (Benaïm, Hofbauer, and Hopkins, 2009).

Table 1: Examples of Symmetric Games

	R	P	S
R	0, 0	-2, 1	1, -2
P	1, -2	0, 0	-2, 1
S	-2, 1	1, -2	0, 0

Rock-Paper-Scissors.

	a	b
a	1, 1	0, 0
b	0, 0	1, 1

Coordination Game

	c	d
c	1, 1	l, $1{+}g$
d	$1{+}g$, l	0, 0

Prisoner's Dilemma $(g > 0 > l)$

Hawk-Dove Game $(g, l > 0)$

This motivates me to present a much weaker stability refinement that focuses only on vulnerability to mutant strategies that are evolutionarily stable. Strategy s^* is vulnerable to strategy s' if strategy s' can take over a population that initially plays s^*. Formally:

Definition 4. Strategy s^* is *vulnerable* to strategy s' if $s' \in B(s^*)$ and $U(s^*,s') < U(s',s')$.

Observe that a neutrally stable strategy is not vulnerable to any other strategy.

Remark 1. All the results in this paper remain the same if one weakens the definition of vulnerability à la Swinkels (1992), and requires s' to be an "equilibrium entrant," i.e., a best reply to the post-entry population for any sufficiently small $\epsilon > 0$: $s' \in B((1-\epsilon)\cdot s^* + \epsilon \cdot s')$.

A symmetric Nash equilibrium s^* is weakly stable if if there does not exist a finite sequence of strategies that starts at s^*, ends in an evolutionarily stable strategy, and each of whose strategies is vulnerable to its successor. Formally:

Definition 5. A symmetric Nash equilibrium s^* is *weakly stable* if there does not exist a finite non-empty sequence of strategies $\left(s^1, ..., s^K\right)$ such that: (1) strategy s^* is vulnerable to s^1, (2) for each $1 \leq k < K$ strategy s^k is vulnerable to s^{k+1}, and (3) strategy s^K is evolutionarily stable.

Observe that:

1. Any neutrally stable strategy is weakly stable.

2. Any game admits a weakly stable strategy.

3. If strategy s^* is not weakly stable, then it is not a plausible prediction of long-run behavior in the population. Even if the population initially plays s^*, then as soon as a small group of agents experiments with playing $\tilde{s}$, the population diverges to $\tilde{s}$. If this is followed by another small group of agents who play s', then the population will converge to s', and will remain there in the long run. Note that our argument relies on the assumption that these experimentations are infrequent enough that strategies that are outperformed following the entry of a group of experimenting agents become sufficiently rare before a new group of agents starts experimenting with a different behavior.

4. The notion of weak stability is able to strictly refine Nash equilibrium only if the game admits an evolutionarily stable strategy.

Remark 2. Definition 5 is structurally similar to Van Veelen's (2012) notion of *robustness against indirect invasions*. A strategy s^* is robust against indirect invasions if there does not exist a sequence of strategies $\left(s_1, ..., s_n\right)$, such that s^* is *weakly* vulnerable to s_1 (i.e., $s_1 \in B\left(s^*\right)$ and $U\left(s^*, s_1\right) \leq U\left(s_1, s_1\right)$), each s_k is weakly vulnerable to s_{k+1}, and s_{K-1} is (strictly) vulnerable to s_K. Note that Van Veelen's notion of robustness refines neutral stability (i.e., it is between evolutionary stability and neutral stability), while weak stability weakens neutral stability (i.e., weak stability is between neutral stability and a symmetric Nash equilibrium).

Remark 3. Definition 5 allows vulnerability to an evolutionarily stable strategy through an arbitrary number of sequential invasions (denoted by K). Almost all the results in this paper rely on at most two sequential invasions (i.e., $K = 2$). The only exception is the case of the one-sided Prisoner's Dilemma in Proposition 5, which requires three steps (i.e., $K = 3$). Moreover, if we focus on the existing belief-free equilibria for the repeated Prisoner's Dilemma in the literature (e.g., Ely and Välimäki, 2002; Piccione, 2002), then most of them are directly vulnerable to an invasion by players who always defect (i.e., $K = 1$).

3 Results for Symmetric Games

Ely, Hörner, and Olszewski (2005) characterize the set of belief-free equilibrium payoffs, and show that such strategies support a large set of payoffs. In what follows, I show that non-trivial belief-free equilibria are unstable in general games, and "extremely" unstable in games with two actions and strict equilibria.

3.1 Instability in General Symmetric Games

My first result shows that any evolutionarily stable belief-free equilibrium must be trivial (i.e., it induces the play of Nash equilibria of the underlying game at all periods)

The sketch of the proof is as follows. As observed by Ely, Hörner, and Olszewski (2005, Section 2.1) at each period t, the set of optimal actions is independent of the private history. This implies that mutants who play a symmetric Nash equilibrium in an auxiliary game in which players are only allowed to choose from the set of optimal actions weakly outperform the incumbents. If the belief-free equilibrium is non-trivial, then the mutants' play differs from the incumbents' play, which implies that the belief-free equilibrium is not evolutionarily stable.

Proposition 1. *Let s^* be a symmetric belief-free equilibrium that is also evolutionarily stable. Then s^* is trivial.*

Proof. A continuation strategy z_i is a *belief-free sequential best reply* to s^* starting from period t if

$$z_i|h_i^{\tilde{t}} \in B_i(s^*|h_{-i}^{\tilde{t}})\,\forall\tilde{t}\geq t \text{ and } h^{\tilde{t}}\in H^{\tilde{t}};$$

the set of belief-free sequential best-replies beginning from period t is denoted by $B_i^t(s^*)$. Let

$$\mathcal{A}_i^t=\left\{a\in A|\exists z_i\in B_i^t(s^*),\,\exists h_i^t \text{ such that } z_i\left(h_i^t\right)(a_i)>0\right\};$$

denote the set of actions in the support of some belief-free sequential best reply starting from period t. As observed by Ely, Hörner, and Olszewski (2005, Section 2.1), $\exists h_i^t$ can be replaced with $\forall h_i^t$, because if z_i is a belief-free sequential best reply to s_{-i} and every continuation strategy $z_i|h_i^t$ gets replaced with the strategy $z_i|\tilde{h}_i^t$ for a given $\tilde{h}_i^t$, then the strategy z_i so obtained is also a belief-free sequential best reply to s_{-i}. Note that the symmetry of the profile (s^*,s^*) implies that $\mathcal{A}^t:=\mathcal{A}_i^t=\mathcal{A}_j^t$. Following Ely, Hörner, and Olszewski (2005), I call $\mathcal{A}^t$ the *regime at period t*.

For each period t, let $\alpha^t\in\Delta\left(\mathcal{A}^t\right)$ be a symmetric Nash equilibrium in the symmetric game $(\mathcal{A}^t,u)$ in which players are restricted to choose actions only in $\mathcal{A}^t\subseteq A$. Let s' be the strategy in which each player plays the mixed action α^t at each period t. The definition of the regimes $\left(\mathcal{A}^t\right)_t$ implies that a mutant player who follows strategy s' best-replies to an incumbent who follows s^*, i.e., $U\left(s',s^*\right)=U\left(s^*,s^*\right)$. The definition of α^t implies that a mutant achieves a weakly higher payoff relative to the incumbents when facing another mutant: $U\left(s',s'\right)\geq U\left(s^*,s'\right)$. This implies that s^* can be evolutionarily stable only if $s'=s^*$, which implies that s^* is trivial. □

As evolutionary stability is a strong refinement, it is desirable to show that non-trivial belief-free equilibria also fail to satisfy weaker notions of stability. In what follows I present a mild restriction on the belief-free equilibrium, which implies that it cannot satisfy neutral stability. Given a belief-free symmetric equilibrium s^*, let $\gamma^t=\gamma^t\left(s^*\right)\in\Delta\left(\mathcal{A}^t\right)$ be the marginal distribution of actions played by each player at period t. I show that if there exists a mixed action $\beta^{t_0}\in\Delta\left(\mathcal{A}^{t_0}\right)$ at some period t_0 such that $u\left(\gamma^{t_0}\left(s^*\right),\beta^{t_0}\right)<u\left(\beta^{t_0},\beta^{t_0}\right)$, then the belief-free equilibrium is not neutrally stable. Formally:

Proposition 2. *Let s^* be a symmetric belief-free equilibrium. Assume that there exists period t_0 and mixed action $\beta^{t_0}\in\Delta\left(\mathcal{A}^{t_0}\right)$ such that $u\left(\gamma^{t_0}\left(s^*\right),\beta^{t_0}\right)<u\left(\beta^{t_0},\beta^{t_0}\right)$. Then s^* is not neutrally stable.*

Proof. Let the mutant strategy s' be defined as follows:

$$s'\left(h_i^t\right) = \begin{cases} s^*\left(h_i^t\right) & t < t_0 \\ \beta^{t_0} & t = t_0 \\ \alpha^t & t > t_0, \end{cases}$$

where $\alpha^t \in \Delta\left(\mathcal{A}^t\right)$ is a symmetric Nash equilibrium in the game $\left(\mathcal{A}^t, u\right)$, as defined in the previous proof. The definition of s' immediately implies by a similar argument to the previous proof that $U\left(s', s^*\right) = U\left(s^*, s^*\right)$ and $U\left(s', s'\right) > U\left(s^*, s'\right)$, which implies that s^* is not neutrally stable. $\square$

Proposition 2 has a simple corollary when the underlying game admits a symmetric strict equilibrium $(\hat{a}, \hat{a})$. If there exists any period at which action $\hat{a}$ is played with a probability strictly between zero and one, then the belief-free equilibrium is not neutrally stable. Formally:

Corollary 1. *Let s^* be a symmetric belief-free equilibrium. Let $\hat{a} \in A$ be a strict symmetric equilibrium of the underlying game. Assume that there exists period t_0 such that $0 < \gamma_{t_0}\left(s^*\right)(\hat{a}) < 1$. Then s^* is not neutrally stable.*

Proof. The result is immediate from the previous proof when taking $\beta^{t_0} \equiv \hat{a}$. $\square$

3.2 Strong Instability for 2 × 2 Games that Admit Symmetric Strict Equilibria

The Prisoner's Dilemma (as described in Table 1) is an important application of belief-free equilibria. In particular, Piccione (2002) and Ely and Välimäki (2002) show a folk theorem result, which relies on belief-free equilibria, for the symmetric Prisoner's Dilemma game with almost perfect monitoring technology and sufficiently patient players. In this section I show that only trivial belief-free equilibria satisfy the mild refinement of weak stability in 2×2 symmetric games that admit a symmetric strict equilibrium. In particular, the symmetric Prisoner's Dilemma game admits a unique weakly stable belief-free equilibrium in which both players defect at all periods.

Proposition 3. *Assume that the symmetric underlying game has two actions $A = \{c, d\}$, and that one of these actions, d, is a symmetric strict equilibrium. Let s^* be a symmetric belief-free equilibrium. If s^* is weakly stable, then it is trivial.*

Proof. Assume first that $\gamma^t\left(s^*\right)$ is pure at all periods t. This implies that s^* induces a deterministic play that is independent of the observed signals. Thus a player's best reply coincides with his myopic best reply, which implies that the pure action profile played at each period must be an equilibrium of the underlying game (i.e., s^* is trivial).

Otherwise, there exists time t such that: (1) both actions are best replies ($\mathcal{A}^t = \{c, d\}$), and (2) there is a private history h_i^t after which player i cooperates with positive probability, i.e., $\exists h_i^t$ s.t. $s_i^*\left(h_i^t\right)(C) > 0$. Let s' be the pure strategy that plays d at all periods except in those in which c is the unique best reply to s^*:

$$s'\left(h_i^t\right) = \begin{cases} c & \mathcal{A}_i^t = \{c\} \\ d & \text{otherwise.} \end{cases}$$

The definition of s' implies that s' is a best reply to s^*, and that $U(s^*, s') < U(s', s')$ (because at any period in which s' and s^* do not coincide, strategy s' plays the unique best reply to itself, namely, action d). If c is also a symmetric strict equilibrium of the underlying game, then s' is evolutionarily stable. Otherwise, let $s_d \equiv d$ be the strategy that always plays action d. The fact that only action d is a symmetric strict equilibrium implies that action d is weakly dominant, strategy s_d is a best reply to s', $U(s_d, s_d) > U(s', s_d)$, and strategy s_d is evolutionarily stable. This implies that s^* is not weakly stable. □

3.3 Strong Instability of Other Equilibria in the Repeated Prisoner's Dilemma

Proposition 3 shows that non-trivial belief-free equilibria do not satisfy weak stability in the repeated Prisoner's Dilemma. In this subsection I extend this result to two other kinds of equilibria in the literature on private monitoring: belief-free review-strategy equilibria and belief-based equilibria.

Instability of Belief-free Review-strategy Equilibria. Matsushima (2004), Yamamoto (2007, 2012), and Deb (2012) use the notion of a *belief-free review-strategy equilibrium* (also called *block equilibrium*) in which (1) the infinite horizon is regarded as a sequence of review phases such that each player chooses a constant action throughout a review phase, and (2) at the beginning of each review phase, a player's continuation strategy is a best reply regardless of the history. A simple adaptation of the proof of Proposition 3 show that defection is the unique weakly stable symmetric belief-free review-strategy equilibrium.[10]

The adaptation of the proof is done as follows. Let s^* be a symmetric belief-free review-strategy equilibrium. Let $(t_l)_{l=1}^{\infty}$ be the increasing sequence of starting times for the review phases. Let $\mathcal{A}_i^l \subseteq \{c, d\}$ denote the set of actions in the support of some sequential best reply that begins from period t_l and plays the same action up to the end of the l-th review phase. Let s' be the deterministic strategy that always defects except in review phases in which $\mathcal{A}_i^l = \{c\}$. An analogous argument to the proof of Proposition 3 shows that s^* is vulnerable to s', and s' is vulnerable to s_d.

Instability of Belief-based Equilibria Bhaskar and Obara (2002) (extending Sekiguchi, 1997) present a folk theorem result for the repeated Prisoner's Dilemma that does not rely on belief-free equilibria. Instead, the best reply of each player depends on his belief about the private history of the opponent ("belief-based equilibria"). In what follows, I show that these equilibria are not weakly stable.

Bhaskar and Obara (2002) consider a symmetric signaling structure with two signals $\Sigma_i = \{C, D\}$, where C (resp., D) is more likely when the opponent plays c (resp., d). Given any action profile, there is a probability of $\epsilon > 0$ that exactly one player receives a wrong signal, and a probability of $\xi > 0$ that both players receive wrong signals. Bhaskar and Obara present for each $0 < x < 1$ a symmetric sequential equilibrium s_x that yields a payoff of at least x whenever ϵ and ξ are sufficiently small. This construction is the key element in their folk theorem result. In what follows I sketch this equilibrium s_x, and then show that it is not weakly stable.

Let s_T be the trigger strategy: cooperate as long as all observed signals are C-s, and defect in the remaining game if signal D is ever observed. The strategy s_x divides the entire game into disjoint subgames (say, into n subgames, where subgame k includes the rounds that are equal to k modulo n), and the play in

[10] Similarly, one can further adapt the proof to show the instability of Sugaya's (2015) equilibria, in which each review phase is divided into several sub-phases, and players may switch their action at the beginning of each sub-phase.

each subgame is independent of the other subgames. Each player mixes at the first round of each subgame: he plays s_T (trigger strategy) with probability π and plays s_d (always defect) with the remaining probability. Bhaskar and Obara show that there exist a division into subgames and a mixing probability π such that (1) the expected discounted symmetric payoff of the game is at least x, and (2) strategy s_x is a sequential equilibrium.[11]

Claim 1. The symmetric sequential equilibrium s_x is not weakly stable.

Sketch of Proof. The fact that s_x mixes between s_d and s_T at the beginning of each subgame, implies that s_d is a best reply to s_x. Recall that s_d is evolutionarily stable and the unique best reply to itself. These observations immediately imply that s_x is not weakly stable. □

4 Analysis of Asymmetric Games

The previous section analyzed stability of equilibria in symmetric games. In this section I show that all of the instability results hold also for the case of asymmetric games. In addition, in this setup I am able to extend the strong instability result also for the Hawk-Dove game (which admits an asymmetric strict equilibrium).

4.1 Evolutionary Stability in Asymmetric Games

In this subsection I adapt the notions of stability to asymmetric games (or symmetric games, in which a player is able to condition on his play on his role, i.e., being Player 1 or Player 2).

A Nash equilibrium is evolutionarily stable if the incumbents achieve a higher payoff against any mutants who best-reply to the incumbents.

Definition 6 (Taylor, 1979; Weibull, 1995, Chapter 5.1)**.** A Nash equilibrium (s_1^*, s_2^*) is neutrally (evolutionarily) stable if $U_1(s_1^*, s_2') + U_2(s_1', s_2^*) \geq U_1(s_1', s_2') + U_2(s_1', s_2')$ $(U_1(s_1^*, s_2') + U_2(s_1', s_2^*) > U_1(s_1', s_2') + U_2(s_1', s_2'))$ for each strategy profile $(s_1', s_2') \neq (s_1^*, s_2^*)$ satisfying $s_1' \in B(s_2^*)$ and $s_2' \in B(s_1^*)$.

In what follows I adapt the notion of weak stability to the setup of asymmetric games. Strategy profile (s_1^*, s_2^*) is vulnerable to strategy profile (s_1', s_2') if the latter can take over a population that initially plays (s_1^*, s_2^*). Formally:

Definition 7. Strategy profile(s_1^*, s_2^*) is *vulnerable* to strategy profile (s_1', s_2') if for each player i, $U_i\left(s_i^*, s_{-i}^*\right) \leq U_i\left(s_i', s_{-i}^*\right)$, and, in addition, $U_1(s_1^*, s_2') + U_2(s_1', s_2^*) < U_1(s_1', s_2') + U_2(s_1', s_2')$.[12]

A Nash equilibrium is weakly stable if there does not exist a sequence of strategy profiles, starting with this equilibrium and ending with an evolutionarily stable strategy profile, such that each profile in the sequence is vulnerable to its successor. Formally:

[11] As observed by Bhaskar (2000) and Bhaskar, Mailath, and Morris (2008), these belief-based equilibria can be purified à la Harsanyi (1973) in a simple way (while this is not the case for belief-free equilibria). Nevertheless, I show they still do not satisfy weak stability.

[12] For asymmetric games we do not require each s_i' to be a best reply to s_{-i}^*. The reason for this is demonstrated in a one-shot coordination game (see Table 1), in which initially $s_1^* = a$ and $s_2^* = b$. A plausible way for a population to diverge from this unstable (and non-equilibrium) initial state, is by having the players in one of the populations change their action. However, a requirement that each s_i' would be a best reply to s^* rules out this plausible change.

Definition 8. A Nash equilibrium (s_1^*, s_2^*) is *weakly stable* if if there does not exist a finite non-empty sequence of strategy profiles $(\vec{s}^1, ..., \vec{s}^K)$ such that: (1) strategy profile (s_1^*, s_2^*) is vulnerable to (s_1^1, s_2^1), (2) for each $1 \leq k < K$, strategy profile (s_1^k, s_2^k) is vulnerable to (s_1^{k+1}, s_2^{k+1}), and (3) strategy profile (s_1^K, s_2^K) is evolutionarily stable.

4.2 Results for General Asymmetric Games

It is well known (see, Weibull 1995, Chapter 5.1, and the citations therein) that a strategy profile is evolutionarily stable if and only if it is a strict equilibrium. This immediately implies that only trivial belief-free equilibria may be evolutionarily stable in this setup.

Fact 1. *Let (s_1^*, s_2^*) be a belief-free equilibrium that is also evolutionarily stable. Then (s_1^*, s_2^*) is a trivial equilibrium.*

Next, we adapt Proposition 2 to this setup, and show that a mild restriction on the belief-free equilibrium implies that it cannot be neutrally stable.

Proposition 4. *Let (s_1^*, s_2^*) be a belief-free equilibrium. Assume that there exists period t_0, and mixed actions $\beta_1^{t_0} \in \Delta(\mathcal{A}_1{}^{t_0})$ and $\beta_2^{t_0} \in \Delta(\mathcal{A}_2{}^{t_0})$, such that*

$$u_1\left(\gamma_1^{t_0}(s_1^*, s_2^*), \beta_2^{t_0}\right) + u_2\left(\beta_1^{t_0}, \gamma_2^{t_0}(s_1^*, s_2^*)\right) < u_1\left(\beta_1^{t_0}, \beta_2^{t_0}\right) + u_2\left(\beta_1^{t_0}, \beta_2^{t_0}\right).$$

Then (s_1^*, s_2^*) is not neutrally stable.

Proof. For each period t, let $\alpha_1^t \in \Delta(\mathcal{A}_1^t)$, $\alpha_2^t \in \Delta(\mathcal{A}_2^t)$ be a Nash equilibrium in the game $((\mathcal{A}_1^t, \mathcal{A}_2^t), (u_1, u_2))$ in which each player is restricted to choose only actions from the regime at period t. Let the mutant strategy profile (s_1', s_2') be defined as follows:

$$s_i'(h_i^t) = \begin{cases} s_i^*(h_i^t) & t < t_0 \\ \beta_i^{t_0} & t = t_0 \\ \alpha_i^t & t > t_0, \end{cases}$$

The definition of s' immediately implies that $s_1' \in B(s_2^*)$, $s_2' \in B(s_1^*)$, and $U_1(s_1^*, s_2') + U_2(s_1', s_2^*) < U_1(s_1', s_2') + U_2(s_1', s_2')$, which implies that s^* is not neutrally stable. □

Proposition 4 has a simple corollary when the underlying game admits a strict equilibrium $(\hat{a}_1, \hat{a}_2)$. If there is any period at which the action profile $(\hat{a}_1, \hat{a}_2)$ is played with probability strictly between zero and one, then the belief-free equilibrium is not neutrally stable. Formally:

Corollary 2. *Let (s_1^*, s_2^*) be a belief-free equilibrium. Let $(\hat{a}_1, \hat{a}_2) \in A_1 \times A_2$ be a strict equilibrium of the underlying game. Assume that there exists period t_0 such that $0 < \gamma_i^{t_0}(s_i^*)(\hat{a}_i) < 1$ and $0 < \gamma_i^{t_0}(s_{-i}^*)(\hat{a}_{-i}) \leq 1$. Then (s_1^*, s_2^*) is not neutrally stable.*

Proof. The result is immediate from the previous proof when taking $\left(\beta_1^{t_0}, \beta_2^{t_0}\right) \equiv (\hat{a}_1, \hat{a}_2)$. □

4.3 Strong Instability for 2×2 Games that Admit Strict Equilibria

The final result shows that only trivial belief-free equilibria satisfy the mild refinement of weak stability in any 2×2 game that admits a strict equilibrium (i.e., any 2×2 generic game except a matching-pennies-like game that admits only a totally mixed equilibrium). Observe, that in contrast to the setup of symmetric games, the result also applies to Hawk-Dove games (see Table 1).

Proposition 5. *Assume that the underlying game has two actions for each player $A_i = \{c_i, d_i\}$, and that (d_1, d_2) is a strict equilibrium of the underlying game. Further assume a mild condition of generic payoffs: for each player i, $u_i(d_i, c_{-i}) \neq u_i(c_i, c_{-i})$. Let (s_1^*, s_2^*) be a belief-free equilibrium. If (s_1^*, s_2^*) is weakly stable, then it is a trivial equilibrium.*

Proof. Assume first that $\gamma_{\hat{i}}^t(s_1^*, s_2^*)$ is pure at all periods t for player $\hat{i}$. This implies that strategy $s_{\hat{i}}^*$ induces a deterministic play that is independent of the observed signals. This implies that the best-reply function of player $s_{-\hat{i}}^*$ coincides with his myopic best reply, which implies due to the mild condition of generic payoffs that $\gamma_{-\hat{i}}^t(s_1^*, s_2^*)$ is also pure at all periods. This implies that the strategy profile induces both players to follow a deterministic play that is independent of the observed signal, and that this deterministic play must be an equilibrium of the underlying game (i.e., (s_1^*, s_2^*) is trivial).

Otherwise, for each player i there exists a period $\hat{t}_i$ such that $0 < \gamma_i^{\hat{t}_i}(s_1^*, s_2^*)(d_i) < 1$. We complete the proof by studying two (exhaustive and mutually exclusive) cases, and showing in both cases that (s_1^*, s_2^*) is not weakly stable.

Case I: (c_1, c_2) is also a strict equilibrium (i.e., the underlying game is a coordination game).[13] We define the profile (s_1', s_2') as follows. For each period t, if the regime is a singleton for both players, the mutant strategies coincide with the incumbent strategies, i.e., if $|\mathcal{A}_1^t| = |\mathcal{A}_2^t| = 1$ then $s_i'(h_i^t) = s_i^*(h_i^t)$ for each history of length t. Otherwise (i.e., if one of the regimes is not a singleton), there exists a strict equilibrium in the regime, and the mutant strategies play this strict equilibrium; i.e., for each history of length t $\left(s_i'(h_i^t), s_{-i}'\left(h_{-i}^t\right)\right) = (c_i, d_{-i}) \in \mathcal{A}_i^t \times \mathcal{A}_{-i}^t$. It is immediate that the profile (s_1^*, s_2^*) is vulnerable to (s_1', s_2'). If the pure profile (s_1', s_2') plays a strict equilibrium at all periods then it is evolutionarily stable. Otherwise, let s_1'' be a strategy that plays at all periods the myopic best reply to the pure strategy s_2', and let $s_2'' \equiv s'_2$. It is immediate that (s_1', s_2') is vulnerable to (s_1'', s_2''), and that (s_1'', s_2'') is evolutionarily stable.

Case II: $d_{\hat{i}}$ is a strictly dominant action for player $\hat{i}$. We define the profile (s_1', s_2') as follows. For each period t, each player chooses a pure action independent of both the history of play and the observed signal:

1. For each player i, if the regime is a singleton ($|\mathcal{A}_i^t| = 1$), then i plays the unique action in $\mathcal{A}_i^t$.
2. If $\mathcal{A}_{\hat{i}}^t = \{c_{\hat{i}}, d_{\hat{i}}\}$, then player $\hat{i}$ plays $d_{\hat{i}}$.
3. If $\mathcal{A}_{-\hat{i}}^t = \left\{c_{-\hat{i}}, d_{-\hat{i}}\right\}$ and $\mathcal{A}_{\hat{i}}^t = \{c_{\hat{i}}\}$, then player $-\hat{i}$ plays the myopic best reply to $c_{\hat{i}}$.
4. If $\mathcal{A}_{-\hat{i}}^t = \left\{c_{-\hat{i}}, d_{-\hat{i}}\right\}$ and $\mathcal{A}_{\hat{i}}^t \neq \{c_{\hat{i}}\}$, then player $-\hat{i}$ plays $d_{-\hat{i}}$.

It is immediate that the profile (s_1^*, s_2^*) is vulnerable to (s_1', s_2'). If $(s_1', s_2') \equiv (d_1, d_2)$ then it is evolutionarily stable. Otherwise:

[13] Hawk-Dove games (Table 1) belong to the this class of coordination games in the setup of asymmetric games (or symmetric games in which a player may condition his play on his role as Player 1 or Player 2).

1. If $s'_{\hat{i}} \not\equiv d_{\hat{i}}$, then let $s''_{\hat{i}} \equiv d_{\hat{i}}$ and $s''_{-\hat{i}} \equiv s'_{-\hat{i}}$. It is immediate that (s'_1, s'_2) is vulnerable to (s''_1, s''_2). If $(s''_1, s''_2) \equiv (d_1, d_2)$, then it is evolutionarily stable. Otherwise, let $(s'''_1, s'''_2) \equiv (d_1, d_2)$, and it is immediate that (s''_1, s''_2) is vulnerable to (s'''_1, s'''_2), and that (s'''_1, s'''_2) is evolutionarily stable.[14]

2. If $s'_{\hat{i}} \equiv d_{\hat{i}}$ and $s'_{-\hat{i}} \not\equiv d_{-\hat{i}}$, then let $(s''_1, s''_2) \equiv (d_1, d_2)$, and it is immediate that (s'_1, s'_2) is vulnerable to (s''_1, s''_2) and that (s''_1, s''_2) is evolutionarily stable.

□

References

Abreu, D., D. Pearce, and E. Stacchetti (1990): "Toward a theory of discounted repeated games with imperfect monitoring," *Econometrica*, 58(5), 1041–1063.

Awaya, Y., and V. Krishna (2015): "On communication and collusion," *American Economic Review*, Forthcoming.

Benaïm, M., J. Hofbauer, and E. Hopkins (2009): "Learning in games with unstable equilibria," *Journal of Economic Theory*, 144(4), 1694–1709.

Bhaskar, V. (2000): "The Robustness of Repeated Game Equilibria to Incomplete Payoff Information," University of Essex.

Bhaskar, V., G. J. Mailath, and S. Morris (2008): "Purification in the infinitely-repeated prisoners' dilemma," *Review of Economic Dynamics*, 11(3), 515–528.

Bhaskar, V., and I. Obara (2002): "Belief-based equilibria in the repeated prisoner's dilemma with private monitoring," *Journal of Economic Theory*, 102(1), 40–69.

Compte, O. (1998): "Communication in repeated games with imperfect private monitoring," *Econometrica*, 66(3), 597–626.

Deb, J. (2012): "Cooperation and community responsibility: A folk theorem for repeated matching games with names," *Available at SSRN 1213102*.

Ely, J. C., J. Hörner, and W. Olszewski (2005): "Belief-free equilibria in repeated games," *Econometrica*, 73(2), 377–415.

Ely, J. C., and J. Välimäki (2002): "A robust folk theorem for the prisoner's dilemma," *Journal of Economic Theory*, 102(1), 84–105.

Fudenberg, D., and E. Maskin (1986): "The folk theorem in repeated games with discounting or with incomplete information," *Econometrica*, pp. 533–554.

Harsanyi, J. C. (1973): "Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points," *International Journal of Game Theory*, 2(1), 1–23.

[14] Note that we need three steps of vulnerability only if the underlying game is a one-sided Prisoner's Dilemma game. If the underlying game is a standard Prisoner's Dilemma, then we can immediately define $\left(s''_1, s''_2\right) \equiv (d_1, d_2)$.

Heller, Y., and E. Mohlin (2015): "Observations on cooperation," Unpublished.

Hörner, J., and W. Olszewski (2009): "How robust is the Folk Theorem?," *The Quarterly Journal of Economics*, 124(4), 1773–1814.

Kandori, M. (2002): "Introduction to repeated games with private monitoring," *Journal of Economic Theory*, 102(1), 1–15.

——— (2009): "Weakly Belief-Free Equilibria in Repeated Games with Private Monitoring," CIRJE F-Series CIRJE-F-491, CIRJE, Faculty of Economics, University of Tokyo.

——— (2011): "Weakly belief-free equilibria in repeated games with private monitoring," *Econometrica*, 79(3), 877–892.

Kandori, M., and H. Matsushima (1998): "Private observation, communication and collusion," *Econometrica*, 66(3), 627–652.

Kandori, M., and I. Obara (2006): "Efficiency in repeated games revisited: The role of private strategies," *Econometrica*, 74(2), 499–519.

Mailath, G. J., and S. Morris (2002): "Repeated games with almost-public monitoring," *Journal of Economic Theory*, 102(1), 189–228.

——— (2006): "Coordination failure in repeated games with almost-public monitoring," *Theoretical Economics*, 1(3), 311–340.

Mailath, G. J., and W. Olszewski (2011): "Folk theorems with bounded recall under (almost) perfect monitoring," *Games and Economic Behavior*, 71(1), 174–192.

Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations*, vol. 2. Oxford University Press.

Matsushima, H. (1991): "On the theory of repeated games with private information: Part I: Anti-folk theorem without communication," *Economics Letters*, 35(3), 253–256.

——— (2004): "Repeated games with private monitoring: Two players," *Econometrica*, 72(3), 823–852.

Matsushima, H., T. Tanaka, and T. Toyama (2013): "Behavioral Approach to Repeated Games with Private Monitoring," CIRJE-F-879, University of Tokyo.

Maynard-Smith, J. (1982): *Evolution and the Theory of Games*. Cambridge University Press.

Maynard-Smith, J., and G. R. Price (1973): "The logic of animal conflict," *Nature*, 246, 15.

Miyagawa, E., Y. Miyahara, and T. Sekiguchi (2008): "The folk theorem for repeated games with observation costs," *Journal of Economic Theory*, 139(1), 192–221.

Obara, I. (2009): "Folk theorem with communication," *Journal of Economic Theory*, 144(1), 120–134.

Piccione, M. (2002): "The repeated prisoner's dilemma with imperfect private monitoring," *Journal of Economic Theory*, 102(1), 70–83.

SEKIGUCHI, T. (1997): "Efficiency in repeated prisoner's dilemma with private monitoring," *Journal of Economic Theory*, 76(2), 345–361.

SUGAYA, T. (2015): "Folk Theorem in Repeated Games with Private Monitoring," Unpublished.

SUGAYA, T., AND S. TAKAHASHI (2013): "Coordination failure in repeated games with private monitoring," *Journal of Economic Theory*, 148(5), 1891–1928.

SWINKELS, J. M. (1992): "Evolutionary stability with equilibrium entrants," *Journal of Economic Theory*, 57(2), 306–332.

TAKAHASHI, S. (2010): "Community enforcement when players observe partners' past play," *Journal of Economic Theory*, 145(1), 42–62.

TAYLOR, P. D. (1979): "Evolutionarily stable strategies with two types of player," *Journal of applied probability*, pp. 76–83.

THOMAS, B. (1985): "On evolutionarily stable sets," *Journal of Mathematical. Biology*, 22(1), 105–115.

VAN VEELEN, M. (2012): "Robustness against indirect invasions," *Games and Economic Behavior*, 74(1), 382–393.

WEIBULL, J. W. (1995): *Evolutionary Game Theory*. The MIT Press.

YAMAMOTO, Y. (2007): "Efficiency results in N player games with imperfect private monitoring," *Journal of Economic Theory*, 135(1), 382–413.

——— (2009): "A limit characterization of belief-free equilibrium payoffs in repeated games," *Journal of Economic Theory*, 144(2), 802–824.

——— (2012): "Characterizing belief-free review-strategy equilibrium payoffs under conditional independence," *Journal of Economic Theory*, 147(5), 1998–2027.

——— (2014): "Individual learning and cooperation in noisy repeated games," *The Review of Economic Studies*, 81(1), 473–500.